

手册页第 1M 部分：系统管理命令

版权所有 © 2011, Oracle 和/或其附属公司。保留所有权利。

This software and related documentation are provided under a license agreement containing restrictions on use and disclosure and are protected by intellectual property laws. Except as expressly permitted in your license agreement or allowed by law, you may not use, copy, reproduce, translate, broadcast, modify, license, transmit, distribute, exhibit, perform, publish or display any part, in any form, or by any means. Reverse engineering, disassembly, or decompilation of this software, unless required by law for interoperability, is prohibited.

The information contained herein is subject to change without notice and is not warranted to be error-free. If you find any errors, please report them to us in writing.

If this is software or related documentation that is delivered to the U.S. Government or anyone licensing it on behalf of the U.S. Government, the following notice is applicable:

U.S. GOVERNMENT RIGHTS

Programs, software, databases, and related documentation and technical data delivered to U.S. Government customers are "commercial computer software" or "commercial technical data" pursuant to the applicable Federal Acquisition Regulation and agency-specific supplemental regulations. As such, the use, duplication, disclosure, modification, and adaptation shall be subject to the restrictions and license terms set forth in the applicable Government contract, and, to the extent applicable by the terms of the Government contract, the additional rights set forth in FAR 52.227-19, Commercial Computer Software License (December 2007). Oracle America, Inc., 500 Oracle Parkway, Redwood City, CA 94065.

This software or hardware is developed for general use in a variety of information management applications. It is not developed or intended for use in any inherently dangerous applications, including applications that may create a risk of personal injury. If you use this software or hardware in dangerous applications, then you shall be responsible to take all appropriate fail-safe, backup, redundancy, and other measures to ensure its safe use. Oracle Corporation and its affiliates disclaim any liability for any damages caused by use of this software or hardware in dangerous applications.

Oracle and Java are registered trademarks of Oracle and/or its affiliates. Other names may be trademarks of their respective owners.

Intel and Intel Xeon are trademarks or registered trademarks of Intel Corporation. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. AMD, Opteron, the AMD logo, and the AMD Opteron logo are trademarks or registered trademarks of Advanced Micro Devices. UNIX is a registered trademark of The Open Group.

This software or hardware and documentation may provide access to or information on content, products, and services from third parties. Oracle Corporation and its affiliates are not responsible for and expressly disclaim all warranties of any kind with respect to third-party content, products, and services. Oracle Corporation and its affiliates will not be responsible for any loss, costs, or damages incurred due to your access to or use of third-party content, products, or services.

目录

前言	23
简介	27
Intro(1M)	28
 系统管理命令—第 1 单元	 31
6to4relay(1M)	32
acct(1M)	35
acctadm(1M)	38
acctcms(1M)	42
acctcon(1M)	44
acctmerg(1M)	46
acctprc(1M)	47
acctsh(1M)	49
acpihpd(1M)	52
adbgen(1M)	53
add_allocatable(1M)	56
addbadsec(1M)	58
add_drv(1M)	60
arp(1M)	67
asradm(1M)	70
asr-notify(1M)	73
atohexlabel(1M)	76
audit(1M)	77
auditconfig(1M)	78
audيتد(1M)	88
auditrecord(1M)	90

auditreduce(1M)	94
auditstat(1M)	103
audit_warn(1M)	105
automount(1M)	108
automountd(1M)	117
autopush(1M)	120
bart(1M)	122
beadm(1M)	128
boot(1M)	135
bootadm(1M)	152
bootconfchk(1M)	155
busstat(1M)	156
captoinfo(1M)	161
catman(1M)	162
cfgadm(1M)	165
cfgadm_ac(1M)	175
cfgadm_cardbus(1M)	179
cfgadm_fp(1M)	180
cfgadm_ib(1M)	187
cfgadm_pci(1M)	196
cfgadm_sata(1M)	201
cfgadm_sbd(1M)	209
cfgadm_scsi(1M)	222
cfgadm_sdcard(1M)	229
cfgadm_shp(1M)	233
cfgadm_sysctrl(1M)	242
cfgadm_usb(1M)	247
chat(1M)	258
check-hostname(1M)	266
check-permissions(1M)	267
chk_encodings(1M)	268
chroot(1M)	270
cimworkshop(1M)	271
clear_locks(1M)	274
clinfo(1M)	275
clri(1M)	276

consadm(1m)	277
console-reset(1M)	279
conv_lp(1M)	280
conv_lpd(1M)	281
coreadm(1M)	282
cpustat(1M)	288
croinfo(1M)	292
cron(1M)	300
cryptoadm(1M)	302
cvcd(1M)	309
datadm(1M)	311
dcs(1M)	314
dd(1M)	316
ddu(1M)	322
ddu-text(1M)	324
devchassisd(1M)	325
devfsadm(1M)	327
device_allocate(1M)	330
device_remap(1M)	331
devinfo(1M)	333
devlinks(1M)	334
devnm(1M)	338
devprop(1M)	339
df(1M)	341
dfmounts(1M)	346
dfmounts_nfs(1M)	348
dfshares(1M)	349
dfshares_nfs(1M)	350
df_ufs(1M)	352
dhcpageant(1M)	353
dhcpconfig(1M)	362
dhcpmgr(1M)	369
dhtadm(1M)	371
dig(1M)	377
directoryserver(1M)	384
disks(1M)	405

diskscan(1M)	409
dispadm(1M)	410
dladm(1M)	413
dlnmgt(1M)	463
dlstat(1M)	464
dmesg(1M)	471
dminfo(1M)	472
dns-sd(1M)	474
dnssec-dsfromkey(1M)	477
dnssec-keyfromlabel(1M)	479
dnssec-keygen(1M)	481
dnssec-makekeyset(1M)	484
dnssec-signkey(1M)	486
dnssec-signzone(1M)	488
domainname(1M)	492
drd(1M)	494
drvconfig(1M)	495
dsbitmap(1M)	497
dscfg(1M)	499
dscfgadm(1M)	502
dscfglockd(1M)	504
dsstat(1M)	505
dsvclockd(1M)	512
dtrace(1M)	513
dumpadm(1M)	520
editmap(1M)	525
edquota(1M)	527
eeprom(1M)	529
efdaemon(1M)	539
embedded_su(1M)	540
etrn(1M)	544
fbconfig(1M)	546
fbconf_xorg(1M)	549
fcinfo(1M)	557
fdetach(1M)	576
fdisk(1M)	577

ff(1M)	584
ff_ufs(1M)	586
fiocompress(1M)	587
flowadm(1M)	588
flowstat(1M)	595
fmadm(1M)	600
fmd(1M)	606
fmdump(1M)	609
fmstat(1M)	616
fmthard(1M)	619
format(1M)	621
fpsd(1M)	626
fruadm(1M)	627
fsck(1M)	629
fsck_pcfs(1M)	633
fsck_udfs(1M)	635
fsck_ufs(1M)	638
fsdb(1M)	641
fsdb_udfs(1M)	642
fsdb_ufs(1M)	650
fsirand(1M)	660
fssnap(1M)	661
fssnap_ufs(1M)	663
fsstat(1M)	669
fstyp(1M)	676
fuser(1M)	678
fwflash(1M)	680
fwtmp(1M)	685
getdevpolicy(1M)	686
getent(1M)	687
gettable(1M)	690
getty(1M)	691
gkadmin(1M)	693
groupadd(1M)	695
groupdel(1M)	697
groupmod(1M)	698

growfs(1M)	700
gsscred(1M)	703
gssd(1M)	705
hald(1M)	706
hal-device(1M)	707
hal-fdi-validate(1M)	708
hal-find(1M)	709
hal-get-property(1M)	710
halt(1M)	712
hextoalabel(1M)	713
host(1M)	714
hostconfig(1M)	717
hotplug(1M)	719
hotplugd(1M)	729
htable(1M)	731
ickey(1M)	732
id(1M)	733
idmap(1M)	736
idmapd(1M)	748
idsconfig(1M)	749
ifconfig(1M)	751
if_mpadm(1M)	773
ifparse(1M)	775
iiadm(1M)	777
iicpbmp(1M)	786
iicpsd(1M)	787
ikeadm(1M)	788
ikecert(1M)	796
ilbadm(1M)	808
inetadm(1M)	823
ilomconfig(1M)	825
imqadmin(1M)	827
imqbrokerd(1M)	828
imqcmd(1M)	833
imqdbmgr(1M)	846
imqkeytool(1M)	849

imqobjmgr(1M)	851
imusermgr(1M)	860
in.chargend(1M)	863
in.comsat(1M)	864
in.daytimed(1M)	865
in.dhcpd(1M)	866
in.discardd(1M)	872
in.echod(1M)	873
inetadm(1M)	874
inetconv(1M)	878
inetd(1M)	880
in.fingerd(1M)	888
infocmp(1M)	890
in.iked(1M)	893
init(1M)	898
init.sma(1M)	903
init.wbem(1M)	904
inityp2l(1M)	906
in.lpd(1M)	908
in.mpathd(1M)	910
in.ndpd(1M)	915
in.rarpd(1M)	918
in.rdisc(1M)	920
in.rexecd(1M)	922
in.ripngd(1M)	924
in.rlogind(1M)	927
in.routed(1M)	931
in.rshd(1M)	937
in.rwhod(1M)	941
install(1M)	943
installboot(1M)	945
installf(1M)	947
installgrub(1M)	951
in.stddiscover(1M)	952
in.stlisten(1M)	953
in.talkd(1M)	954

in.telnetd(1M)	955
in.tftpd(1M)	960
in.timed(1M)	962
intrd(1M)	963
intrstat(1M)	964
in.uucpd(1M)	967
iostat(1M)	969
ipaddrsel(1M)	975
ipadm(1M)	979
ipf(1M)	1007
ipfs(1M)	1011
ipfstat(1M)	1013
ipmgmt(1M)	1016
ipmon(1M)	1017
ipmpstat(1M)	1020
ipnat(1M)	1027
ippool(1M)	1029
ipqosconf(1M)	1032
ipsecacls(1M)	1043
ipseccomp(1M)	1050
ipseckey(1M)	1074
iscsiadm(1M)	1089
isns(1M)	1101
isnsadm(1M)	1104
itadm(1M)	1111
itu(1M)	1119
k5srvutil(1M)	1122
kadb(1M)	1124
kadmin(1M)	1126
kadmind(1M)	1141
kcfd(1M)	1145
kclient(1M)	1146
kdb5_ldap_util(1M)	1152
kdb5_util(1M)	1163
kdcmgr(1M)	1169
kernel(1M)	1173

keyserv(1M)	1177
killall(1M)	1179
kmscfg(1M)	1180
kprop(1M)	1182
kpropd(1M)	1184
kproplog(1M)	1186
krb5kdc(1M)	1188
ksslcfg(1M)	1190
kstat(1M)	1195
ktkt_warnd(1M)	1199
labeld(1M)	1200
labelit(1M)	1201
labelit_hsf(1M)	1203
labelit_udfs(1M)	1204
labelit_ufs(1M)	1206
latencytop(1M)	1207
ldapaddent(1M)	1210
ldap_cachemgr(1M)	1216
ldapclient(1M)	1218
ldmad(1M)	1229
link(1M)	1230
llc2_loop(1M)	1232
lldpadm(1M)	1234
lldpd(1M)	1249
lms(1M)	1250
localectr(1M)	1251
locator(1M)	1255
lockd(1M)	1256
lockfs(1M)	1259
lockstat(1M)	1263
lofiadm(1M)	1273
logadm(1M)	1280
logins(1M)	1289
lpget(1M)	1291
lpset(1M)	1293
luxadm(1M)	1295

系统管理命令—第2单元	1307
mail.local(1M)	1308
makedbm(1M)	1310
makemap(1M)	1312
makeuuid(1M)	1314
masfcnv(1M)	1316
mdlogd(1M)	1322
mdmonitord(1M)	1324
mdnsd(1M)	1325
medstat(1M)	1327
metaclear(1M)	1329
metadb(1M)	1332
metadevadm(1M)	1338
metahs(1M)	1341
metaimport(1M)	1345
metainit(1M)	1347
metaoffline(1M)	1358
metaparam(1M)	1360
metarecover(1M)	1362
metarename(1M)	1365
metareplace(1M)	1368
metaset(1M)	1371
metassist(1M)	1380
metastat(1M)	1385
metasync(1M)	1390
metattach(1M)	1392
mib2mof(1M)	1397
mibiisa(1M)	1399
mkbootmedia(1M)	1423
mkdevalloc(1M)	1424
mkdevmaps(1M)	1425
mkfifo(1M)	1426
mkfile(1M)	1428
mkfs(1M)	1429
mkfs_pcfs(1M)	1431
mkfs_udfs(1M)	1435

mkfs_ufs(1M)	1437
mknod(1M)	1442
mkntfs(1M)	1443
mkpwdict(1M)	1446
modinfo(1M)	1447
modload(1M)	1449
modunload(1M)	1450
mofcomp(1M)	1451
mofreg(1M)	1454
monitor(1M)	1457
mount(1M)	1469
mountall(1M)	1474
mountd(1M)	1476
mount_hfs(1M)	1478
mount_nfs(1M)	1481
mount_pcfs(1M)	1491
mount_smbfs(1M)	1494
mount_tmpfs(1M)	1499
mount_udfs(1M)	1501
mount_ufs(1M)	1503
mpathadm(1M)	1507
mpstat(1M)	1514
msgid(1M)	1517
mvdir(1M)	1518
named(1M)	1519
named-checkconf(1M)	1525
named-checkzone(1M)	1526
ncaconfd(1M)	1529
ncheck(1M)	1530
ncheck_ufs(1M)	1532
ndd(1M)	1533
ndmpadm(1M)	1535
ndmpd(1M)	1539
ndmpstat(1M)	1540
netadm(1M)	1543
netcfg(1M)	1548

netcfgd(1M)	1562
netsservices(1M)	1563
netstat(1M)	1564
netstrategy(1M)	1573
newaliases(1M)	1574
newfs(1M)	1576
newkey(1M)	1581
nfs4cbd(1M)	1582
nfsd(1M)	1583
nfslogd(1M)	1587
nfsmapid(1M)	1590
nfsref(1M)	1593
nfsstat(1M)	1595
nscadm(1M)	1601
nscd(1M)	1602
nscfg(1M)	1604
nslookup(1M)	1608
nsupdate(1M)	1612
ntfscat(1M)	1616
ntfscclone(1M)	1618
ntfsccluster(1M)	1622
ntfscmp(1M)	1624
ntfscp(1M)	1625
ntfsfix(1M)	1627
ntfsinfo(1M)	1628
ntfslabel(1M)	1630
ntfsls(1M)	1632
ntfsprogs(1M)	1634
ntfsresize(1M)	1636
ntfsundelete(1M)	1639
nwamd(1M)	1643
obpsym(1M)	1644
oplhpd(1M)	1646
parted(1M)	1647
pbind(1M)	1651
pcitool(1M)	1654

pginfo(1M)	1657
pgstat(1M)	1661
picld(1M)	1667
ping(1M)	1669
pkg2du(1M)	1674
pkgadd(1M)	1676
pkgadm(1M)	1683
pkgask(1M)	1688
pkgchk(1M)	1690
pkgcond(1M)	1694
pkgrm(1M)	1696
plockstat(1M)	1699
pntadm(1M)	1701
polkit-is-privileged(1M)	1708
pooladm(1M)	1709
poolbind(1M)	1711
poolcfg(1M)	1713
poold(1M)	1717
poolstat(1M)	1719
ports(1M)	1723
poweradm(1M)	1727
powertop(1M)	1731
pppd(1M)	1733
pppoec(1M)	1757
pppoed(1M)	1760
pppstats(1M)	1765
praudit(1M)	1768
projadd(1M)	1770
projdel(1M)	1773
projmod(1M)	1775
prstat(1M)	1780
prtconf(1M)	1787
prtdiag(1M)	1789
prtdscp(1M)	1790
prtfpu(1M)	1792
prtpicl(1M)	1793

prvtoc(1M)	1794
psradm(1M)	1797
psrinfo(1M)	1800
psrset(1M)	1802
pwck(1M)	1807
pwconv(1M)	1808
quot(1M)	1810
quota(1M)	1812
quotacheck(1M)	1813
quotaon(1M)	1815
raidctl(1M)	1817
ramdiskadm(1M)	1826
rcapadm(1M)	1828
rcapd(1M)	1831
rctladm(1M)	1833
rdate(1M)	1835
reboot(1M)	1836
rem_drv(1M)	1840
remove_allocatable(1M)	1842
removef(1M)	1844
reparsed(1M)	1846
repquota(1M)	1847
rmmount(1M)	1848
rmt(1M)	1850
rmvolmgr(1M)	1852
rndc(1M)	1854
rndc-confgen(1M)	1856
roleadd(1M)	1858
roledel(1M)	1863
rolemod(1M)	1865
root_archive(1M)	1870
route(1M)	1871
routeadm(1M)	1879
rpcbind(1M)	1884
rpc.bootparamd(1M)	1887
rpcinfo(1M)	1888

rpc.mdcommd(1M)	1892
rpc.metad(1M)	1893
rpc.metamedd(1M)	1894
rpc.metamhd(1M)	1895
rpc.rexd(1M)	1896
rpc.rstatd(1M)	1898
rpc.rusersd(1M)	1899
rpc.rwalld(1M)	1900
rpc.smserverd(1M)	1901
rpc.sprayd(1M)	1902
rpc.yppasswdd(1M)	1903
rpc.yppupdated(1M)	1906
rquotad(1M)	1907
rsh(1M)	1908
rtc(1M)	1910
rtquery(1M)	1911
runacct(1M)	1913
rwall(1M)	1916
sar(1M)	1917
sasinfo(1M)	1919
savecore(1M)	1929
sbdadm(1M)	1931
scadm(1M)	1933
sckmd(1M)	1942
scmadm(1M)	1943
sconadm(1M)	1945
sdpadm(1M)	1950
sendmail(1M)	1951
sf880drd(1M)	1978
sftp-server(1M)	1979
shadowd(1M)	1981
shadowstat(1M)	1982
share(1M)	1983
shareall(1M)	1985
sharectl(1M)	1986
share_nfs(1M)	1989

share_smb(1M)	1996
showmount(1M)	2002
shutdown(1M)	2003
slpd(1M)	2005
smattrpop(1M)	2007
smbadm(1M)	2011
smbd(1M)	2018
smbiod(1M)	2019
smbios(1M)	2020
smbstat(1M)	2022
smrsh(1M)	2026
smtnrhdb(1M)	2027
smtnrhtp(1M)	2032
smtnzonecfg(1M)	2037
sasinfo(1M)	2042
sndradm(1M)	2044
sndrd(1M)	2052
sndrsyncd(1M)	2054
snmpdx(1M)	2055
snmp-notify(1M)	2058
snmpXwbemd(1M)	2060
snoop(1M)	2063
soconfig(1M)	2073
soladdapp(1M)	2076
soldelapp(1M)	2077
solstice(1M)	2078
sppptun(1M)	2079
spray(1M)	2082
srptadm(1M)	2083
sshd(1M)	2086
ssh-keysign(1M)	2102
statd(1M)	2104
stclient(1M)	2106
stmfadm(1M)	2111
stmsboot(1M)	2120
strace(1M)	2125

strclean(1M)	2127
strerr(1M)	2128
sttydefs(1M)	2130
su(1M)	2132
sulogin(1M)	2135
svadm(1M)	2136
svcadm(1M)	2138
svccfg(1M)	2144
svc.configd(1M)	2154
svc.ipfd(1M)	2155
svc.startd(1M)	2160
swap(1M)	2166
sync(1M)	2169
syncinit(1M)	2170
syncloop(1M)	2173
syncstat(1M)	2176
sysconfig(1M)	2179
sysdef(1M)	2183
syseventadm(1M)	2185
syseventconfd(1M)	2190
syseventd(1M)	2191
syslogd(1M)	2193
tapes(1M)	2197
tcxconfig(1M)	2201
th_define(1M)	2202
th_manage(1M)	2213
tic(1M)	2215
tncfg(1M)	2216
tnchkdb(1M)	2223
tnctl(1M)	2225
tnd(1M)	2227
tninfo(1M)	2229
tpmadm(1M)	2231
traceroute(1M)	2234
trapstat(1M)	2241
ttymon(1M)	2252

tunefs(1M)	2256
txzonemgr(1M)	2258
tzreload(1M)	2260
tzselect(1M)	2261
uadmin(1M)	2262
ucodeadm(1M)	2263
ufsdump(1M)	2265
ufsrestore(1M)	2272
unshare(1M)	2279
unshare_nfs(1M)	2280
 系统管理命令—第 3 单元	 2281
update_drv(1M)	2282
useradd(1M)	2286
userdel(1M)	2292
usermod(1M)	2294
utmpd(1M)	2300
uuccheck(1M)	2302
uucico(1M)	2303
uucleanup(1M)	2305
uusched(1M)	2307
Uutry(1M)	2308
uuxqt(1M)	2309
vbiosd(1M)	2310
vdiskadm(1M)	2311
virt-convert(1M)	2320
virtinfo(1M)	2323
vmstat(1M)	2326
vntsd(1M)	2330
volcopy(1M)	2334
volcopy_ufs(1M)	2336
vrrpadm(1M)	2337
vrrpd(1M)	2342
vscanadm(1M)	2343
vs cand(1M)	2348

vt daemon(1M)	2350
wall(1M)	2351
wanboot_keygen(1M)	2353
wanboot_keymgmt(1M)	2355
wanboot_p12split(1M)	2356
wanbootutil(1M)	2357
wbemadmin(1M)	2358
wbemconfig(1M)	2361
wbemlogviewer(1M)	2362
wcadmin(1M)	2364
whodo(1M)	2369
wpad(1M)	2371
wracct(1M)	2373
wusbadm(1M)	2375
ypbind(1M)	2381
ypinit(1M)	2384
ypmake(1M)	2386
ypmap2src(1M)	2388
yppoll(1M)	2390
yppush(1M)	2391
ypserv(1M)	2393
ypserv_resolv(1M)	2397
ypset(1M)	2398
ypstart(1M)	2400
ypxfr(1M)	2401
zdb(1M)	2403
zdump(1M)	2404
zfs(1M)	2405
zfs_allow(1M)	2438
zfs_encrypt(1M)	2444
zfs_share(1M)	2452
zic(1M)	2458
zoneadm(1M)	2463
zoneadmd(1M)	2470
zonecfg(1M)	2471
zonep2vchk(1M)	2495

zonestatd(1M) 2508

zpool(1M) 2509

zstreamdump(1M) 2531

前言

无论是初次使用 SunOS 操作系统的用户还是熟悉该操作系统的用户，均可通过联机手册页获取与系统及其功能有关的信息。手册页用于简要回答关于用途/用法的问题。手册页通常是参考手册的组成部分，并不用作教程。

概述

以下内容包含对手册页各部分及其所引用信息的简要说明：

- 第 1 部分按照字母顺序介绍了操作系统中提供的各种命令。
- 第 1M 部分按照字母顺序介绍了主要用于系统维护和管理的各种命令。
- 第 2 部分介绍了所有的系统调用。其中的大多数调用可能返回一个或多个错误。如果返回一个异常值，则表明有错误情况。
- 第 3 部分介绍了各种库中包含的函数，这些函数不属于第 2 部分介绍的那些直接调用 UNIX 系统原语 (primitive) 的函数。
- 第 4 部分简要介绍了各种文件的格式，并在适当之处给出了文件格式的 C 结构声明。
- 第 5 部分包含其他文档，如字符集表。
- 第 7 部分介绍了涉及特定硬件外围设备和设备驱动程序的各种特殊文件。还介绍了 STREAMS 软件驱动程序、模块和 STREAMS 通用的一组系统调用。
- 第 9E 部分介绍了 DDI (Device Driver Interface, 设备驱动程序接口) /DKI (Driver/Kernel Interface, 驱动程序/内核接口)、仅 DDI 和仅 DKI 入口点例程，开发者可以将这些例程包含在设备驱动程序中。
- 第 9F 部分介绍了可供设备驱动程序使用的内核函数。
- 第 9S 部分介绍了驱动程序用来在驱动程序和内核之间共享信息的数据结构。

下面是手册页的通用格式。每个手册的手册页部分通常遵循该顺序，但只包括需要的标题。例如，如果未报告任何已知问题，则不包括“已知问题”部分。有关每一部分的更多详细信息，请参见 [intro](#) 页；有关手册页的更多一般信息，请参见 [man\(1\)](#)。

名称 本部分提供了记录的命令或函数的名称，后跟其用途的简要说明。

用法概要	本部分说明了命令或函数的语法。如果命令或文件不存在于标准路径中，则显示其全路径名。除非要求使用不同的参数顺序，否则选项和参数均按字母顺序排列，首先是单个字母的参数，接下来是带有参数的选项。 本部分使用以下特殊字符： [] 方括号。括在方括号中的选项或参数是可选的。如果没有方括号，则必须指定参数。 ... 省略号。可以为该符号前面的参数提供多个值，或者可以多次指定该参数，例如“filename...”。 - 分隔符。一次只能指定一个由该字符分隔的参数。 { } 大括号。括在大括号内的选项和/或参数是相互依赖的，因此必须将大括号中的所有内容视为一个单元。
协议	本部分仅出现在第 3R 子部分，用于指示协议说明文件。
描述	本部分定义了服务的功能和行为。因此，它简明地介绍了命令执行哪些操作。它不讨论“选项”或引用“示例”。在“用法”下介绍了交互式命令、子命令、请求、宏和函数。
IOctl	本部分仅出现在第 7 部分的手册页中。只有为 <code>ioctl(2)</code> 系统调用提供了适当参数的设备类才被称为 <code>ioctl</code> ，并生成自己的标题。特定设备的 <code>ioctl</code> 调用按字母顺序显示在该特定设备的手册页中。 <code>ioctl</code> 调用用于特殊类别的设备。所有这些调用都以 <code>io</code> 结尾，例如 <code>mtio(7I)</code> 。
选项	本部分列出了各命令选项及每个选项用途的简明摘要。逐个列出各个选项，并以它们在“用法概要”部分显示的顺序排列。在选项下讨论各个选项可能的参数，还提供缺省值（如果适用）。
操作数	本部分列出了命令操作数，并介绍它们对命令操作的影响。
输出	本部分介绍了命令所生成的输出（标准输出、标准错误或输出文件）。
返回值	如果手册页记录返回值的函数，则本部分列出这些值并介绍返回这些值应满足的条件。如果函数只能返回常量值（例如 0 或 -1），则将在标记的段落中列出这些值。否则，会有单个段落介绍每个函数的返回值。声明为 <code>void</code> 的函数不返回值，因此不会在“返回值”中讨论这些函数。
错误	对于故障，大多数函数将指出它们出现故障的原因的错误代码置于全局变量 <code>errno</code> 中。本部分按字母顺序列出了函数可以生成的所有错误代码，并介绍了导致每个错误的条件。如果多个条件可以导致同一错误，则在错误代码下以单独的段落介绍每个条件。
用法	本部分列出了需要详细说明的特殊规则、功能和命令。此处列出的子部分用于说明内置功能：
	命令

	修饰符 变量 表达式 输入语法
示例	本部分提供了用法的示例，或者如何使用命令或函数的示例。会尽可能显示包括命令行输入和计算机响应的完整示例。只要给出了示例，就会显示 example% 提示，如果用户必须为超级用户，则提示显示为 example# 。示例后面跟有说明、变量替换规则或返回值。大部分示例说明了“用法概要”、“说明”、“选项”和“用法”部分的概念。
环境变量	本部分列出了命令或函数影响的所有环境变量，其后附加了关于影响的简要说明。
退出状态	本部分列出了命令返回到调用程序或 shell 中的值以及导致返回这些值的条件。通常，返回零表示成功完成，返回非零值表示各种错误条件。
文件	本部分列出了手册页引用的所有文件名称、相关文件以及命令创建或所需的文件。每个文件名称后面都具有描述性摘要或说明。
属性	本部分通过定义属性类型及其相应的值列出了命令、实用程序和设备驱动程序的特征。有关详细信息，请参见 attributes(5) 。
另请参见	本部分列出了对其他手册页、内部文档和外部出版物的引用。
诊断	本部分列出了诊断消息以及导致错误的条件的简要说明。
警告	本部分列出了有关特殊条件的警告，这些条件可能会严重影响您的工作状况。此部分不是诊断列表。
附注	本部分列出了不属于页面任何部分的其他信息。它采用对用户旁白提示的形式，包含用户特别关注的要点。此处不包含关键信息。
已知问题	本部分介绍了已知问题，并尽可能给出解决方法。

参 考 文 档

简介

引用名	Intro – introduction to maintenance commands and application programs												
描述	This section describes, in alphabetical order, commands that are used chiefly for system maintenance and administration purposes. Because of command restructuring for the Virtual File System architecture, there are several instances of multiple manual pages that begin with the same name. For example, the mount, pages – mount(1M), mount_hfs(1M), mount_nfs(1M), mount_tmpfs(1M), and mount_ufs(1M). In each such case the first of the multiple pages describes the syntax and options of the generic command, that is, those options applicable to all FSTypes (file system types). The succeeding pages describe the functionality of the FSType-specific modules of the command. These pages list the command followed by an underscore (_) and the FSType to which they pertain. Note that the administrator should not attempt to call these modules directly. The generic command provides a common interface to all of them. Thus the FSType-specific manual pages should not be viewed as describing distinct commands, but rather as detailing those aspects of a command that are specific to a particular FSType.												
Command Syntax	Unless otherwise noted, commands described in this section accept options and other arguments according to the following syntax: <i>name</i> [<i>option</i>(<i>s</i>)] [<i>cmdarg</i>(<i>s</i>)] where: <table><tr><td><i>name</i></td><td>The name of an executable file.</td></tr><tr><td><i>option</i></td><td>– <i>noargletter</i>(<i>s</i>) or, – <i>argletter</i>< ><i>optarg</i> where < > is optional white space.</td></tr><tr><td><i>noargletter</i></td><td>A single letter representing an option without an argument.</td></tr><tr><td><i>argletter</i></td><td>A single letter representing an option requiring an argument.</td></tr><tr><td><i>optarg</i></td><td>Argument (character string) satisfying preceding <i>argletter</i>.</td></tr><tr><td><i>cmdarg</i></td><td>Pathname (or other command argument) <i>not</i> beginning with – or, – by itself indicating the standard input.</td></tr></table>	<i>name</i>	The name of an executable file.	<i>option</i>	– <i>noargletter</i> (<i>s</i>) or, – <i>argletter</i> < > <i>optarg</i> where < > is optional white space.	<i>noargletter</i>	A single letter representing an option without an argument.	<i>argletter</i>	A single letter representing an option requiring an argument.	<i>optarg</i>	Argument (character string) satisfying preceding <i>argletter</i> .	<i>cmdarg</i>	Pathname (or other command argument) <i>not</i> beginning with – or, – by itself indicating the standard input.
<i>name</i>	The name of an executable file.												
<i>option</i>	– <i>noargletter</i> (<i>s</i>) or, – <i>argletter</i> < > <i>optarg</i> where < > is optional white space.												
<i>noargletter</i>	A single letter representing an option without an argument.												
<i>argletter</i>	A single letter representing an option requiring an argument.												
<i>optarg</i>	Argument (character string) satisfying preceding <i>argletter</i> .												
<i>cmdarg</i>	Pathname (or other command argument) <i>not</i> beginning with – or, – by itself indicating the standard input.												
属性	See attributes(5) for a discussion of the attributes listed in this section.												
Acknowledgments	Oracle America, Inc. gratefully acknowledges The Open Group for permission to reproduce portions of its copyrighted documentation. Original documentation from The Open Group can be obtained online at http://www.opengroup.org/bookstore/. The Institute of Electrical and Electronics Engineers and The Open Group, have given us permission to reprint portions of their documentation.												

In the following statement, the phrase "this text" refers to portions of the system documentation.

Portions of this text are reprinted and reproduced in electronic form in the SunOS Reference Manual, from IEEE Std 1003.1, 2004 Edition, Standard for Information Technology -- Portable Operating System Interface (POSIX), The Open Group Base Specifications Issue 6, Copyright (C) 2001-2004 by the Institute of Electrical and Electronics Engineers, Inc and The Open Group. In the event of any discrepancy between these versions and the original IEEE and The Open Group Standard, the original IEEE and The Open Group Standard is the referee document. The original Standard can be obtained online at <http://www.opengroup.org/unix/online.html>.

This notice shall appear on any product containing this material.

另请参见

[getopt\(1\)](#), [getopt\(3C\)](#), [attributes\(5\)](#)

诊断

Upon termination, each command returns 0 for normal termination and non-zero to indicate troubles such as erroneous parameters, bad or inaccessible data, or other inability to cope with the task at hand. It is called variously "exit code," "exit status," or "return code," and is described only where special conventions are involved.

附注

Unfortunately, not all commands adhere to the standard syntax.

参 考 文 档

系统管理命令—第 1 单元

示例 1 Printing the In-Kernel Configuration Status (续)

If 6to4 relay router communication is disabled, the administrator will see the following message:

6to4relay: 6to4 Relay Router communication support is disabled.

If 6to4 router communication is enabled, the user will see this message:

6to4relay: 6to4 Relay Router communication support is enabled.
IPv4 destination address of Relay Router = 192.88.99.1

退出状态

The following exit values are returned:

0 Successful completion.

>0 An error occurred.

文件

/usr/sbin/6to4relay The default installation root

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[ifconfig\(1M\)](#), [attributes\(5\)](#)

Huitema, C. RFC 3068, An Anycast Prefix for 6to4 Relay Routers. Network Working Group. June, 2001.

Carpenter, B. and Moore, K. RFC 3056, Connection of IPv6 Domains via IPv4 Clouds. Network Working Group. February, 2001.

诊断

The 6to4relay reports the following messages:

6to4relay: input (0.0.0.0) is not a valid IPv4 unicast address

示例: The following example provides an incorrect unicast address.

example# **6to4relay -e -a 0.0.0.0**

描述: The address specified with the -a option must be a valid unicast address.

6to4relay: option requires an argument -a

usage:

6to4relay

6to4relay -e [-a <addr>]

```
6to4relay -d
```

```
6to4relay -h
```

示例: The following example does not include an argument for the `-a` option.

```
example# 6to4relay -e -a
```

描述: The `-a` option requires an argument.

usage:

```
6to4relay
```

```
6to4relay -e [-a <addr>]
```

```
6to4relay -d
```

```
6to4relay -h
```

示例: The following example specifies options that are not permitted.

```
example# 6to4relay -e -d
```

描述: The options specified are not permitted. A usage message is output to the screen.

usage:

```
6to4relay
```

```
6to4relay -e [-a <addr>]
```

```
6to4relay -d
```

```
6to4relay -h
```

示例: The following example specifies the `-a` option without specifying the `-e` option.

```
example# 6to4relay -a 1.2.3.4
```

描述: The `-e` option is required in conjunction with the `-a` option. A usage message is output to the screen.

```
6to4relay: ioctl (I_STR) : Invalid argument
```

示例: The following example specifies an invalid address.

```
example# 6to4relay -e -a 239.255.255.255
```

描述: The address specified with the `-a` option must not be a class *d* *addr*.

引用名	acct, acctdisk, acctdusg, accton, acctwtmp, closewtmp, utmp2wtmp – overview of accounting and miscellaneous accounting commands
用法概要	<pre> /usr/lib/acct/acctdisk /usr/lib/acct/acctdusg [-u filename] [-p filename] /usr/lib/acct/accton [filename] /usr/lib/acct/acctwtmp reason filename /usr/lib/acct/closewtmp /usr/lib/acct/utmp2wtmp </pre>
描述	Accounting software is structured as a set of tools (consisting of both C programs and shell procedures) that can be used to build accounting systems. acctsh(1M) describes the set of shell procedures built on top of the C programs. Connect time accounting is handled by various programs that write records into <code>/var/adm/wtmpx</code>, as described in utmpx(4). The programs described in acctcon(1M) convert this file into session and charging records, which are then summarized by acctmerg(1M). Process accounting is performed by the system kernel. Upon termination of a process, one record per process is written to a file (normally <code>/var/adm/pacct</code>). The programs in acctprc(1M) summarize this data for charging purposes; acctcms(1M) is used to summarize command usage. Current process data may be examined using acctcom(1). Process accounting records and connect time accounting records (or any accounting records in the <code>tacct</code> format described in acct.h(3HEAD)) can be merged and summarized into total accounting records by <code>acctmerg</code> (see <code>tacct</code> format in acct.h(3HEAD)). <code>prtacct</code> (see acctsh(1M)) is used to format any or all accounting records. <code>acctdisk</code> reads lines that contain user ID, login name, and number of disk blocks and converts them to total accounting records that can be merged with other accounting records. <code>acctdisk</code> returns an error if the input file is corrupt or improperly formatted. <code>acctdusg</code> reads its standard input (usually from <code>find / -print</code>) and computes disk resource consumption (including indirect blocks) by login. <code>accton</code> without arguments turns process accounting off. If <i>filename</i> is given, it must be the name of an existing file, to which the kernel appends process accounting records (see acct(2) and acct.h(3HEAD)). <code>acctwtmp</code> writes a utmpx(4) record to <i>filename</i>. The record contains the current time and a string of characters that describe the <i>reason</i>. A record type of <code>ACCOUNTING</code> is assigned (see utmpx(4)) <i>reason</i> must be a string of 11 or fewer characters, numbers, \$, or spaces. For example, the following are suggestions for use in reboot and shutdown procedures, respectively: <pre> acctwtmp "acctg on" /var/adm/wtmpx acctwtmp "acctg off" /var/adm/wtmpx </pre>

For each user currently logged on, `closewtmp` puts a false `DEAD_PROCESS` record in the `/var/adm/wtmpx` file. `runacct` (see [runacct\(1M\)](#)) uses this false `DEAD_PROCESS` record so that the connect accounting procedures can track the time used by users logged on before `runacct` was invoked.

For each user currently logged on, `runacct` uses `utmp2wtmp` to create an entry in the file `/var/adm/wtmpx`, created by `runacct`. Entries in `/var/adm/wtmpx` enable subsequent invocations of `runacct` to account for connect times of users currently logged in.

选项

The following options are supported:

- u *filename* Places in *filename* records consisting of those filenames for which `acctdusg` charges no one (a potential source for finding users trying to avoid disk charges).
- p *filename* Specifies a password file, *filename*. This option is not needed if the password file is `/etc/passwd`.

环境变量

If any of the `LC_*` variables (`LC_TYPE`, `LC_MESSAGES`, `LC_TIME`, `LC_COLLATE`, `LC_NUMERIC`, and `LC_MONETARY`) (see [environ\(5\)](#)) are not set in the environment, the operational behavior of `acct` for each corresponding locale category is determined by the value of the `LANG` environment variable. If `LC_ALL` is set, its contents are used to override both the `LANG` and the other `LC_*` variables. If none of the above variables are set in the environment, the "C" (U.S. style) locale determines how `acct` behaves.

`LC_CTYPE` Determines how `acct` handles characters. When `LC_CTYPE` is set to a valid value, `acct` can display and handle text and filenames containing valid characters for that locale. `acct` can display and handle Extended Unix Code (EUC) characters where any character can be 1, 2, or 3 bytes wide. `acct` can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.

`LC_TIME` Determines how `acct` handles date and time formats. In the "C" locale, date and time handling follows the U.S. rules.

文件

- `/etc/passwd` Used for login name to user ID conversions.
- `/usr/lib/acct` Holds all accounting commands listed in sub-class 1M of this manual.
- `/var/adm/pacct` Current process accounting file.
- `/var/adm/wtmpx` History of user access and administration information..

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/accounting/legacy-accounting

另请参见

acctcom(1), acctcms(1M), acctcon(1M), acctmerg(1M), acctprc(1M), acctsh(1M),
fwtmp(1M), runacct(1M), acct(2), acct.h(3HEAD), passwd(4), utmpx(4), attributes(5),
environ(5)

《Oracle Solaris 管理：常见任务》

引用名	acctadm – configure extended accounting facility
用法概要	<pre>/usr/sbin/acctadm [-DEsrux] [-d <i>resource_list</i>] [-e <i>resource_list</i>] [-f <i>filename</i>] [task process flow net]</pre>
描述	acctadm configures various attributes of the extended accounting facility. Without arguments, acctadm displays the current status of the extended accounting facility.
选项	The following options are supported: -d <i>resource_list</i> Disable reporting of resource usage for resource. Specify <i>resource_list</i> as a comma-separated list of resources or resource groups. This option requires an operand. See OPERANDS. -D Disable accounting of the given operand type without closing the accounting file. This option can be used to temporarily stop writing accounting records to the accounting file without closing it. To close the file use the -x option. See -x. -e <i>resource_list</i> Enable reporting of resource usage for resource. Specify <i>resource_list</i> as a comma-separated list of resources or resource groups. This option requires an operand. See OPERANDS. -E Enable accounting of the given operand type without sending the accounting output to a file. This option requires an operand. See OPERANDS. -f <i>filename</i> Send the accounting output for the given operand type to <i>filename</i>. If <i>filename</i> exists, its contents must be of the given accounting type. This option requires an operand. See OPERANDS. -r Display available resource groups. When this option is used with an operand, it displays resource groups available for a given accounting type. When no operand is specified, this option displays resource groups for all available accounting types. See OPERANDS. -s Start method for the smf(5) instance. This option is used to restore the extended accounting configuration at boot.

-x

Deactivate accounting of the given operand type. This option also closes the accounting file for the given accounting type if it is currently open.

This option requires an operand. See OPERANDS.

操作数

The -d, -D, -e, -E, -f, and -x options require an operand.

The following operands are supported:

process

Run acctadm on the process accounting components of the extended accounting facility.

task

Run acctadm on the task accounting components of the extended accounting facility.

flow

Run acctadm on the IPQoS accounting components of the extended accounting facility.

net

Run acctadm on links and flows administered by [dladm\(1M\)](#) and [flowadm\(1M\)](#), respectively. Basic network accounting relates only to links, while extended network accounting includes both link and flow accounting.

The optional final parameter to acctadm represents whether the command should act on the process, system task, IPQoS, or network accounting components of the extended accounting facility.

示例

示例 1 Displaying the Current Status

The following command displays the current status. In this example, system task accounting is active and tracking only CPU resources. Process and flow accounting are not active.

```
$ acctadm
    Task accounting: active
    Task accounting file: /var/adm/exacct/task
    Tracked task resources: extended
    Untracked task resources: host
    Process accounting: inactive
    Process accounting file: none
    Tracked process resources: none
    Untracked process resources: extended,host
    Flow accounting: inactive
    Flow accounting file: none
    Tracked flow resources: none
    Untracked flow resources: extended
    Net accounting: inactive
    Net accounting file: none
    Tracked Net resources: none
    Untracked Net resources: extended
```

示例 2 Activating Basic Process Accounting

The following command activates basic process accounting:

```
$ acctadm -e basic -f /var/adm/exacct/proc process
```

示例 3 Displaying Available Resource Groups

The following command displays available resource groups:

```
$ acctadm -r
process:
extended pid,uid,gid,cpu,time,command,tty,projid, \
taskid,ancpid,wait-status,zone,flag,memory,mstate
basic    pid,uid,gid,cpu,time,command,tty,flag
task:
extended taskid,projid,cpu,time,host,mstate,anctaskid,zone
basic    taskid,projid,cpu,time
flow:
extended saddr,daddr,sport,dport,proto,dsfield,nbytes,npkts, \
action,ctime,lseen,projid,uid
basic    saddr,daddr,sport,dport,proto,nbytes,npkts,action
net:
extended name,devname,edest,vlan_tpid,vlan_tci,sap,cpuid, \
priority,bwlimit,curtime,ibytes,obytes,ipkts,opks,ierrpks \
oerrpks,saddr,daddr,sport,dport,protocol,dsfield
basic    name,devname,edest,vlan_tpid,vlan_tci,sap,cpuid, \
priority,bwlimit,curtime,ibytes,obytes,ipkts,opks,ierrpks \
oerrpks
```

In the output above, the lines beginning with `extended` are shown with a backslash character. In actual `acctadm` output, these lines are displayed as unbroken, long lines.

示例 4 Displaying Resource Groups for Task Accounting

The following command displays resource groups for task accounting:

```
$ acctadm -r task
extended taskid,projid,cpu,time,host,mstate,anctaskid,zone
basic    taskid,projid,cpu,time
```

退出状态

The following exit values are returned:

0

Successful completion.

The modifications to the current configuration were valid and made successfully.

1

An error occurred.

A fatal error occurred either in obtaining or modifying the accounting configuration.

- 2
Invalid command line options were specified.
- 95
A fatal, non-configuration error occurred during the start of the `smf(5)` service instance.
- 96
A fatal configuration error occurred during the start of the `smf(5)` service instance.

属性 See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 `dladm(1M)`, `flowadm(1M)`, `acct(2)`, `attributes(5)`, `smf(5)`, `ipqos(7ipp)`

附注 Both extended accounting and regular accounting can be active.

Available resources can vary from system to system, and from platform to platform.

Extended accounting configuration is stored in the service management facility (`smf(5)`) repository. The configuration is restored at boot by a transient service instance, one per accounting type:

<code>svc:/system/extended-accounting:flow</code>	Flow accounting
<code>svc:/system/extended-accounting:process</code>	Process accounting
<code>svc:/system/extended-accounting:task</code>	Task accounting
<code>svc:/system/extended-accounting:net</code>	Network accounting

The instances are enabled or disabled by `acctadm` as needed. Configuration changes are made using `acctadm`; service properties should not be modified directly using `svccfg(1M)`.

Users can manage extended accounting (start accounting, stop accounting, change accounting configuration parameters) if they have the appropriate RBAC Rights profile for the accounting type to be managed:

- Extended Accounting Flow Management
- Extended Accounting Process Management
- Extended Accounting Task Management
- Extended Accounting Network Management

The preceding profiles are for, respectively, flow accounting, process accounting, task accounting, and network accounting.

引用名	acctcms – command summary from process accounting records
用法概要	<pre>/usr/lib/acct/acctcms [-a [-o] [-p]] [-c] [-j] [-n] [-s] [-t] filename...</pre>
描述	acctcms reads one or more <i>filenames</i> , normally in the form described in acct.h(3HEAD) . It adds all records for processes that executed identically named commands, sorts them, and writes them to the standard output, normally using an internal summary format.
选项	-a Print output in ASCII rather than in the internal summary format. The output includes command name, number of times executed, total kcore-minutes, total CPU minutes, total real minutes, mean size (in K), mean CPU minutes per invocation, "hog factor," characters transferred, and blocks read and written, as in acctcom(1). Output is normally sorted by total kcore-minutes. Use the following options only with the -a option: -o Output a (non-prime) offshift-time-only command summary. -p Output a prime-time-only command summary. When -o and -p are used together, a combination prime-time and non-prime-time report is produced. All the output summaries are total usage except number of times executed, CPU minutes, and real minutes, which are split into prime and non-prime. -c Sort by total CPU time, rather than total kcore-minutes. -j Combine all commands invoked only once under "***other". -n Sort by number of command invocations. -s Any file names encountered hereafter are already in internal summary format. -t Process all records as total accounting records. The default internal summary format splits each field into prime and non-prime-time parts. This option combines the prime and non-prime time parts into a single field that is the total of both, and provides upward compatibility with old style acctcms internal summary format records.
示例	示例 1 Using the acctcms command. A typical sequence for performing daily command accounting and for maintaining a running total is: <pre>example% acctcms filename ... > today example% cp total previoustotal example% acctcms -s today previoustotal > total example% acctcms -a -s today</pre>
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/accounting/legacy-accounting

另请参见 [acctcom\(1\)](#), [acct\(1M\)](#), [acctcon\(1M\)](#), [acctmerg\(1M\)](#), [acctprc\(1M\)](#), [acctsh\(1M\)](#), [fwtmp\(1M\)](#), [runacct\(1M\)](#), [acct\(2\)](#), [acct.h\(3HEAD\)](#), [utmpx\(4\)](#), [attributes\(5\)](#)

附注 Unpredictable output results if -t is used on new style internal summary format files, or if it is not used with old style internal summary format files.

引用名	acctcon, acctcon1, acctcon2 – connect-time accounting								
用法概要	<pre>/usr/lib/acct/acctcon [-l <i>lineuse</i>] [-o <i>reboot</i>] /usr/lib/acct/acctcon1 [-p] [-t] [-l <i>lineuse</i>] [-o <i>reboot</i>] /usr/lib/acct/acctcon2</pre>								
描述	acctcon converts a sequence of login/logoff records to total accounting records (see the tacct format in acct.h(3HEAD)). The login/logoff records are read from standard input. The file /var/adm/wtmpx is usually the source of the login/logoff records; however, because it might contain corrupted records or system date changes, it should first be fixed using wtmpfix. The fixed version of file /var/adm/wtmpx can then be redirected to acctcon. The tacct records are written to standard output. acctcon is a combination of the programs acctcon1 and acctcon2. acctcon1 converts login/logoff records, taken from the fixed /var/adm/wtmpx file, to ASCII output. acctcon2 reads the ASCII records produced by acctcon1 and converts them to tacct records. acctcon1 can be used with the -l and -o options, described below, as well as with the -p and -t options.								
选项	<table><tr><td>-p</td><td>Print input only, showing line name, login name, and time (in both numeric and date/time formats).</td></tr><tr><td>-t</td><td>acctcon1 maintains a list of lines on which users are logged in. When it reaches the end of its input, it emits a session record for each line that still appears to be active. It normally assumes that its input is a current file, so that it uses the current time as the ending time for each session still in progress. The -t flag causes it to use, instead, the last time found in its input, thus assuring reasonable and repeatable numbers for non-current files.</td></tr><tr><td>-l <i>lineuse</i></td><td><i>lineuse</i> is created to contain a summary of line usage showing line name, number of minutes used, percentage of total elapsed time used, number of sessions charged, number of logins, and number of logoffs. This file helps track line usage, identify bad lines, and find software and hardware oddities. Hangup, termination of login(1) and termination of the login shell each generate logoff records, so that the number of logoffs is often three to four times the number of sessions. See init(1M) and utmpx(4).</td></tr><tr><td>-o <i>reboot</i></td><td>reboot is filled with an overall record for the accounting period, giving starting time, ending time, number of reboots, and number of date changes.</td></tr></table>	-p	Print input only, showing line name, login name, and time (in both numeric and date/time formats).	-t	acctcon1 maintains a list of lines on which users are logged in. When it reaches the end of its input, it emits a session record for each line that still appears to be active. It normally assumes that its input is a current file, so that it uses the current time as the ending time for each session still in progress. The -t flag causes it to use, instead, the last time found in its input, thus assuring reasonable and repeatable numbers for non-current files.	-l <i>lineuse</i>	<i>lineuse</i> is created to contain a summary of line usage showing line name, number of minutes used, percentage of total elapsed time used, number of sessions charged, number of logins, and number of logoffs. This file helps track line usage, identify bad lines, and find software and hardware oddities. Hangup, termination of login(1) and termination of the login shell each generate logoff records, so that the number of logoffs is often three to four times the number of sessions. See init(1M) and utmpx(4) .	-o <i>reboot</i>	reboot is filled with an overall record for the accounting period, giving starting time, ending time, number of reboots, and number of date changes.
-p	Print input only, showing line name, login name, and time (in both numeric and date/time formats).								
-t	acctcon1 maintains a list of lines on which users are logged in. When it reaches the end of its input, it emits a session record for each line that still appears to be active. It normally assumes that its input is a current file, so that it uses the current time as the ending time for each session still in progress. The -t flag causes it to use, instead, the last time found in its input, thus assuring reasonable and repeatable numbers for non-current files.								
-l <i>lineuse</i>	<i>lineuse</i> is created to contain a summary of line usage showing line name, number of minutes used, percentage of total elapsed time used, number of sessions charged, number of logins, and number of logoffs. This file helps track line usage, identify bad lines, and find software and hardware oddities. Hangup, termination of login(1) and termination of the login shell each generate logoff records, so that the number of logoffs is often three to four times the number of sessions. See init(1M) and utmpx(4) .								
-o <i>reboot</i>	reboot is filled with an overall record for the accounting period, giving starting time, ending time, number of reboots, and number of date changes.								
示例	示例 1 Using the acctcon command. The acctcon command is typically used as follows: <pre>example% acctcon -l lineuse -o reboots < tmpwtmp > ctacct</pre> The acctcon1 and acctcon2 commands are typically used as follows:								

示例 1 Using the acctcon command. (续)

```
example% acctcon1 -l lineuse -o reboots < tmpwtmp | sort +1n +2 > ctmp
example% acctcon2 < ctmp > ctacct
```

文件 /var/adm/wtmpx History of user access and administration information

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/accounting/legacy-accounting

另请参见 [acctcom\(1\)](#), [login\(1\)](#), [acct\(1M\)](#), [acctcms\(1M\)](#), [acctmerg\(1M\)](#), [acctprc\(1M\)](#), [acctsh\(1M\)](#), [fwtmp\(1M\)](#), [init\(1M\)](#), [runacct\(1M\)](#), [acct\(2\)](#), [acct.h\(3HEAD\)](#), [utmpx\(4\)](#), [attributes\(5\)](#)

《Oracle Solaris 管理：常见任务》

附注 The line usage report is confused by date changes. Use `wtmpfix` (see [fwtmp\(1M\)](#)), with the `/var/adm/wtmpx` file as an argument, to correct this situation.

During a single invocation of any given command, the `acctcon`, `acctcon1`, and `acctcon2` commands can process a maximum of:

- 6000 distinct session
- 1000 distinct terminal lines
- 2000 distinct login names

If at some point the actual number of any one of these items exceeds the maximum, the command will not succeed.

引用名	acctprc, acctprc1, acctprc2 – process accounting
用法概要	<pre>/usr/lib/acct/acctprc /usr/lib/acct/acctprc1 [ctmp] /usr/lib/acct/acctprc2</pre>
描述	acctprc reads the standard input and converts it to total accounting records (see the tacct record in acct.h(3HEAD)). acctprc divides CPU time into prime time and non-prime time and determines mean memory size (in memory segment units). acctprc then summarizes the tacct records, according to user IDs, and adds login names corresponding to the user IDs. The summarized records are then written to the standard output. acctprc1 reads input in the form described by acct.h(3HEAD), adds login names corresponding to user IDs, then writes for each process an ASCII line giving user ID, login name, prime CPU time (tics), non-prime CPU time (tics), and mean memory size (in memory segment units). If <i>ctmp</i> is given, it should contain a list of login sessions sorted by user ID and login name. If this file is not supplied, it obtains login names from the password file, just as acctprc does. The information in <i>ctmp</i> helps it distinguish between different login names that share the same user ID. From the standard input, acctprc2 reads records in the form written by acctprc1, summarizes them according to user ID and name, then writes the sorted summaries to the standard output as total accounting records.
示例	示例1 Examples of acctprc. The acctprc command is typically used as shown below: <pre>example% acctprc < /var/adm/pacct > ptacct</pre> The acctprc1 and acctprc2s commands are typically used as shown below: <pre>example% acctprc1 ctmp </var/adm/pacct example% acctprc2 > ptacct</pre>
文件	/etc/passwd system password file
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/accounting/legacy-accounting

另请参见 [acctcom\(1\)](#), [acct\(1M\)](#), [acctcms\(1M\)](#), [acctcon\(1M\)](#), [acctmerg\(1M\)](#), [acctsh\(1M\)](#), [cron\(1M\)](#), [fwtmp\(1M\)](#), [runacct\(1M\)](#), [acct\(2\)](#), [acct.h\(3HEAD\)](#), [utmpx\(4\)](#), [attributes\(5\)](#)

附注 Although it is possible for acctprc1 to distinguish among login names that share user IDs for commands run from a command line, it is difficult for acctprc1 to make this distinction for commands invoked in other ways. A command run from [cron\(1M\)](#) is an example of where

acctprc1 might have difficulty. A more precise conversion can be done using the acctwtmp program in [acct\(1M\)](#). acctprc does not distinguish between users with identical user IDs.

A memory segment of the mean memory size is a unit of measure for the number of bytes in a logical memory segment on a particular processor.

During a single invocation of any given command, the acctprc, acctprc1, and acctprc2 commands can process a maximum of

- 6000 distinct sessions
- 1000 distinct terminal lines
- 2000 distinct login names

If at some point the actual number of any one of these items exceeds the maximum, the command will not succeed.

引用名	acctsh, chargefee, ckpacct, dodisk, lastlogin, monacct, nulladm, prctmp, prdaily, prtacct, shutacct, startup, turnacct – shell procedures for accounting
用法概要	<pre>/usr/lib/acct/chargefee login-name number /usr/lib/acct/ckpacct [blocks] /usr/lib/acct/dodisk [-o] [filename]... /usr/lib/acct/lastlogin /usr/lib/acct/monacct number /usr/lib/acct/nulladm filename... /usr/lib/acct/prctmp filename /usr/lib/acct/prdaily [-c] [-l] [mdd] /usr/lib/acct/prtacct filename ['' heading ''] /usr/lib/acct/shutacct ['' reason ''] /usr/lib/acct/startup /usr/lib/acct/turnacct on off switch</pre>
描述	
chargefee Command	chargefee can be invoked to charge a <i>number</i> of units to <i>login-name</i> . A record is written to <i>/var/adm/fee</i> , to be merged with other accounting records by runacct(1M) .
ckpacct Command	ckpacct should be initiated using cron(1M) to periodically check the size of <i>/var/adm/pacct</i> . If the size exceeds <i>blocks</i> , 500 by default, turnacct will be invoked with argument switch. To avoid a conflict with turnacct switch execution in runacct, do not run ckpacct and runacct simultaneously. If the number of free disk blocks in the <i>/var</i> file system falls below 500, ckpacct will automatically turn off the collection of process accounting records via the off argument to turnacct. When at least 500 blocks are restored, the accounting will be activated again on the next invocation of ckpacct. This feature is sensitive to the frequency at which ckpacct is executed, usually by the cron(1M) command.
dodisk Command	dodisk should be invoked by cron(1M) to perform the disk accounting functions.
lastlogin Command	lastlogin is invoked by runacct(1M) to update <i>/var/adm/acct/sum/loginlog</i> , which shows the last date on which each person logged in.
monacct Command	monacct should be invoked once each month or each accounting period. <i>number</i> indicates which month or period it is. If <i>number</i> is not given, it defaults to the current month (01–12). This default is useful if monacct is to be executed using cron(1M) on the first day of each month. monacct creates summary files in <i>/var/adm/acct/fiscal</i> and restarts the summary files in <i>/var/adm/acct/sum</i> .

nulladm Command	nulladm creates <i>filename</i> with mode 664 and ensures that owner and group are adm. It is called by various accounting shell procedures.
prctmp Command	prctmp can be used to print the session record file (normally /var/adm/acct/nite/ctmp created by acctcon1 (see acctcon(1M))).
prdaily Command	prdaily is invoked by runacct(1M) to format a report of the previous day's accounting data. The report resides in /var/adm/acct/sum/rprt/mmdd where <i>mmdd</i> is the month and day of the report. The current daily accounting reports may be printed by typing prdaily. Previous days' accounting reports can be printed by using the <i>mmdd</i> option and specifying the exact report date desired.
prtacct Command	prtacct can be used to format and print any total accounting (tacct)file.
shutacct Command	shutacct is invoked during a system shutdown to turn process accounting off and append a <i>reason</i> record to /var/adm/wtmpx.
startup Command	startup can be invoked when the system is brought to a multi-user state to turn process accounting on.
turnacct Command	turnacct is an interface to accton (see acct(1M)) to turn process accounting on or off. The switch argument moves the current /var/adm/pacct to the next free name in /var/adm/pacct. <i>incr</i> (where <i>incr</i> is a number starting with 0 and incrementing by one for each additional pacct file), then turns accounting back on again. This procedure is called by ckpacct and thus can be taken care of by the cron(1M) command and used to keep pacct to a reasonable size. shutacct uses turnacct to stop process accounting. startup uses turnacct to start process accounting.

选项

- The following options are supported:
- c This option prints a report of exceptional resource usage by command, and may be used on current day's accounting data only.
 - l This option prints a report of exceptional usage by login id for the specified date. Previous daily reports are cleaned up and therefore inaccessible after each invocation of monacct.
 - o This option uses acctdusg (see [acct\(1M\)](#)) to do a slower version of disk accounting by login directory. *filenames* specifies the one or more filesystem names where disk accounting will be done. If *filenames* are used, disk accounting will be done on these filesystems only. If the -o option is used, *filenames* should be mount points of mounted filesystems. If the -o option is omitted, *filenames* should be the special file names of mountable filesystems.

文件	/etc/logadm.conf	Configuration file for the logadm(1M) command
	/usr/lib/acct	Holds all accounting commands listed in section 1M of this manual

/usr/lib/acct/ptecms.awk	Contains the limits for exceptional usage by command name
/usr/lib/acct/ptelus.awk	Contains the limits for exceptional usage by login ID
/var/adm/acct/fiscal	Fiscal reports directory
/var/adm/acct/nite	Working directory
/var/adm/acct/sum	Summary directory that contains information for monacct
/var/adm/acct/sum/loginlog	File updated by last login
/var/adm/fee	Accumulator for fees
/var/adm/pacct	Current file for per-process accounting
/var/adm/pacctincr	Used if pacct gets large and during execution of daily accounting procedure
/var/adm/wtmpx	History of user access and administration information

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/accounting/legacy-accounting

另请参见 [acctcom\(1\)](#), [acct\(1M\)](#), [acctcms\(1M\)](#), [acctcon\(1M\)](#), [acctmerg\(1M\)](#), [acctprc\(1M\)](#), [cron\(1M\)](#), [fwtmp\(1M\)](#), [logadm\(1M\)](#), [runacct\(1M\)](#), [acct\(2\)](#), [acct.h\(3HEAD\)](#), [utmpx\(4\)](#), [attributes\(5\)](#)

附注 See [runacct\(1M\)](#) for the main daily accounting shell script, which performs the accumulation of connect, process, fee, and disk accounting on a daily basis. It also creates summaries of command usage.

引用名 acpihp – ACPI（高级配置和电源接口）热插拔守护进程

用法概要 /usr/platform/i86pc/lib/acpihp

描述 ACPI 热插拔守护进程 `acpihp` 是一个在 x86 平台上运行的守护进程。该守护进程由服务管理工具启动。`acpihp` 与 `syseventd(1M)` 和 `cfgadm(1M)` 子系统进行通信，处理热插拔设备的 ACPI 事件。

`acpihp` 的服务 FMRI（故障管理资源标识符）如下：

`svc:/platform/i86pc/acpihp:default`

请注意，`acpihp` 是一个专用接口。

选项 `acpihp` 守护进程不支持任何选项。

错误 `acpihp` 使用 `syslog(3C)` 报告状态和错误消息。所有的消息由 `LOG_DAEMON` 工具记录。

错误消息记录的优先级为 `LOG_ERR` 和 `LOG_NOTICE`，提示性消息记录的优先级为 `LOG_DEBUG`。`/etc/syslog.conf` 文件中的缺省项将所有的错误消息记录到 `/var/adm/messages` 日志中。

属性 有关以下属性的说明，请参见 `attributes(5)`：

属性类型	属性值
可用性	system/kernel/dynamic-reconfiguration/i86pc
接口稳定性	Private（专用）

另请参见 `svcs(1)`、`cfgadm(1M)`、`inetadm(1M)`、`svcadm(1M)`、`syseventd(1M)`、`syslog(3C)`、`syslog.conf(4)`、

附注 `acpihp` 服务由服务管理工具 `smf(5)` 管理，其服务标识符为：

`svc:/platform/i86pc/acpihp:default`

引用名

adbgen – generate adb script

用法概要

```
/usr/lib/adb/adbgen [-m model] filename.adb . . .
```

描述

adbgen makes it possible to write [adb\(1\)](#) scripts that do not contain hard-coded dependencies on structure member offsets. The input to adbgen is a file named *filename.adb* that contains header information, then a null line, then the name of a structure, and finally an adb script. adbgen only deals with one structure per file; all member names are assumed to be in this structure. The output of adbgen is an adb script in *filename.adb*. adbgen operates by generating a C program which determines structure member offsets and sizes, which in turn generate the adb script.

The header lines, up to the null line, are copied verbatim into the generated C program. Typically, these are `#include` statements, which include the headers containing the relevant structure declarations.

The adb script part may contain any valid adb commands (see [adb\(1\)](#)), and may also contain adbgen requests, each enclosed in braces (`{ }`). Request types are:

- Print a structure member. The request form is `{member, format}`. *member* is a member name of the *structure* given earlier, and *format* is any valid adb format request or any of the adbgen format specifiers (such as `{POINTER}`) listed below. For example, to print the `p_pid` field of the *proc* structure as a decimal number, you would write `{p_pid, d}`.
- Print the appropriate adb format character for the given adbgen format specifier. This action takes the data model into consideration. The request form is `{format specifier}`. The valid adbgen format specifiers are:

<code>{POINTER}</code>	pointer value in hexadecimal
<code>{LONGDEC}</code>	long value in decimal
<code>{ULONGDEC}</code>	unsigned long value in decimal
<code>{ULONGHEX}</code>	unsigned long value in hexadecimal
<code>{LONGOCT}</code>	long value in octal
<code>{ULONGOCT}</code>	unsigned long value in octal

- Reference a structure member. The request form is `{*member, base}`. *member* is the member name whose value is desired, and *base* is an adb register name which contains the base address of the structure. For example, to get the `p_pid` field of the *proc* structure, you would get the *proc* structure address in an adb register, for example `<f`, and write `{*p_pid, <f}`.
- Tell adbgen that the offset is valid. The request form is `{OFFSETOK}`. This is useful after invoking another adb script which moves the adb *dot*.
- Get the size of the *structure*. The request form is `{SIZEOF}`. adbgen replaces this request with the size of the structure. This is useful in incrementing a pointer to step through an array of structures.

- Calculate an arbitrary C expression. The request form is {EXPR, *expression*}. adbggen replaces this request with the value of the expression. This is useful when more than one structure is involved in the script.
- Get the offset to the end of the structure. The request form is {END}. This is useful at the end of the structure to get adb to align the *dot* for printing the next structure member.

adbgen keeps track of the movement of the adb *dot* and generates adb code to move forward or backward as necessary before printing any structure member in a script. adbggen's model of the behavior of adb's *dot* is simple: it is assumed that the first line of the script is of the form *struct_address/adb text* and that subsequent lines are of the form *+/adb text*. The adb *dot* then moves in a sane fashion. adbggen does not check the script to ensure that these limitations are met. adbggen also checks the size of the structure member against the size of the adb format code and warns if they are not equal.

选项

The following option is supported:

-m *model* Specifies the data type model to be used by adbggen for the macro. This affects the outcome of the {*format specifier*} requests described under DESCRIPTION and the offsets and sizes of data types. *model* can be *ilp32* or *lp64*. If the -m option is not given, the data type model defaults to *ilp32*.

操作数

The following operand is supported:

***filename*.adb** Input file that contains header information, followed by a null line, the name of the structure, and finally an adb script.

示例

示例 1 A sample adbggen file.

For an include file *x.h* which contained

```
struct x {
    char    *x_cp;
    char    x_c;
    int     x_i;
};
```

then, an adbggen file (call it *script.adb*) to print the file *x.h* would be:

```
#include "x.h"
x
./"x_cp"16t"x_c"8t"x_i"n{x_cp,{POINTER}}{x_c,C}{x_i,D}
```

After running adbggen as follows,

```
% /usr/lib/adb/adbggen script.adb
```

the output file *script* contains:

```
./"x_cp"16t"x_c"8t"x_i"nXC3+D
```

示例1 A sample adbgen file. (续)

For a macro generated for a 64-bit program using the lp64 data model as follows,

```
% /usr/lib/adb/adbgen/ -m lp64 script.adb
```

the output file script would contain:

```
./"x_cp"16t"x_c"8t"x_i"nJC3+D
```

To invoke the script, type:

```
example% adb program
x$<script
```

文件

```
/usr/platform/platform-name/lib/adb/*
platform-specific adb scripts for debugging the 32-bit kernel

/usr/platform/platform-name/lib/adb/sparcv9/*
platform-specific adb scripts for debugging the 64-bit SPARC V9 kernel

/usr/lib/adb/*
adb scripts for debugging the 32-bit kernel

/usr/lib/adb/sparcv9/*
adb scripts for debugging the 64-bit SPARC V9 kernel
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	developer/debug/mdb

另请参见

[adb\(1\)](#), [kmdb\(1\)](#), [uname\(1\)](#), [attributes\(5\)](#)

诊断

Warnings are given about structure member sizes not equal to adb format items and about badly formatted requests. The C compiler complains if a structure member that does not exist is referenced. It also complains about an ampersand before array names; these complaints may be ignored.

附注

platform-name can be found using the -i option of [uname\(1\)](#).

已知问题

adb syntax is ugly; there should be a higher level interface for generating scripts.

Structure members which are bit fields cannot be handled because C will not give the address of a bit field. The address is needed to determine the offset.

引用名	add_allocatable – add entries to allocation databases
用法概要	<pre>/usr/sbin/add_allocatable [-f] [-s] [-d] -n name -t type -l device-list [-a authorization] [-c clean] [-o key=value]</pre>
描述	add_allocatable creates new entries for user allocatable devices that are to be managed by the device allocation mechanism. add_allocatable can also be used to update existing entries of such devices. add_allocatable can also create and update entries for non-allocatable devices, such as printers, whose label range is managed by the device allocation mechanism. add_allocatable can be used in shell scripts, such as installation scripts for driver packages, to automate the administrative work of setting up a new device. Use list_devices(1) to see the names and types of allocatable devices, their attributes, and device paths.
选项	-f Force an update of an already-existing entry with the specified information. add_allocatable exits with an error if this option is not specified when an entry with the specified device name already exists. -s Turn on silent mode. add_allocatable does not print any error or warning messages. -d If this option is present, add_allocatable updates the system-supplied default attributes of the device type specified with -t. -n name Adds or updates an entry for device that is specified by name. -t type Adds or updates device entries that are of a type that are specified by type. -l device-list Adds or updates device paths to the device that is specified with -n. Multiple paths in device-list must be separated by white spaces and the list must be quoted. -a authorization Adds or updates the authorization that is associated with either the device that is specified with -n or with devices of the type that is specified with -t. When more than one authorization is specified, the list must be separated by commas and must be quoted. When the device is not allocatable, authorization is specified with an asterisk (*) and must be quoted. When the device is allocatable by any user, authorization is specified with the at sign (@) and must be quoted. Default authorization is '@'. -c clean Specifies the device_clean(5) program clean to be used with the device that is specified with -n or with devices of the type that is specified with -t. The default clean program is /bin/true.

- o *key=value*

Accepts a string of colon-separated *key=value* pairs for a device that is specified with -n or with devices of the type that is specified with -t. The following keys are currently interpreted by the system:
- minlabel

The minimum label at which the device can be used.
- maxlabel

The maximum label at which the device can be used.
- class

Specifies a logical grouping of devices. For example, all Sun Ray devices of all device types is a logical grouping. The class keyword has no default value.
- xdpi

Specifies the display name of the X session. This keyword is used to identify devices that are associated with the X session. The xdpi keyword has no default value.

退出状态

When successful, add_allocate returns an exit status of 0 (true). add_allocate returns a nonzero exit status in the event of an error. The exit codes are as follows:

1

Invocation syntax error

2

Unknown system error

3

An entry already exists for the specified device. This error occurs only when the -f option is not specified.

4

Permission denied. User does not have DAC or MAC access record updates.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/trusted
Interface Stability	See below.

The invocation is Uncommitted. The options are Uncommitted. The output is Not-an-Interface.

另请参见

[allocate\(1\)](#), [deallocate\(1\)](#), [list_devices\(1\)](#), [remove_allocatable\(1M\)](#), [attributes\(5\)](#), [device_clean\(5\)](#)

附注

The functionality described on this manual page is available only if the system is configured with Trusted Extensions.

引用名

addbadsec – map out defective disk blocks

用法概要

addbadsec [-p] [-a blkno [blkno]...] [-f filename] raw_device

描述

addbadsec is used by the system administrator to map out bad disk blocks. Normally, these blocks are identified during surface analysis, but occasionally the disk subsystem reports unrecoverable data errors indicating a bad block. A block number reported in this way can be fed directly into addbadsec, and the block will be remapped. addbadsec will first attempt hardware remapping. This is supported on SCSI drives and takes place at the disk hardware level. If the target is an IDE drive, then software remapping is used. In order for software remapping to succeed, the partition must contain an alternate slice and there must be room in this slice to perform the mapping.

It should be understood that bad blocks lead to data loss. Remapping a defective block does not repair a damaged file. If a bad block occurs to a disk-resident file system structure such as a superblock, the entire slice might have to be recovered from a backup.

选项

The following options are supported:

- a Adds the specified blocks to the hardware or software map. If more than one block number is specified, the entire list should be quoted and block numbers should be separated by white space.
- f Adds the specified blocks to the hardware or software map. The bad blocks are listed, one per line, in the specified file.
- p Causes addbadsec to print the current software map. The output shows the defective block and the assigned alternate. This option cannot be used to print the hardware map.

操作数

The following operand is supported:

raw_device The address of the disk drive (see FILES).

文件

The raw device should be /dev/rdisk/c?[t?]d?p0. See [disks\(1M\)](#) for an explanation of SCSI and IDE device naming conventions.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86
Availability	system/core-os

另请参见

[disks\(1M\)](#), [diskscan\(1M\)](#), [fdisk\(1M\)](#), [fmthard\(1M\)](#), [format\(1M\)](#), [attributes\(5\)](#)

附注

The `format(1M)` utility is available to format, label, analyze, and repair SCSI disks. This utility is included with the `addbadsec`, `diskscan(1M)`, `fdisk(1M)`, and `fmthard(1M)` commands available for x86. To format an IDE disk, use the DOS "format" utility; however, to label, analyze, or repair IDE disks on x86 systems, use the Solaris `format(1M)` utility.

引用名

add_drv – add a new device driver to the system

用法概要

```
add_drv [-b basedir] [-c class_name]
        [-i 'identify_name...'] [-m 'permission', '...']
        [-p 'policy'] [-P privilege] [-n] [-f] [-u] [-v] device_driver
```

描述

The add_drv command is used to inform the system about newly installed device drivers.

Each device on the system has a name associated with it. This name is represented by the name property for the device. Similarly, the device may also have a list of driver names associated with it. This list is represented by the compatible property for the device.

The system determines which devices will be managed by the driver being added by examining the contents of the name property and the compatible property (if it exists) on each device. If the value in the name property does not match the driver being added, each entry in the compatible property is tried, in order, until either a match occurs or there are no more entries in the compatible property.

In some cases, adding a new driver may require a reconfiguration boot. See the NOTES section.

Aliases might require quoting (with double-quotes) if they contain numbers. See EXAMPLES.

The /etc/minor_perm
File

add_drv and [update_drv\(1M\)](#) read the /etc/minor_perm file to obtain permission information. The permission specified is applied to matching minor nodes created when a device bound to the driver is attached. A minor node's permission may be manually changed by [chmod\(1\)](#). For such nodes, the specified permissions apply, overriding the default permissions specified via add_drv or [update_drv\(1M\)](#).

The format of the /etc/minor_perm file is as follows:

name:*minor_name permissions owner group*

minor_name may be the actual name of the minor node, or contain shell metacharacters to represent several minor nodes (see [sh\(1\)](#)).

For example:

```
sd:* 0640 root sys
zs:[a-z],cu 0600 uucp uucp
mm:kmem 0640 root bin
```

The first line sets all devices exported by the sd node to 0640 permissions, owned by root, with group sys. In the second line, devices such as a, cu and z, cu exported by the zs driver are set to 0600 permission, owned by uucp, with group uucp. In the third line the kmem device exported by the mm driver is set to 0640 permission, owned by root, with group bin.

Running add_drv from
a postinstall Script

When running add_drv from within the context of a package's postinstall script, you must consider whether the package is being added to a system image or to a running system. When a package is being installed on a system image, the BASEDIR variable refers to the image's base

directory. In this situation, `add_drv` should be invoked with `-b $BASEDIR`. This causes `add_drv` only to update the image's system files; a reboot of the system or client would be required to make the driver operational.

When a package is being installed on the running system itself, the system files need to be updated, as in the case above. However, the running kernel can be informed of the existence of the new driver without requiring a reboot. To accomplish this, the `postinstall` script must invoke `add_drv` without the `-b` option. Accordingly, `postinstall` scripts invoking `add_drv` should be written thusly:

```
if [ "${BASEDIR:=/}" = "/" ]
then
    ADD_DRV="add_drv"
else
    ADD_DRV="add_drv -b ${BASEDIR}"
fi
$ADD_DRV [<options>] <driver>
```

...or, alternatively:

```
if [ "${BASEDIR:=/}" != "/" ]
then
    BASEDIR_OPT="-b $BASEDIR"
fi
add_drv $BASEDIR_OPT [<options>] <driver>
```

The `-b` option is described below.

选项

`-b basedir`

Installs the driver on the system with a root directory of *basedir* rather than installing on the system executing `add_drv`. This option is typically used in package post-installation scripts when the package is not being installed on the system executing the `pkgadd` command. The system using *basedir* as its root directory must reboot to complete the driver installation.

注 – The root file system of any non-global zones must not be referenced with the `-b` option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See [zones\(5\)](#).

`-c class_name`

The driver being added to the system exports the class *class_name*.

`-f`

Normally if a reconfiguration boot is required to complete the configuration of the driver into the system, `add_drv` will not add the driver. The force flag forces `add_drv` to add the driver even if a reconfiguration boot is required. See the `-v` flag.

`-i 'identify_name'`

A white-space separated list of aliases for the driver *device_driver*.

-m '*permission*'

Specify the file system permissions for device nodes created by the system on behalf of *device_driver*.

-n

Do not try to load and attach *device_driver*, just modify the system configuration files for the *device_driver*.

-p '*policy*'

Specify an additional device security policy.

The device security policy consists of several whitespace separated tokens:

{minorspec {token=value}+}+

minorspec is a simple wildcard pattern for a minor device. A single * matches all minor devices. Only one * is allowed in the pattern.

Patterns are matched in the following order:

- entries without a wildcard
- entries with wildcards, longest wildcard first

The following tokens are defined: *read_priv_set* and *write_priv_set*. *read_priv_set* defines the privileges that need to be asserted in the effective set of the calling process when opening a device for reading. *write_priv_set* defines the privileges that need to be asserted in the effective set of the calling process when opening a device for writing. See [privileges\(5\)](#).

A missing minor spec is interpreted as a *.

-P '*privilege*'

Specify additional, comma separated, privileges used by the driver. You can also use specific privileges in the device's policy.

-u

Add the driver to the system, leaving it in an inactive state for later configuration with [devfsadm\(1M\)](#) -u. The -u behavior differs from -n in that -n only updates the system files, requiring a reboot to attach the driver. Drivers added with -u can be attached by running `devfsadm -u` without rebooting. Driver writers should verify their driver with this behavior. See NOTES for additional considerations. The -u option cannot be used together with -n or -b.

-v

The verbose flag causes `add_drv` to provide additional information regarding the success or failure of a driver's configuration into the system. See the EXAMPLES section.

示例

示例 1 Adding SUNW Example Driver to the System

The following example adds the SUNW,example driver to a 32-bit system, with an alias name of SUNW,alias. It assumes the driver has already been copied to /usr/kernel/drv.

```
example# add_drv -m '* 0666 bin bin', 'a 0644 root sys' \
    -p 'a write_priv_set=sys_config * write_priv_set=none' \
    -i 'SUNW,alias' SUNW,example
```

Every minor node created by the system for the SUNW,example driver will have the permission 0666, and be owned by user bin in the group bin, except for the minor device a, which will be owned by root, group sys, and have a permission of 0644. The specified device policy requires no additional privileges to open all minor nodes, except minor device a, which requires the sys_config privilege when opening the device for writing.

示例 2 Adding Driver to the Client /export/root/sun1

The following example adds the driver to the client /export/root/sun1. The driver is installed and loaded when the client machine, sun1, is rebooted. This second example produces the same result as the first, except the changes are on the diskless client, sun1, and the client must be rebooted for the driver to be installed.

```
example# add_drv -m '* 0666 bin bin', 'a 0644 root sys' \
    -i 'SUNW,alias' -b /export/root/sun1 \
    SUNW,example
```

See the note in the description of the -b option, above, specifying the caveat regarding the use of this option with the Solaris zones feature.

示例 3 Adding Driver for a Device Already Managed by an Existing Driver

The following example illustrates the case where a new driver is added for a device that is already managed by an existing driver. Consider a device that is currently managed by the driver dumb_framebuffer. The name and compatible properties for this device are as follows:

```
name="display"
compatible="whizzy_framebuffer", "dumb_framebuffer"
```

If add_drv is used to add the whizzy_framebuffer driver, the following will result.

```
example# add_drv whizzy_framebuffer
Error: Could not install driver (whizzy_framebuffer)
Device managed by another driver.
```

If the -v flag is specified, the following will result.

```
example# add_drv -v whizzy_framebuffer
Error: Could not install driver (whizzy_framebuffer)
Device managed by another driver.
Driver installation failed because the following
entries in /devices would be affected:
```

示例 3 Adding Driver for a Device Already Managed by an Existing Driver (续)

```
/devices/iommu@f,e0000000/sbus@f,e0001000/display[:*]
(Device currently managed by driver "dumb_framebuffer")
```

The following entries in /dev would be affected:

```
/dev/fbs/dumb_framebuffer0
```

If the -v and -f flags are specified, the driver will be added resulting in the following.

```
example# add_drv -vf whizzy_framebuffer
A reconfiguration boot must be performed to complete the
installation of this driver.
```

The following entries in /devices will be affected:

```
/devices/iommu@f,e0000000/sbus@f,e0001000/display[:*]
(Device currently managed by driver "dumb_framebuffer")
```

The following entries in /dev will be affected:

```
/dev/fbs/dumb_framebuffer0
```

The above example is currently only relevant to devices exporting a generic device name.

示例 4 Use of Double Quotes in Specifying Driver Alias

The following example shows the use of double quotes in specifying a driver alias that contains numbers.

```
example# add_drv -i "pci10c5,25" smc
```

add_drv returns 0 on success and 1 on failure.

```
/kernel/drv
  32-bit boot device drivers
/kernel/drv/sparcv9
  64-bit SPARC boot device drivers
```

```
/kernel/drv/amd64
  64-bit x86 boot device drivers
```

```
/usr/kernel/drv
  other 32-bit drivers that could potentially be shared between platforms
```

```
/usr/kernel/drv/sparcv9
  other 64-bit SPARC drivers that could potentially be shared between platforms
```

退出状态
文件

`/usr/kernel/drv/amd64`
other 64-bit x86 drivers that could potentially be shared between platforms

`/platform/'uname -i'/kernel/drv`
32-bit platform-dependent drivers

`/platform/'uname -i'/kernel/drv/sparcv9`
64-bit SPARC platform-dependent drivers

`/platform/'uname -i'/kernel/drv/amd64`
64-bit x86 platform-dependent drivers

`/etc/driver_aliases`
driver aliases file

`/etc/driver_classes`
driver classes file

`/etc/minor_perm`
minor node permissions

`/etc/name_to_major`
major number binding

`/etc/security/device_policy`
device policy

`/etc/security/extra_privs`
device privileges

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [boot\(1M\)](#), [chmod\(1\)](#), [devfsadm\(1M\)](#), [kernel\(1M\)](#), [modinfo\(1M\)](#), [rem_drv\(1M\)](#), [update_drv\(1M\)](#), [driver.conf\(4\)](#), [system\(4\)](#), [attributes\(5\)](#), [privileges\(5\)](#), [devfs\(7FS\)](#), [ddi_create_minor_node\(9F\)](#)

《Writing Device Drivers》

附注 It is possible to add a driver for a device already being managed by a different driver, where the driver being added appears in the device's `compatible` list before the current driver. In such cases, a reconfiguration boot is required (see [boot\(1M\)](#) and [kernel\(1M\)](#)). After the reconfiguration boot, device links in `/dev` and references to these files may no longer be valid (see the `-v` flag). If a reconfiguration boot would be required to complete the driver installation, `add_drv` will fail unless the `-f` option is specified. See Example 3 in the EXAMPLES section.

With the introduction of the device policy several drivers have had their minor permissions changed and a device policy instated. The typical network driver should use the following device policy:

```
add_drv -p 'read_priv_set=net_rawaccess\  
    write_priv_set=net_rawaccess' -m '* 666 root sys'\  
mynet
```

This document does not constitute an API. `/etc/minor_perm`, `/etc/name_to_major`, `/etc/driver_classes`, and `/devices` may not exist or may have different contents or interpretations in a future release. The existence of this notice does not imply that any other documentation that lacks this notice constitutes an API.

`/etc/minor_perm` can only be updated by `add_drv(1M)`, [rem_drv\(1M\)](#) or [update_drv\(1M\)](#).

In the current version of `add_drv`, the use of double quotes to specify an alias is optional when used from the command line. However, when using `add_drv` from packaging scripts, you should continue to use double quotes to specify an alias.

Some drivers should not be added and configured on the system directly, but should only be configured as the system boots. The reasons for this restriction include, but are not limited to, a driver dependency on configuration early during boot or a dependency on some kernel component being installed or updated at the same time as the driver is being added. Such drivers should only be added to the system with the `-n` flag, so the driver is only loaded and configured when the system is rebooted, thus assuring an environment in which the driver can be configured properly.

已知问题

Previous versions of `add_drv` accepted a pathname for *device_driver*. This feature is no longer supported and results in failure.

引用名	arp – address resolution display and control
用法概要	<pre>arp <i>hostname</i> arp -a [-n] arp -d <i>hostname</i> arp -f <i>filename</i> arp -s <i>hostname ether_address</i> [temp] [pub] [trail] [permanent]</pre>
描述	The <code>arp</code> program displays and modifies the Internet-to-MAC address translation tables used by the address resolution protocol (see arp(7P)). With no flags, the program displays the current ARP entry for <i>hostname</i>. The host may be specified by name or by number, using Internet dot notation. Options that modify the ARP translation tables (<code>-d</code>, <code>-f</code>, and <code>-s</code>) can be used only when the invoked command is granted the <code>sys_net_config</code> privilege. See privileges(5).
选项	<code>-a</code> Display all of the current ARP entries. The definition for the flags in the table are: <code>d</code> Unverified; this is a local IP address that is currently undergoing Duplicate Address Detection. ARP will not respond to requests for this address until Duplicate Address Detection completes. <code>o</code> Old; this entry is aging away. If IP requests it again, a new ARP query will be generated. This state is used for detecting peer address changes. <code>y</code> Delayed; periodic address defense and conflict detection was unable to send a packet due to internal network use limits for non-traffic-related messages (100 packets per hour per interface). This occurs only on interfaces with very large numbers of aliases. <code>A</code> Authority; this machine is authoritative for this IP address. ARP will not accept updates from other machines for this entry. <code>L</code> Local; this is a local IP address configured on one of the machine's logical interfaces. ARP will defend this address if another node attempts to claim it. <code>M</code> Mapping; only used for the multicast entry for 224.0.0.0 <code>P</code> Publish; includes IP address for the machine and the addresses that have explicitly been added by the <code>-s</code> option. ARP will respond to ARP requests for this address. <code>S</code> Static; entry cannot be changed by learned information. This indicates that the permanent flag was used when creating the entry. <code>U</code> Unresolved; waiting for ARP response.

You can use the `-n` option with the `-a` option to disable the automatic numeric IP address-to-name translation. Use `arp -an` or `arp -na` to display numeric IP addresses. The `arp -a` option is equivalent to:

```
# netstat -p -f inet
```

...and `-an` and `-na` are equivalent to:

```
# netstat -pn -f inet
```

- d Delete an entry for the host called *hostname*.

Note that ARP entries for IPMP (IP Network Multipathing) data and test addresses are managed by the kernel and thus cannot be deleted.

- f Read the file named *filename* and set multiple entries in the ARP tables. Entries in the file should be of the form:

```
hostname MACaddress [temp] [pub] [trail] [permanent]
```

See the `-s` option for argument definitions.

- s Create an ARP entry for the host called *hostname* with the MAC address *MACaddress*. For example, an Ethernet address is given as six hexadecimal bytes separated by colons. The entry will not be subject to deletion by aging unless the word `temp` is specified in the command. If the word `pub` is specified, the entry will be published, which means that this system will respond to ARP requests for *hostname* even though the *hostname* is not its own. The word `permanent` indicates that the system will not accept MAC address changes for *hostname* from the network.

Solaris does not implement trailer encapsulation, and the word `trail` is accepted on entries for compatibility only.

`arp -s` can be used for a limited form of proxy ARP when a host on one of the directly attached networks is not physically present on a subnet. Another machine can then be configured to respond to ARP requests using `arp -s`. This is useful in certain SLIP configurations.

Non-temporary proxy ARP entries for an IPMP (IP Network Multipathing) group are automatically managed by the kernel. Specifically, if the hardware address in an entry matches the hardware address of an IP interface in an IPMP group, and the IP address is not local to the system, this will be regarded as an IPMP proxy ARP entry. This entry will have its hardware address automatically adjusted in order to keep the IP address reachable so long as the IPMP group has not entirely failed.

ARP entries must be consistent across an IPMP group. Therefore, ARP entries cannot be associated with individual underlying IP interfaces in an IPMP group, and must instead be associated with the corresponding IPMP IP interface.

Note that ARP entries for IPMP data and test addresses are managed by the kernel and thus cannot be changed.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [ifconfig\(1M\)](#), [netstat\(1M\)](#), [attributes\(5\)](#), [privileges\(5\)](#), [arp\(7P\)](#)

引用名

asradm – Auto Service Request registration utility

用法概要

```
asradm list

asradm send [-n] [activate|audit|deactivate|fault|heartbeat|test]

asradm register -u username [-p password-file] [-e endpoint_URL]

asradm set-proxy [-i] | [-h host[:port]] [-u username] [-p password-file]

asradm unregister
```

描述

The `asradm` utility is used by a privileged system administrator to register hosts for enrollment in the Oracle Auto Service Request (ASR) for Oracle Sun systems. Using auto-case generation improves system availability and expedites the diagnostic process when specific hardware faults occur.

Once the system has been registered with a valid My Oracle Support (MOS) account the [smf\(5\)](#) service, `asr-notify`, will start to send HTTPS/XML telemetry either directly to the configured endpoint or through an optional HTTP proxy.

The `asradm` utility can also list the current registration state as seen from the registered server. This only confirms that messages are being sent to the Oracle ASR service and does not mean the system is actively being monitored. The user must log into their MOS account and activate the service for the system. An email should be sent to the user after registration that will describe the specific details needed to complete enrollment.

ASR messages can be sent manually using the `send` subcommand. The content of ASR messages can also be viewed without actually sending the messages as well. The generation of messages can be done even if the system has not been registered.

The system can also be unregistered from the ASR service, which will remove the system from being monitored and will disable all telemetry sent by the system.

选项

The following options are supported:

`-e endpoint_URL`

Sets the endpoint URL used for registration and all message telemetry. The default value is `transport.sun.com` sends all telemetry directly to Oracle service. A different URL can be used to support a local instance of an ASR Manager solution. The local ASR Manager can be used to aggregate telemetry from many hosts instances.

`-h host[:port]`

Sets the HTTPS proxy host and optional port number to use for connecting to the internet. If a port is not specified, the default port value of 80 is used.

`-i`

Sets the HTTPS connection to be a direct internet connection and thus not use any proxy host. This option will clear any previously set HTTPS proxy information.

- n
Do a dry run of sending an event, which displays the message data to stdout that would have been sent.
- u *username*
When used with the `register` subcommand, this specifies the MOS user name to be associated with registrations of products on this system. When used with the `set - proxy` subcommand this specifies the HTTPS proxy user name used for message transport.
- p *password-file*
This should be a single-line file containing a password value. It can be immediately removed after running this command. When used with the `register` subcommand, this specifies the password associated with the MOS user name. When used with the `set - proxy` subcommand, this option specifies the password associated with the HTTPS proxy used to connect to the internet. If the `- p` option is not entered and the `- u` option was set, the user will be prompted for the password.

子命令

The `asradm` subcommands are described as follows.

`register`

Registers the system with MOS using the supplied authentication credentials. The MOS password will be used only for initial registration, to obtain a token used for all future telemetry and will not be stored anywhere on the system.

If the host system is behind a firewall, then the HTTP proxy settings must be set using the `set-proxy` command before registering the system.

Once registered, an ASR activation message will be sent that will request support for automated support call generation. Subsequently, an email will be sent to the registered user giving the status of the service request.

`list`

Shows the current authenticated MOS user name and the network connectivity information needed for HTTPS communication with MOS. If no registration has yet been done (by means of the `register` command), then the status of `Unregistered` is displayed.

`set-proxy`

Sets up the HTTPS connection information to be used for sending all ASR messages. Either a direct connection can be made or an HTTPS proxy can be defined.

`unregister`

Sends an ASR deactivation event and removes all configured registration information. No further telemetry will be sent and the Oracle ASR service will no longer generate any automated support calls.

`send`

Manually sends a specified ASR message to the Oracle ASR service or, with the dry run option (`- n`), display ASR messages.

- 示例

示例 1 Setup Internet Connection to Use an HTTPS Proxy

The following command will route all messages through an HTTPS proxy host `webproxy.example.com` on port 8080.

asradm set-proxy -h webproxy.example.com:8080
- 示例 2 Registering an MOS ID

The following command is used to interactively authenticate and register this system with the given MOS ID for use with ASR. Following this, you will be prompted for your support user name and password.

asradm register
- 示例 3 Authenticating Non-interactively

This is similar to `set-proxy` example, above. The difference is that the MOS user name and password are specified by means of the command line.

asradm register -u joe.admin@example.com -p mypassword
- 示例 4 Viewing Contents of Audit Message

The command below will display an audit message without sending the event. This will work even if the ASR service has not been registered.

asradm send -n audit

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/fault-management/asr-notify
Interface Stability	Uncommitted

另请参见

[svcs\(1\)](#), [asr-notify\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [syslogd\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

引用名	asr-notify – auto service request notification daemon for Fault Management events
用法概要	<code>/usr/lib/fm/notify/asr-notify [-]</code>
描述	<code>asr-notify</code> is a daemon that subscribes to Fault Management life cycle events and produces HTTPS/XML notifications based on a set of notification preferences that are stored in the SMF service configuration repository. The messages will be sent to the Oracle Auto Service Request (ASR) service offering once the service has been registered using the asradm(1M) command or by setting SMF auto registration properties. If the auto registration properties are set, the service will attempt to register on either startup or refresh. If registration fails, the service will go into maintenance mode. If registration is successful, the SMF auto registration properties will be deleted. <code>asr-notify</code> is managed by the service management facility, smf(5), under the service FMRI: <pre>svc:/system/fm/asr-notify:default</pre> Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M). The service's status can be queried using the svcs(1) command. ASR notification preferences are set using svccfg(1M).
Properties	The following service properties can be set: <i>config/rootdir</i> This is an <code>string</code> property that defaults to <code>/</code>. When set, the specified root directory will be used for all pathnames evaluated by <code>asr-notify</code>. <i>config/http_timeout</i> This is an <code>integer</code> property that sets the number of seconds to wait for HTTP connections before generating an error. <i>config/endpoint</i> This is an <code>string</code> property that sets the endpoint URL used for registration and all message telemetry. The endpoint URL can be set to a local ASR Manager that will aggregate telemetry from many hosts instances. If not set, then all telemetry will be sent directly to the Oracle Auto Service Request service. <i>autoreg/user</i> This is an <code>string</code> property that is used to set the My Oracle Support user name used for auto registration. <i>autoreg/password</i> This is an <code>string</code> property that is used to set the My Oracle Support password used for auto registration. <i>autoreg/proxy-host</i> This is an <code>string</code> property that is used to set the HTTP proxy <i>hostname:port</i> and needs to be set only if an HTTP proxy is used to connect to the internet.

autoreg/proxy-user

This is an astring property that is used to set the HTTP proxy user name and needs to be set only if the HTTP proxy requires credentials.

autoreg/proxy-password

This is an astring property that is used to set the HTTP proxy password and needs to be set only if the HTTP proxy requires credentials.

Messages

The service will send the following types of messages:

- fault messages when an FMA fault is created
- daily heartbeat messages to indicate that the system is still active
- running and weekly audit messages to provide a full inventory of operational components

The contents of each message are an XML document that adheres to the ASR message XML schema. Each message starts with a common message header followed by event specific content.

ASR event header information:

XML Element	Description
-----	-----
site-id	the serial number of the host sending the event
host-id	the hostname(1) of the system that the message is about
message-uuid	a unique id generated for each message
message-time	local time on the system that generated the message
system-id	the serial number of the system that the message is about
asset-id	the registered service tag for the system
product-id	the unique product identifier for the system
product-name	the name of the product for the system
event	element containing fault specific event information

The fault-specific content in the event will be contained within the primary-event-information XML element.

XML Element	Description
-----	-----
message-id	the knowledge article message ID
event-uuid	the FMA event UUID
event-time	the time that the FMA event was generated
severity	the severity of the FMA event
component	element containing list of suspects
summary	a short summary of the event

description event description
payload the raw FMA event contents

For each suspect FRU associated with a fault event, a hardware-component element will be included within the component.

XML Element	Description
-----	-----
name	the component path name of the FRU
serial	the FRU serial number
part	the FRU part number
revision	the FRU revision level
additional-information	a list of additional properties containing values such as FRU manufacturer and FRU model

Some faults might be associated with a software issue instead of a hardware issue. In this case, the fault event will contain a software-module element within the component list.

XML Element	Description
-----	-----
name	the name or FMRI of the software module in error
description	a description of the software module error

Heartbeat events do not contain any additional data and audit events contain a list of all hardware-component, software-package, and software-module elements available to the system.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/fault-management/asr-notify
Interface Stability	Uncommitted

另请参见

[svcs\(1\)](#), [asradm\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [syslogd\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

引用名

用法概要

描述

选项

退出状态

文件

属性

atohexlabel – convert a human readable label to its internal text equivalent

/usr/sbin/atohexlabel [human-readable-sensitivity-label]

/usr/sbin/atohexlabel -c [human-readable-clearance]

atohexlabel converts a human readable label into an internal text representation that is safe for storing in a public object. If no option is supplied, the label is assumed to be a sensitivity label.

Internal conversions can later be parsed to their same value. This internal form is often hexadecimal. The converted label is written to the standard output file. If no human readable label is specified, the label is read from the standard input file. The expected use of this command is emergency repair of labels that are stored in internal databases.

-c Identifies the human readable label as a clearance.

The following exit values are returned:

0 On success.

1 On failure, and writes diagnostics to the standard error file.

/etc/security/tsol/label_encodings

The label encodings file contains the classification names, words, constraints, and values for the defined labels of this system.

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/trusted
Interface Stability	See below.

The command output is Committed for systems with the same label_encodings file. The command invocation is Committed for systems that implement the DIA MAC policy.

[hextoalabel\(1M\)](#), [label_to_str\(3TSOL\)](#), [str_to_label\(3TSOL\)](#), [label_encodings\(4\)](#), [attributes\(5\)](#)

《Trusted Extensions 配置和管理》

附注

The functionality described on this manual page is available only if the system is configured with Trusted Extensions.

This file is part of the Defense Intelligence Agency (DIA) Mandatory Access Control (MAC) policy. This file might not be applicable to other MAC policies that might be developed for future releases of Solaris Trusted Extensions software.

引用名	audit – control the behavior of the audit daemon
用法概要	audit -n -s -t -v
描述	The <code>audit</code> command is the system administrator's interface to start, stop, and refresh the audit service auditd(1M) . Refreshing the audit service rereads the system and plugin configuration.
选项	-n Notify the audit service <code>audit_binfile</code> plugin to close the current audit file and open a new audit file in the current audit directory. -s Start (enable) the audit service if it is not running, or refresh the audit service, if it is currently running. -t Terminate (disable) the audit service. The audit service will close out the active plugins, stop auditing and exit. Use -s to restart auditing. -v Verify that at least one plugin is active. Verify the directory parameters of audit_binfile(5).
诊断	The <code>audit</code> command will exit with 0 upon success and a positive integer upon failure.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [auditconfig\(1M\)](#), [auditd\(1M\)](#), [attributes\(5\)](#), [audit_binfile\(5\)](#)

See the section on Solaris Auditing in 《Oracle Solaris 管理：安全服务》.

附注 The `audit` command does not modify a process's preselection mask. Its functions are limited to performing control actions of the auditing subsystem. See [auditconfig\(1M\)](#) for configuration.

The -s option validates the audit plugin configuration. If it is not valid an error message is displayed and the audit service is not started or refreshed. The -v option may be used to validate the audit plugin configuration before using the -s option to start or refresh the audit service.

The -s option also checks state of the audit service. In case the audit service is found in the maintenance state (thus not able to be enabled or refreshed) the `audit` command returns with an appropriate message and exit code.

All options are valid in the global zone. Unless per-zone is enabled, only the -v option is valid in a local zone. See [auditconfig\(1M\)](#) for per-zone audit configuration.

引用名 auditconfig – configure auditing

用法概要 auditconfig *subcommand*...

描述 auditconfig provides a command line interface to get and set kernel audit parameters.

Except for getting or setting the persistent audit service values, this functionality is available only if the Solaris Auditing feature has been enabled.

A zero (0) queue value indicates that the system default is in effect.

The setting of the perzone policy determines the scope of the audit setting controlled by auditconfig. If perzone is set, then the values reflect the local zone except as noted. Otherwise, the settings are for the entire system. Any restriction based on the perzone setting is noted for each option to which it applies.

A non-global zone administrator can set all audit policy options except perzone and ahlt. perzone and ahlt apply only to the global zone; setting these policies requires the privileges of a global zone administrator. perzone and ahlt are described under the -setpolicy option, below.

This command is available to administrators who have been granted the Audit Control Rights Profile.

选项 The following option is supported:

-t
Display or set the values on the running system in addition to the persistent values of the audit service.

This option is available only for the subcommands that list it below.

子命令

-aconf
Set the configured non-attributable audit mask. For example:

```
# auditconfig -aconf  
Configured non-attributable event mask.
```

-audit event *sorffretval string*
This command constructs an audit record for audit event *event* using the process's audit characteristics containing a text token *string*. The return token is constructed from the *sorff* (success/failure flag) and the *retval* (return value). The event is type char*, the *sorff* is 0/1 for success/failure, *retval* is an errno value, *string* is type *char. This command is useful for constructing an audit record with a shell script. An example of this option:

```
# auditconfig -audit AUE_ftpd 0 0 "test string"  
#
```

```
audit record from audit trail:  
header,76,2,ftp access,,Fri Dec 08 08:44:02 2000, + 669 msec
```

```
subject,abc,root,other,root,other,104449,102336,235 197121 elbow
text,test string
return,success,0
```

-chkaconf

Checks the configuration of the non-attributable events set in the kernel against the entries configured in the audit service (-setnaflags). If the active class mask of a kernel audit event does not match the configured class mask, a mismatch is reported.

-chkconf

Check the configuration of kernel audit event to class mappings. If the runtime class mask of a kernel audit event does not match the configured class mask, a mismatch is reported.

-conf

Configure kernel audit event to class mappings. Runtime class mappings are changed to match those in the audit event to class database file.

-getasid

Prints the audit session ID of the current process. For example:

```
# auditconfig -getasid
audit session id = 102336
```

-getaudit

Returns the audit characteristics of the current process.

```
# auditconfig -getaudit
audit id = abc(666)
process preselection mask = lo(0x1000,0x1000)
terminal id (maj,min,host) = 235,197121,elbow(172.146.89.77)
audit session id = 102336
```

-getaid

Prints the audit ID of the current process. For example:

```
# auditconfig -getaid
audit id = abc(666)
```

-getcar

Prints current active root location (anchored from root [or local zone root] at system boot). For example:

```
# auditconfig -getcar
current active root = /
```

-getclass *event*

Display the preselection mask associated with the specified kernel audit event. *event* is the kernel event number or event name.

-getcond

Display the kernel audit condition. The condition displayed is the literal string `auditing` meaning auditing is enabled and turned on (the kernel audit module is constructing and

queuing audit records, audit daemon is running); `noaudit`, meaning auditing is enabled but turned off (the kernel audit module is not constructing and queuing audit records, audit daemon is not running); `disabled`, meaning that the audit module has not been enabled (the module has been excluded in [system\(4\)](#)). See [auditd\(1M\)](#) for further information.

-getestate *event*

For the specified event (string or event number), print out classes *event* has been assigned. For example:

```
# auditconfig -getestate 20
audit class mask for event AUE_REBOOT(20) = 0x800
# auditconfig -getestate AUE_RENAME
audit class mask for event AUE_RENAME(42) = 0x30
```

-getflags

Display the user default audit preselection flags.

-getkaudit

Get audit characteristics of the current zone. For example:

```
# auditconfig -getkaudit
audit id = unknown(-2)
process preselection mask = lo,na(0x1400,0x1400)
terminal id (maj,min,host) = 0,0,(0.0.0.0)
audit session id = 0
```

If the audit policy `perzone` is not set, the terminal id is that of the global zone. Otherwise, it is the terminal id of the local zone.

-getmask

Get non-attributable pre-selection mask for the current zone. For example:

```
# auditconfig -getmask
audit flags for non-attributable events = lo,na(0x1400,0x1400)
```

If the audit policy `perzone` is not set, the kernel mask is that of the global zone. Otherwise, it is that of the local zone.

-getnaflags

Display the non-attributable audit flags.

-getpinfo *pid*

Display the audit ID, preselection mask, terminal ID, and audit session ID for the specified process.

-getplugin [*name*]

Display information about the plugin name. If *name* is not specified, display all plugins.

[-t] -getpolicy

Display the kernel audit policy. The `ahlt` and `perzone` policies reflect the settings from the global zone. If `perzone` is set, all other policies reflect the local zone's settings. If `perzone` is not set, the policies are machine-wide.

-getcwd

Prints current working directory (anchored from zone root at system boot). For example:

```
# cd /usr/tmp
# auditconfig -getcwd
current working directory = /var/tmp
```

[-t] -getqbufsz

Get audit queue write buffer size. For example:

```
# auditconfig -getqbufsz
no configured audit queue size
audit queue buffer size (bytes) = 1024
```

[-t] -getqctrl

Get audit queue write buffer size, audit queue hiwater mark, audit queue lowater mark, audit queue prod interval (ticks).

```
# auditconfig -getqctrl
no configured audit queue lowater mark
no configured audit queue hiwater mark
no configured audit queue size
no configured audit queue delay
audit queue hiwater mark (records) = 100
audit queue lowater mark (records) = 10
audit queue buffer size (bytes) = 1024
audit queue delay (ticks) = 20

# auditconfig -setqbufsz 8192
# auditconfig -t -setqbufsz 12288
# auditconfig -setqdelay 20
# auditconfig -t -setqdelay 25
# auditconfig -getqctrl
no configured audit queue lowater mark
no configured audit queue hiwater mark
configured audit queue buffer size (bytes) = 8192
configured audit queue delay (ticks) = 20
active audit queue hiwater mark (records) =    100
active audit queue lowater mark (records) =     10
active audit queue buffer size (bytes) = 12288
active audit queue delay (ticks) = 25
```

[-t] -getqdelay

Get interval at which audit queue is prodded to start output. For example:

```
# auditconfig -getqdelay
no configured audit queue delay
audit queue delay (ticks) = 20
```

[-t] -getqhiwater

Get high water point in undelivered audit records when audit generation will block. For example:

```
# ./auditconfig -getqhiwater
no configured audit queue hiwater mark
audit queue hiwater mark (records) = 100
```

[-t] -getqlowater

Get low water point in undelivered audit records where blocked processes will resume. For example:

```
# auditconfig -getqlowater
no configured audit queue lowater mark
audit queue lowater mark (records) = 10
```

-getstat

Print current audit statistics information. For example:

```
# auditconfig -getstat
gen nona kern aud  ctl  enq wrtn wblk rblk drop  tot  mem
910   1  725  184   0  910  910   0  231   0   88   48
```

See [auditstat\(1M\)](#) for a description of the headings in `-getstat` output.

-gettid

Print audit terminal ID for current process. For example:

```
# auditconfig -gettid
terminal id (maj,min,host) = 235,197121,elbow(172.146.89.77)
```

-lsevent

Display the currently configured (runtime) kernel and user level audit event information.

-lspolicy

Display the kernel audit policies with a description of each policy.

-setasid *session-ID* [*cmd*]

Execute shell or *cmd* with specified *session-ID*. For example:

```
# ./auditconfig -setasid 2000 /bin/ksh
#
# ./auditconfig -getpinfo 104485
audit id = abc(666)
process preselection mask = lo(0x1000,0x1000)
terminal id (maj,min,host) = 235,197121,elbow(172.146.89.77)
audit session id = 2000
```

- setaudit *audit-ID preselect_flags term-ID session-ID [cmd]*
Execute shell or *cmd* with the specified audit characteristics.
- setaudit *audit-ID [cmd]*
Execute shell or *cmd* with the specified *audit-ID*.
- setclass *event audit_flag[,audit_flag...]*
Map the kernel event *event* to the classes specified by *audit_flag* list. *event* is an event number or name. An *audit_flag* is a character string representing an audit class. See [audit_flags\(5\)](#) for further information. If *perzone* is not set, this option is valid only in the global zone.
- setflags *audit_flags*
Set the default user audit preselection flags; see [audit_flags\(5\)](#). The default preselection flags are combined with the user's specific audit flags to form the user's audit preselection mask.
- setkaudit *IP-address_type IP_address*
Set IP address of machine to specified values. *IP-address_type* is *ipv6* or *ipv4*.

If *perzone* is not set, this option is valid only in the global zone.
- setkmask *audit_flags*
Set non-attributable preselection flags of machine.

If *perzone* is not set, this option is valid only in the global zone.
- setnaflags *audit_flags*
Set the non-attributable audit flags; see [audit_flags\(5\)](#). Non-attributable audit flags define which classes of events are to be audited when the action cannot be attributed to an authenticated user. Failed login is an example of an event that is non-attributable.
- setpmask *pid flags*
Set the preselection mask of the specified process. *flags* is the ASCII representation of the flags similar to that in [audit_flags\(5\)](#).

If *perzone* is not set, this option is valid only in the global zone.
- setplugin *name active|inactive [attributes [qsize]]*
Configure the plugin *name* to be *active* or *inactive*. Optionally configure the *attributes* and number of unprocessed audit records to queue for the plugin. See the relevant audit plugin man pages and [auditd\(1M\)](#).
- [-t] -setpolicy [+|-]*policy_flag[,policy_flag...]*
Set the kernel audit policy. A policy *policy_flag* is literal strings that denotes an audit policy. A prefix of + adds the policies specified to the current audit policies. A prefix of - removes the policies specified from the current audit policies. No policies can be set from a local zone unless the *perzone* policy is first set from the global zone. The following are the valid policy flag strings (`auditconfig -lspolicy` also lists the current valid audit policy flag strings):

all	Include all policies that apply to the current zone.
ahlt	Panic is called and the system dumps core if an asynchronous audit event occurs that cannot be delivered because the audit queue has reached the high-water mark or because there are insufficient resources to construct an audit record. By default, records are dropped and a count is kept of the number of dropped records.
arge	Include the execv(2) system call environment arguments to the audit record. This information is not included by default.
argv	Include the execv(2) system call parameter arguments to the audit record. This information is not included by default.
cnt	Do not suspend processes when audit resources are exhausted. Instead, drop audit records and keep a count of the number of records dropped. By default, process are suspended until audit resources become available.
group	Include the supplementary group token in audit records. By default, the group token is not included.
none	Include no policies. If used in other than the global zone, the ahl t and perzone policies are not changed.
path	Add secondary path tokens to audit record. These are typically the pathnames of dynamically linked shared libraries or command interpreters for shell scripts. By default, they are not included.
perzone	Maintain separate configuration, queues, and logs for each zone and execute a separate version of auditd(1M) for each zone.
public	Audit public files. By default, read-type operations are not audited for certain files which meet <i>public</i> characteristics: owned by root, readable by all, and not writable by all.
trail	Include the trailer token in every audit record. By default, the trailer token is not included.
seq	Include the sequence token as part of every audit record. By default, the sequence token is not included. The sequence token attaches a sequence number to every audit record.
windata_down	Include in an audit record any downgraded data moved between windows. This policy is available only if the system is configured with Trusted Extensions. By default, this information is not included.
windata_up	Include in an audit record any upgraded data moved between windows. This policy is available only if the system is configured with Trusted Extensions. By default, this information is not included.

- zonename Include the zonename token as part of every audit record. By default, the zonename token is not included. The zonename token gives the name of the zone from which the audit record was generated.
- [-t] -setqbufsz *buffer_size***
Set the audit queue write buffer size (bytes). Zero (0), indicates reset to no configured value.
- [-t] -setqctrl *hiwater lowater bufsz interval***
Set the audit queue write buffer size (bytes), hiwater audit record count, lowater audit record count, and wakeup interval (ticks). Valid within a local zone only if perzone is set. Zero (0), indicates reset to no configured value.
- [-t] -setqdelay *interval***
Set the audit queue wakeup interval (ticks). This determines the interval at which the kernel pokes the audit queue, to write audit records to the audit trail. Valid within a local zone only if perzone is set. Zero (0), indicates reset to no configured value.
- [-t] -setqhiwater *hiwater***
Set the number of undelivered audit records in the audit queue at which audit record generation blocks. Valid within a local zone only if perzone is set. Zero (0), indicates reset to no configured value.
- [-t] -setqlowater *lowater***
Set the number of undelivered audit records in the audit queue at which blocked auditing processes unblock. Valid within a local zone only if perzone is set. Zero (0), indicates reset to no configured value.
- setsmask *asid flags***
Set the preselection mask of all processes with the specified audit session ID. Valid within a local zone only if perzone is set.
- setstat**
Reset audit statistics counters. Valid within a local zone only if perzone is set.
- setumask *username|audit flags***
Set the preselection mask of all processes with the specified username or audit ID. Valid within a local zone only if perzone is set.

示例

示例1 Using auditconfig

The following is an example of an auditconfig command:

```
#
# map kernel audit event number 10 to the "fr" audit class
#
% auditconfig -setclass 10 fr

#
# turn on inclusion of exec arguments in exec audit records
#
```

示例1 Using auditconfig (续)

```
% auditconfig -setpolicy +argv
```

示例2 Setting Only the Number of Unprocessed Audit Records

The following sequence of commands sets only the number of unprocessed audit records to queue for the audit_binfile plugin:

```
#
# see if audit_binfile is active
#
% auditconfig -getplugin audit_binfile

#
# set to queue 20 unprocessed audit records
#
% auditconfig -setplugin active "" 20
```

示例3 Resetting Queue Control Parameters

The following commands reset active and configured queue control parameters:

```
#
# reset configured audit queue control parameters
#
% auditconfig -setqctrl 0 0 0 0

#
# reset active audit queue control parameters
#
% auditconfig -t -setqctrl 0 0 0 0
```

退出状态

0 Successful completion.

1 An error occurred.

文件

/etc/security/audit_event Stores event definitions used in the audit system.

/etc/security/audit_class Stores class definitions used in the audit system.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[audit\(1M\)](#), [auditd\(1M\)](#), [auditstat\(1M\)](#), [praudit\(1M\)](#), [execv\(2\)](#), [audit_class\(4\)](#), [audit_event\(4\)](#), [system\(4\)](#), [attributes\(5\)](#), [audit_binfile\(5\)](#), [audit_flags\(5\)](#), [audit_remote\(5\)](#), [audit_syslog\(5\)](#)

See the section on Solaris Auditing in [《Oracle Solaris 管理：安全服务》](#).

附注

If plugin output is selected using the `-setplugin` option, the behavior of the system with respect to the `-setpolicy +cnt` and the `-setqhiwater` options is modified slightly. If `-setpolicy +cnt` is set, data will continue to be sent to the selected plugin, even though output of the `audit_binary` plugin is stopped, pending the freeing of disk space. If `-setpolicy -cnt` is used, the blocking behavior is as described under SUBCOMMANDS, above. The queue high water mark value is used within `auditd` as the upper bound for its queue limits unless overridden by means of the `qsize` attribute, as described in the explanation of the `-setplugin` option, above.

The `auditconfig` options that modify or display process-based information are not affected by the `perzone` policy. Those that modify system audit data such as the terminal id and audit queue parameters are valid only in the global zone, unless the `perzone` policy is set. The display of a system audit reflects the local zone if `perzone` is set. Otherwise, it reflects the settings of the global zone.

The change to plugins (`-setplugin`) settings do not take effect (such as becoming active or inactive, or changing the active attribute or queue size values) until the audit service is refreshed. Use [audit\(1M\)](#) to refresh the audit service.

引用名 auditd – audit daemon

用法概要 /usr/sbin/auditd

描述 The audit daemon, `auditd`, controls the generation and location of audit trail files and the generation of `syslog` messages based on its configuration (see [auditconfig\(1M\)](#)). If auditing is enabled, `auditd` reads its configuration to do the following:

- Configure audit policy.
- Configure the audit queue control parameters.
- Configure the event-to-class mappings.
- Set the default audit masks.
- Load one or more plugins.

Solaris provides three plugins. [audit_binfile\(5\)](#) writes binary audit data to a file. [audit_remote\(5\)](#) sends binary audit data to an authenticated server with privacy and integrity protection. [audit_syslog\(5\)](#) sends text summaries of audit records to the `syslog` daemon.

- Read audit data from the kernel and pass that data to each of the active plugins.
- Execute the [audit_warn\(1M\)](#) script to warn of various conditions.

[audit\(1M\)](#) is used to control `auditd`. It can cause `auditd` to:

- Close the current audit file and open a new one.
- Start and refresh the service based on the current properties.
- Close the audit trail and terminate auditing.

[auditconfig\(1M\)](#) is used to configure `auditd`. It can configure the active and permanent:

- audit policy
- audit queue control parameters
- default audit masks
- which plugins are to be loaded
- plugin attributes

Audit Record Queue The maximum number of records to queue for audit data sent to the plugin is specified by the `qsize` parameter specified for the plugin. If omitted, the current `hiwater` mark is used. See the `-getqctrl` option in [auditconfig\(1M\)](#). When this maximum is reached, `auditd` will either block processes or discard data, depending on the `cnt` audit policy as described in [auditconfig\(1M\)](#).

Auditing Conditions The audit daemon and audit plugins invoke the program [audit_warn\(1M\)](#) under certain conditions. See [audit_warn\(1M\)](#) for more information.

文件

- `etc/security/audit/audit_class`
- `etc/security/audit/audit_event`

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[audit\(1M\)](#), [audit_warn\(1M\)](#), [auditconfig\(1M\)](#), [praudit\(1M\)](#), [audit_class\(4\)](#), [audit_class\(4\)](#), [audit_event\(4\)](#), [attributes\(5\)](#), [audit_binfile\(5\)](#), [audit_flags\(5\)](#), [audit_remote\(5\)](#), [audit_syslog\(5\)](#), [smf\(5\)](#)

See the section on Solaris Auditing in 《Oracle Solaris 管理：安全服务》.

附注

`auditd` is loaded in the global zone at boot time if auditing is enabled.

If the audit policy `perzone` is set, `auditd` runs in each zone, starting automatically when the local zone boots. If a zone is running when the `perzone` policy is set, auditing must be started manually in local zones. It is not necessary to reboot the system or the local zone to start auditing in a local zone. `auditd` can be started with `audit -s` and will start automatically with future boots of the zone.

When `auditd` runs in a local zone, the configuration is taken from the local zone's [smf\(5\)](#) repository and the `/etc/security` directory's files: `audit_class`, `user_attr`, and `audit_event`.

Configuration changes do not affect audit sessions that are currently running, as the changes do not modify a process's preselection mask. To change the preselection mask on a running process, use the `-setpmask` option of the `auditconfig` command (see [auditconfig\(1M\)](#)). If the user logs out and logs back in, the new configuration changes will be reflected in the next audit session.

The audit service FMRI is `svc:/system/auditd:default`.

引用名 auditrecord – display Solaris audit record formats

用法概要 /usr/sbin/auditrecord [-d] [[-a] | [-e *string*] | [-c *class*] |
 [-i *id*] | [-p *programname*] | [-s *systemcall*] | [-h]]

描述 The auditrecord utility displays the event ID, audit class and selection mask, and record format for audit record event types defined in [audit_event\(4\)](#). You can use auditrecord to generate a list of all audit record formats, or to select audit record formats based on event class, event name, generating program name, system call name, or event ID.

There are two output formats. The default format is intended for display in a terminal window; the optional HTML format is intended for viewing with a web browser.

Tokens contained in square brackets ([]) are optional and might not be present in every record.

选项 The following options are supported:

- a
List all audit records.
- c *class*
List all audit records selected by *class*. *class* is one of the two-character class codes from the file /etc/security/audit_class.
- d
Debug mode. Display number of audit records that are defined in audit_event, the number of classes defined in audit_class, any mismatches between the two files, and report which defined events do not have format information available to auditrecord.
- e *string*
List all audit records for which the event ID label contains the string *string*. The match is case insensitive.
- h
Generate the output in HTML format.
- i *id*
List the audit records having the numeric event ID *id*.
- p *programname*
List all audit records generated by the program *programname*, for example, audit records generated by a user-space program.
- s *systemcall*
List all audit records generated by the system call *systemcall*, for example, audit records generated by a system call.

The -p and -s options are different names for the same thing and are mutually exclusive. The -a option is ignored if any of -c, -e, -i, -p, or -s are given. Combinations of -c, -e, -i, and either -p or -s are ANDed together.

示例

示例 1 Displaying an Audit Record with a Specified Event ID

The following example shows how to display the contents of a specified audit record.

```
% auditrecord -i 6152
terminal login
program      /usr/sbin/login      see login(1)
              /usr/dt/bin/dtlogin See dtlogin
event ID     6152                 AUE_login
class        lo                  (0x00001000)
  header
  subject
  [text]
  return
                                error message
```

示例 2 Displaying an Audit Record with an Event ID Label that Contains a Specified String

The following example shows how to display the contents of a audit record with an event ID label that contains the string login.

```
# auditrecord -e login
terminal login
program      /usr/sbin/login      see login(1)
              /usr/dt/bin/dtlogin See dtlogin
event ID     6152                 AUE_login
class        lo                  (0x00001000)
  header
  subject
  [text]
  return
                                error message

rlogin
program      /usr/sbin/login      see login(1) - rlogin
event ID     6155                 AUE_rlogin
class        lo                  (0x00001000)
  header
  subject
  [text]
  return
                                error message
```

退出状态

```
0
  Successful operation

non-zero
  Error
```

文件

```
/etc/security/audit_class
  Provides the list of valid classes and the associated audit mask.
```

`/etc/security/audit_event`
Provides the numeric event ID, the literal event name, and the name of the associated system call or program.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
CSI	Enabled
Interface Stability	Uncommitted

另请参见 [auditconfig\(1M\)](#), [praudit\(1M\)](#), [audit.log\(4\)](#), [audit_class\(4\)](#), [audit_event\(4\)](#), [attributes\(5\)](#)

See the section on Solaris Auditing in 《Oracle Solaris 管理：安全服务》.

诊断 If unable to read either of its input files or to write its output file, `auditrecord` shows the name of the file on which it failed and exits with a non-zero return.

If no options are provided, if an invalid option is provided, or if both `-s` and `-p` are provided, an error message is displayed and `auditrecord` displays a usage message then exits with a non-zero return.

附注 This command was formerly known as `bsmrecord`.

If `/etc/security/audit_event` has been modified to add user-defined audit events, `auditrecord` displays the record format as undefined.

The audit records displayed by `auditrecord` are the core of the record that can be produced. Various audit policies and optional tokens, such as those shown below, might also be present.

The following is a list of [praudit\(1M\)](#) token names with their descriptions.

`group`

Present if the group audit policy is set.

`sensitivity label`

Present when Trusted Extensions is enabled and represents the label of the subject or object with which it is associated. The `mandatory_label` token is noted in the basic audit record where a label is explicitly part of the record.

`sequence`

Present when the seq audit policy is set.

`trailer`

Present when the trail audit policy is set.

zone

The name of the zone generating the record when the zonename audit policy is set. The zonename token is noted in the basic audit record where a zone name is explicitly part of the record.

引用名	auditreduce – merge and select audit records from audit trail files
用法概要	auditreduce [options] [audit-trail-file]...
描述	auditreduce allows you to select or merge records from audit trail files. Audit files can be from one or more machines. The merge function merges together audit records from one or more input audit trail files into a single output file. The records in an audit trail file are assumed to be sorted in chronological order (oldest first) and this order is maintained by auditreduce in the output file. Unless instructed otherwise, auditreduce will merge the entire audit trail, which consists of all the audit trail files in the directory structure <i>audit_root_dir</i>/* (see audit.log(4) for details of the structure of the audit root). Unless specified with the -R or -S option, <i>audit_root_dir</i> defaults to /var/audit. By using the file selection options it is possible to select some subset of these files, or files from another directory, or files named explicitly on the command line. The select function allows audit records to be selected on the basis of numerous criteria relating to the record's content (see audit.log(4) for details of record content). A record must meet all of the <i>record-selection-option</i> criteria to be selected.
Audit Trail Filename Format	Any audit trail file not named on the command line must conform to the audit trail filename format. Files produced by the audit system already have this format. Output file names produced by auditreduce are in this format. It is: <i>start-time . end-time . suffix</i> where <i>start-time</i> is the 14-character timestamp of when the file was opened, <i>end-time</i> is the 14-character timestamp of when the file was closed, and <i>suffix</i> is the name of the machine which generated the audit trail file, or some other meaningful suffix (for example, <i>all</i>, if the file contains a combined group of records from many machines). The <i>end-time</i> can be the literal string <i>not_terminated</i>, to indicate that the file is still being written to by the audit system. Timestamps are of the form <i>yyyymmddhhmmss</i> (year, month, day, hour, minute, second). The timestamps are in Coordinated Universal Time (UTC).
选项	
File Selection Options	The file selection options indicate which files are to be processed and certain types of special treatment. -A All of the records from the input files will be selected regardless of their timestamp. This option effectively disables the -a, -b, and -d options. This is useful in preventing the loss of records if the -D option is used to delete the input files after they are processed. Note, however, that if a record is <i>not</i> selected due to another option, then -A will not override that.

-C

Only process complete files. Files whose filename *end-time* timestamp is not_terminated are not processed (such a file is currently being written to by the audit system). This is useful in preventing the loss of records if -D is used to delete the input files after they are processed. It does not apply to files specified on the command line.

-D *suffix*

Delete input files after they are read if the entire run is successful. If auditreduce detects an error while reading a file, then that file is not deleted. If -D is specified, -A, -C and -O are also implied. *suffix* is given to the -O option. This helps prevent the loss of audit records by ensuring that all of the records are written, only complete files are processed, and the records are written to a file before being deleted. Note that if both -D and -O are specified in the command line, the order of specification is significant. The *suffix* associated with the latter specification is in effect.

-M *machine*

Allows selection of records from files with *machine* as the filename suffix. If -M is not specified, all files are processed regardless of suffix. -M can also be used to allow selection of records from files that contain combined records from many machines and have a common suffix (such as `all`).

-N

Select objects in *new mode*. This flag is off by default, thus retaining backward compatibility. In the existing, *old mode*, specifying the -e, -f, -g, -r, or -u flags would select not only actions taken with those IDs, but also certain objects owned by those IDs. When running in *new mode*, only actions are selected. In order to select objects, the -o option must be used.

-O *suffix*

Direct output stream to a file in the current `audit_root_dir` with the indicated suffix. *suffix* can alternatively contain a full pathname, in which case the last component is taken as the suffix, ahead of which the timestamps will be placed, ahead of which the remainder of the pathname will be placed. If the -O option is not specified, the output is sent to the standard output. When auditreduce places timestamps in the filename, it uses the times of the first and last records in the merge as the *start-time* and *end-time*.

-Q

Quiet. Suppress notification about errors with input files.

-R *pathname*

Specify the pathname of an alternate audit root directory *audit_root_dir* to be *pathname*. Therefore, rather than using `/var/audit` by default, `pathname/*` will be examined instead.

注 – The root file system of any non-global zones must not be referenced with the -R option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See [zones\(5\)](#).

-S *specific_directory*

This option causes `auditreduce` to read audit trail files from a specific location (*specific_directory*). *specific_directory* is normally interpreted as the name of a subdirectory of the audit root, therefore `auditreduce` will look in *audit_root_dir/specific_directory* for the audit trail files. But if *specific_directory* contains any backslash characters (`/`), it is the name of a directory not necessarily contained in the audit root. In this case, *specific_directory* will be consulted. This option allows archived files to be manipulated easily, without requiring that they be physically located in a directory structure like that of `/var/audit`.

-V

Verbose. Display the name of each file as it is opened, and how many records total were written to the output stream.

Record Selection
Options

The record selection options listed below are used to indicate which records are written to the output file produced by `auditreduce`.

Multiple arguments of the same type are not permitted.

-a *date-time*

Select records that occurred at or after *date-time*. The *date-time* argument is described under Option Arguments, below. *date-time* is in local time. The `-a` and `-b` options can be used together to form a range.

-b *date-time*

Select records that occurred before *date-time*.

-c *audit-classes*

Select records by audit class. Records with events that are mapped to the audit classes specified by *audit-classes* are selected. Audit class names are defined in `audit_class(4)`. The *audit-classes* can be a comma separated list of `audit_flags` like those described in `audit_flags(5)`. Using the `audit_flags`, one can select records based upon success and failure criteria.

-d *date-time*

Select records that occurred on a specific day (a 24-hour period beginning at 00:00:00 of the day specified and ending at 23:59:59). The day specified is in local time. The time portion of the argument, if supplied, is ignored. Any records with timestamps during that day are selected. If any hours, minutes, or seconds are given in *time*, they are ignored. `-d` can not be used with `-a` or `-b`.

-e *effective-user*

Select records with the specified *effective-user*.

-f *effective-group*

Select records with the specified *effective-group*.

-g *real-group*

Select records with the specified *real-group*.

-j *subject-ID*

Select records with the specified *subject-ID* where *subject-ID* is a process ID.

-l *label*

Select records with the specified label (or label range), as explained under “Option Arguments,” below. This option is available only if the system is configured with Trusted Extensions.

-m *event*

Select records with the indicated *event*. The *event* is the literal string or the *event* number.

-o *object_type=objectID_value*

Select records by object type. A match occurs when the record contains the information describing the specified *object_type* and the object ID equals the value specified by *objectID_value*. The allowable object types and values are as follows:

file=pathname

Select records containing file system objects with the specified pathname, where pathname is a comma separated list of regular expressions. If a regular expression is preceded by a tilde (~), files matching the expression are excluded from the output. For example, the option *file=~ /usr/openwin, /usr, /etc* would select all files in /usr or /etc except those in /usr/openwin. The order of the regular expressions is important because auditreduce processes them from left to right, and stops when a file is known to be either selected or excluded. Thus the option *file= /usr, /etc, ~ /usr/openwin* would select all files in /usr and all files in /etc. Files in /usr/openwin are not excluded because the regular expression /usr is matched first. Care should be given in surrounding the *pathname* with quotes so as to prevent the shell from expanding any tildes.

filegroup=group

Select records containing file system objects with *group* as the owning group.

fileowner=user

Select records containing file system objects with *user* as the owning user.

fmri=service_instance

Select records containing fault management resource identifier (FMRI) objects with the specified *service instance*. See [smf\(5\)](#).

msgqid=ID

Select records containing message queue objects with the specified *ID* where *ID* is a message queue ID.

msgqgroup=group

Select records containing message queue objects with *group* as the owning or creating group.

msgqowner=user

Select records containing message queue objects with *user* as the owning or creating user.

pid=ID

Select records containing process objects with the specified *ID* where *ID* is a process ID. Process are objects when they are receivers of signals.

procgroup=group

Select records containing process objects with *group* as the real or effective group.

procowner=user

Select records containing process objects with *user* as the real or effective user.

semid=ID

Select records containing semaphore objects with the specified *ID* where *ID* is a semaphore ID.

semgroup=group

Select records containing semaphore objects with *group* as the owning or creating group.

semowner=user

Select records containing semaphore objects with *user* as the owning or creating user.

shmid=ID

Select records containing shared memory objects with the specified *ID* where *ID* is a shared memory ID.

shmgroup=group

Select records containing shared memory objects with *group* as the owning or creating group.

shmowner=user

Select records containing shared memory objects with *user* as the owning or creating user.

sock=port_number|machine

Select records containing socket objects with the specified *port_number* or the specified *machine* where *machine* is a machine name as defined in [hosts\(4\)](#).

user=user_name

Select records containing the user object whose name is specified. User objects are generally specified for administrative actions on a user.

-r real-user

Select records with the specified *real-user*.

-s session-id

Select audit records with the specified *session-id*.

-u audit-user

Select records with the specified *audit-user*.

-z zone-name

Select records from the specified zone name. The zone name selection is case-sensitive.

When one or more *filename* arguments appear on the command line, only the named files are processed. Files specified in this way need not conform to the audit trail filename format. However, -M, -S, and -R must not be used when processing named files. If the *filename* is "-" then the input is taken from the standard input.

Option Arguments

audit-trail-file

An audit trail file as defined in [audit.log\(4\)](#). An audit trail file not named on the command line must conform to the audit trail file name format. Audit trail files produced as output of `auditreduce` are in this format as well. The format is:

start-time . end-time . suffix

start-time is the 14 character time stamp denoting when the file was opened. *end-time* is the 14 character time stamp denoting when the file was closed. *end-time* can also be the literal string `not_terminated`, indicating the file is still be written to by the audit daemon or the file was not closed properly (a system crash or abrupt halt occurred). *suffix* is the name of the machine that generated the audit trail file (or some other meaningful suffix; for example, `all` would be a good suffix if the audit trail file contains a combined group of records from many machines).

date-time

The *date-time* argument to -a, -b, and -d can be of two forms: An absolute *date-time* takes the form:

yyyymmdd [*hh* [*mm* [*ss*]]]

where *yyyy* specifies a year (with 1970 as the earliest value), *mm* is the month (01-12), *dd* is the day (01-31), *hh* is the hour (00-23), *mm* is the minute (00-59), and *ss* is the second (00-59). The default is 00 for *hh*, *mm* and *ss*.

An offset can be specified as: `+n d|h|m|s` where *n* is a number of units, and the tags *d*, *h*, *m*, and *s* stand for days, hours, minutes and seconds, respectively. An offset is relative to the starting time. Thus, this form can only be used with the -b option.

event

The literal string or ordinal event number as found in [audit_event\(4\)](#). If *event* is not found in the `audit_event` file it is considered invalid.

group

The literal string or ordinal group ID number as found in [group\(4\)](#). If *group* is not found in the group file it is considered invalid. *group* can be negative.

label

The literal string representation of a MAC label or a range of two valid MAC labels. To specify a range, use `x;y` where *x* and *y* are valid MAC labels. Only those records that are fully bounded by *x* and *y* will be selected. If *x* or *y* is omitted, the default uses `ADMIN_LOW` or `ADMIN_HIGH` respectively. Notice that quotes must be used when specifying a range.

pathname

A regular expression describing a pathname.

user

The literal username or ordinal user ID number as found in [passwd\(4\)](#). If the username is not found in the `passwd` file it is considered invalid. *user* can be negative.

示例

示例1 Using `auditreduce`

`praudit(1M)` is available to display audit records in a human-readable form.

This will display the entire audit trail in a human-readable form:

```
% auditreduce | praudit
```

If all the audit trail files are being combined into one large file, then deleting the original files could be desirable to prevent the records from appearing twice:

```
% auditreduce -V -D /var/audit/combined/all
```

This displays what user `milner` did on April 13, 1988. The output is displayed in a human-readable form to the standard output:

```
% auditreduce -d 19880413 -u milner | praudit
```

The above example might produce a large volume of data if `milner` has been busy. Perhaps looking at only login and logout times would be simpler. The `-c` option will select records from a specified class:

```
% auditreduce -d 19880413 -u milner -c lo | praudit
```

To see `milner`'s login/logout activity for April 13, 14, and 15, the following is used. The results are saved to a file in the current working directory. Notice that the name of the output file will have `milnerlo` as the *suffix*, with the appropriate timestamp prefixes. Notice also that the long form of the name is used for the `-c` option:

```
% auditreduce -a 19880413 -b +3d -u milner -c login_logout -O milnerlo
```

To follow `milner`'s movement about the file system on April 13, 14, and 15 the `chdir` record types could be viewed. Notice that in order to get the same time range as the above example we needed to specify the `-b` time as the day *after* our range. This is because 19880416 defaults to midnight of that day, and records before that fall on 0415, the end-day of the range.

```
% auditreduce -a 19880413 -b 19880416 -u milner -m AUE_CHDIR | praudit
```

In this example, the audit records are being collected in summary form (the login/logout records only). The records are being written to a summary file in a different directory than the normal audit root to prevent the selected records from existing twice in the audit root.

```
% auditreduce -d 19880330 -c lo -O /var/audit/audit_summary/logins
```


示例1 Using auditreduce (续)

If activity for user ID 9944 has been observed, but that user is not known to the system administrator, then the command in the following example searches the entire audit trail for any records generated by that user. `auditreduce` queries the system about the current validity of ID 9944 and displays a warning message if it is not currently active:

```
% auditreduce -O /var/audit/audit_suspect/user9944 -u 9944
```

To get an audit log of only the global zone:

```
% auditreduce -z global
```

`/var/audit/*` default location of audit trails, when stored

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	See below.

The command invocation is Stable. The binary file format is Stable. The binary file contents is Unstable.

[praudit\(1M\)](#), [audit.log\(4\)](#), [audit_class\(4\)](#), [group\(4\)](#), [hosts\(4\)](#), [passwd\(4\)](#), [attributes\(5\)](#), [audit_flags\(5\)](#), [smf\(5\)](#)

See the section on Solaris Auditing in 《Oracle Solaris 管理：安全服务》.

`auditreduce` displays error messages if there are command line errors and then exits. If there are fatal errors during the run, `auditreduce` displays an explanatory message and exits. In this case, the output file might be in an inconsistent state (no trailer or partially written record) and `auditreduce` displays a warning message before exiting. Successful invocation returns 0 and unsuccessful invocation returns 1.

Since `auditreduce` might be processing a large number of input files, it is possible that the machine-wide limit on open files will be exceeded. If this happens, `auditreduce` displays a message to that effect, give information on how many file there are, and exit.

If `auditreduce` displays a record's timestamp in a diagnostic message, that time is in local time. However, when filenames are displayed, their timestamps are in UTC.

Conjunction, disjunction, negation, and grouping of record selection options should be allowed.

文件
属性

另请参见

诊断

已知问题

附注

The -z option should be used only if the audit policy zonename is set. If there is no zonename token, then no records will be selected.

引用名

auditstat – display kernel audit statistics

用法概要

auditstat [-c count] [-h numlines] [-i interval] [-n]
[-T u | d] [-v]

描述

auditstat displays kernel audit statistics. The fields displayed are as follows:

aud	The total number of audit records processed by the userland audit.
ctl	This field is obsolete.
drop	The total number of audit records that have been dropped. Records are dropped according to the kernel audit policy. See auditconfig(1M) , AUDIT_CNT policy for details.
enq	The total number of audit records put on the kernel audit queue.
gen	The total number of audit records that have been constructed (not the number written).
kern	The total number of audit records produced by user processes (as a result of system calls).
mem	The total number of Kbytes of memory currently in use by the kernel audit module.
nona	The total number of non-attributable audit records that have been constructed. These are audit records that are not attributable to any particular user.
rb1k	The total number of times that the audit queue has blocked waiting to process audit data.
tot	The total number of Kbytes of audit data written to the audit trail.
wb1k	The total number of times that user processes blocked on the audit queue at the high water mark.
wrtn	The total number of audit records written. The difference between enq and wrtn is the number of outstanding audit records on the audit queue that have not been written.

选项

-c count	Display the statistics a total of count times. If count is equal to zero, statistics are displayed indefinitely. A time interval must be specified.
-h numlines	Display a header for every numlines of statistics printed. The default is to display the header every 20 lines. If numlines is equal to zero, the header is never displayed.
-i interval	Display the statistics every interval where interval is the number of seconds to sleep between each collection.
-n	Display the number of kernel audit events currently configured.
-T u d	Display a time stamp.

Specify u for a printed representation of the internal representation of time. See [time\(2\)](#). Specify d for standard date format. See [date\(1\)](#).

-v Display the version number of the kernel audit module software.

退出状态 auditstat returns 0 upon success and 1 upon failure.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [auditconfig\(1M\)](#), [praudit\(1M\)](#), [attributes\(5\)](#)

引用名	audit_warn – audit daemon warning script												
用法概要	<code>/etc/security/audit_warn [option [arguments]]</code>												
描述	The <code>audit_warn</code> utility processes warning or error messages from the audit daemon. When a problem is encountered, the audit daemon, auditd(1M) calls <code>audit_warn</code> with the appropriate arguments. The <i>option</i> argument specifies the error type. The system administrator can specify a list of mail recipients to be notified when an <code>audit_warn</code> situation arises by defining a mail alias called <code>audit_warn</code> in aliases(4). The users that make up the <code>audit_warn</code> alias are typically the <code>audit</code> and <code>root</code> users.												
选项	The following options are supported: <table> <tr> <td><code>allhard count</code></td><td>Indicates that the hard limit for all filesystems has been exceeded <i>count</i> times. The default action for this option is to send mail to the <code>audit_warn</code> alias only if the <i>count</i> is 1, and to write a message to the machine console every time. It is recommended that mail <i>not</i> be sent every time as this could result in a the saturation of the file system that contains the mail spool directory.</td></tr> <tr> <td><code>allsoft</code></td><td>Indicates that the soft limit for all filesystems has been exceeded. The default action for this option is to send mail to the <code>audit_warn</code> alias and to write a message to the machine console.</td></tr> <tr> <td><code>auditoff</code></td><td>Indicates that someone other than the audit daemon changed the system audit state to something other than <code>AUC_AUDITING</code>. The audit daemon will have exited in this case. The default action for this option is to send mail to the <code>audit_warn</code> alias and to write a message to the machine console.</td></tr> <tr> <td><code>ebusy</code></td><td>Indicates that the audit daemon is already running. The default action for this option is to send mail to the <code>audit_warn</code> alias and to write a message to the machine console.</td></tr> <tr> <td><code>hard filename</code></td><td>Indicates that the hard limit for the file has been exceeded. The default action for this option is to send mail to the <code>audit_warn</code> alias and to write a message to the machine console.</td></tr> <tr> <td><code>hostname</code></td><td>Indicates that the audit service could not find an IP address to associate with the local hostname. It has fallen back to using the “loopback” address. The audit trail might not</td></tr> </table>	<code>allhard count</code>	Indicates that the hard limit for all filesystems has been exceeded <i>count</i> times. The default action for this option is to send mail to the <code>audit_warn</code> alias only if the <i>count</i> is 1, and to write a message to the machine console every time. It is recommended that mail <i>not</i> be sent every time as this could result in a the saturation of the file system that contains the mail spool directory.	<code>allsoft</code>	Indicates that the soft limit for all filesystems has been exceeded. The default action for this option is to send mail to the <code>audit_warn</code> alias and to write a message to the machine console.	<code>auditoff</code>	Indicates that someone other than the audit daemon changed the system audit state to something other than <code>AUC_AUDITING</code> . The audit daemon will have exited in this case. The default action for this option is to send mail to the <code>audit_warn</code> alias and to write a message to the machine console.	<code>ebusy</code>	Indicates that the audit daemon is already running. The default action for this option is to send mail to the <code>audit_warn</code> alias and to write a message to the machine console.	<code>hard filename</code>	Indicates that the hard limit for the file has been exceeded. The default action for this option is to send mail to the <code>audit_warn</code> alias and to write a message to the machine console.	<code>hostname</code>	Indicates that the audit service could not find an IP address to associate with the local hostname. It has fallen back to using the “loopback” address. The audit trail might not
<code>allhard count</code>	Indicates that the hard limit for all filesystems has been exceeded <i>count</i> times. The default action for this option is to send mail to the <code>audit_warn</code> alias only if the <i>count</i> is 1, and to write a message to the machine console every time. It is recommended that mail <i>not</i> be sent every time as this could result in a the saturation of the file system that contains the mail spool directory.												
<code>allsoft</code>	Indicates that the soft limit for all filesystems has been exceeded. The default action for this option is to send mail to the <code>audit_warn</code> alias and to write a message to the machine console.												
<code>auditoff</code>	Indicates that someone other than the audit daemon changed the system audit state to something other than <code>AUC_AUDITING</code> . The audit daemon will have exited in this case. The default action for this option is to send mail to the <code>audit_warn</code> alias and to write a message to the machine console.												
<code>ebusy</code>	Indicates that the audit daemon is already running. The default action for this option is to send mail to the <code>audit_warn</code> alias and to write a message to the machine console.												
<code>hard filename</code>	Indicates that the hard limit for the file has been exceeded. The default action for this option is to send mail to the <code>audit_warn</code> alias and to write a message to the machine console.												
<code>hostname</code>	Indicates that the audit service could not find an IP address to associate with the local hostname. It has fallen back to using the “loopback” address. The audit trail might not												

translate properly. See `/var/audit/debug` for more information. The audit service can be refreshed (`audit -s`) to retry to find an IP address.

`nostart` Indicates that auditing could not be started. The default action for this option is to send mail to the `audit_warn` alias and to write a message to the machine console. Some administrators may prefer to modify `audit_warn` to reboot the system when this error occurs.

`plugin name error count text` Indicates that an error occurred during execution of the `auditd` plugin *name*. The default action for this option is to send mail to the `audit_warn` alias only if *count* is 1, and to write a message to the machine console every time. (Separate counts are kept for each error type.) It is recommended that mail not be sent every time as this could result in the saturation of the file system that contains the mail spool directory. The *text* field provides the detailed error message passed from the plug-in. The *error* field is one of the following strings:

- `load_error` Unable to load the plugin *name*.
- `sys_error` The plugin *name* is not executing due to a system error such as a lack of resources.
- `config_error` No plug-ins loaded (including the binary file plug-in, `audit_binfile(5)`) due to configuration errors (see the `-setplugin` option of the `auditconfig(1M)` command). The name string is `--`, to indicate that no plug-in name applies.
- `retry` The plugin *name* reports it has encountered a temporary failure. For example, the `audit_binfree.so` plugin uses `retry` to indicate that all directories are full.
- `no_memory` The plugin *name* reports a failure due to lack of memory.
- `invalid` The plugin *name* reports it received an invalid input.
- `failure` The plugin *name* has reported an error as described in *text*.

postsigterm	Indicates that an error occurred during the orderly shutdown of the audit daemon. The default action for this option is to send mail to the audit_warn alias and to write a message to the machine console.
soft filename	Indicates that the soft limit for filename has been exceeded. The default action for this option is to send mail to the audit_warn alias and to write a message to the machine console.
tmpfile	Indicates that the temporary audit file already exists indicating a fatal error. The default action for this option is to send mail to the audit_warn alias and to write a message to the machine console.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

The interface stability is Evolving. The file content is Unstable.

另请参见 [audit\(1M\)](#), [auditconfig\(1M\)](#), [auditd\(1M\)](#), [aliases\(4\)](#), [audit.log\(4\)](#), [attributes\(5\)](#)

See the section on Solaris Auditing in 《[Oracle Solaris 管理：安全服务](#)》.

附注 This functionality is available only if the Solaris Auditing feature has been enabled.

If the audit policy perzone is set, the /etc/security/audit_warn script for the local zone is used for notifications from the local zone's instance of auditd. If the perzone policy is not set, all auditd errors are generated by the global zone's copy of /etc/security/audit_warn.

引用名

automount – install automatic mount points

用法概要

```
/usr/sbin/automount [-t duration] [-v]
```

描述

The automount utility installs autofs mount points and associates an automount map with each mount point. It starts the [automountd\(1M\)](#) daemon if it finds any non-trivial entries in either local or distributed automount maps and if the daemon is not already running. The autofs file system monitors attempts to access directories within it and notifies the [automountd\(1M\)](#) daemon. The daemon uses the map to locate a file system, which it then mounts at the point of reference within the autofs file system. A map can be assigned to an autofs mount using an entry in the /etc/auto_master map or a direct map.

If the file system is not accessed within an appropriate interval (10 minutes by default), the automountd daemon unmounts the file system.

The file /etc/auto_master determines the locations of all autofs mount points. By default, this file contains three entries:

```
# Master map for automounter
#
+auto_master
/net          -hosts    -nosuid
/home         auto_home
```

The +auto_master entry is a reference to an external NIS master map. If one exists, then its entries are read as if they occurred in place of the +auto_master entry. The remaining entries in the master file specify a directory on which an autofs mount will be made followed by the automounter map to be associated with it. Optional mount options may be supplied as an optional third field in the each entry. These options are used for any entries in the map that do not specify mount options explicitly. The automount command is usually run without arguments. It compares the entries /etc/auto_master with the current list of autofs mounts in /etc/mnttab and adds, removes or updates autofs mounts to bring the /etc/mnttab up to date with the /etc/auto_master. At boot time it installs all autofs mounts from the master map. Subsequently, it may be run to install autofs mounts for new entries in the master map or the direct map, or to perform unmounts for entries that have been removed from these maps.

SMF Management

The automount service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/system/filesystem/autofs:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

Startup autofs SMF parameters for automount can be manipulated using [sharectl\(1M\)](#). Refer to [sharectl\(1M\)](#) for instructions for viewing and updating these parameters. Supported parameters are:

`timeout=num`
Specifies a duration, in seconds, that a file system is to remain mounted when not in use. The default value is 600 (10 minutes). Equivalent to the `-t` option in automount.

`automount_verbose=TRUE | FALSE`
Verbose mode. Causes you to be notified of non-critical events, such as `autofs` mounts and unmounts. The default value is FALSE. Equivalent to the `-v` option in automount.

Automount with Oracle Solaris Trusted Extensions
If a system is configured with Oracle Solaris Trusted Extensions, users have home directories at each label within their clearance. Therefore a home directory needs to be available in each corresponding labeled zone. Instead of using the `auto_home` map, a special map is automatically generated, using the zone's name as a suffix. By default the map contains the single entry:

`-fstype=lofs :/export/home/&`

When a home directory is referenced and the name does not match any other keys in the zone's `auto_home_zonename` map, it will match this loopback mount specification. If this loopback match occurs and the name corresponds to a valid user whose home directory does not exist in the zone, the directory is automatically created on behalf of the user.

It is also possible to share home directories in a zone, in read-only mode, with higher-level zones, using NFS. In this case, the higher-level zone needs to have an automap entry for each lower-level zone that is to be imported. A typical map entry for the public zone, to be interpreted in the internal zone, would be called `auto_home_public`, and would look like this:

`+auto_home_public`

`public-zone-IP-address:/export/home/&`

This automap entry would then be included in `/etc/auto_master`, as follows:

`/zone/public/home auto_home_public -nobrowse`

Users in higher-level zones can use the [updatehome\(1\)](#) utility to synchronize specific startup files using their minimum labeled zone as the source.

选项
The following options are supported:

- `-t duration` Specifies a *duration*, in seconds, that a file system is to remain mounted when not in use. The default is 10 minutes.
- `-v` Verbose mode. Notifies of `autofs` mounts, unmounts, or other non-essential information.

用法

Map Entry Format

A simple map entry (mapping) takes the form:

```
key [ -mount-options ] location . . .
```

where *key* is the full pathname of the directory to mount when used in a direct map, or the simple name of a subdirectory in an indirect map. *mount-options* is a comma-separated list of mount options, and *location* specifies a file system from which the directory may be mounted. In the case of a simple NFS mount, the options that can be used are as specified in [mount_nfs\(1M\)](#), and *location* takes the form:

```
host: pathname
```

host is the name of the host from which to mount the file system, and *pathname* is the absolute pathname of the directory to mount.

Options to other file systems are documented on the other `mount_*` reference manual pages, for example, [mount_nfs\(1M\)](#).

Replicated File Systems

Multiple *location* fields can be specified for replicated NFS file systems, in which case automount and the kernel will each try to use that information to increase availability. If the read-only flag is set in the map entry, automountd mounts a list of locations that the kernel may use, sorted by several criteria. Only locations available at mount time will be mounted, and thus be available to the kernel. When a server does not respond, the kernel will switch to an alternate server. The sort ordering of automount is used to determine how the next server is chosen. If the read-only flag is not set, automount will mount the best single location, chosen by the same sort ordering, and new servers will only be chosen when an unmount has been possible, and a remount is done. Servers on the same local subnet are given the strongest preference, and servers on the local net are given the second strongest preference. Among servers equally far away, response times will determine the order if no weighting factors (see below) are used.

If the list includes server locations using both the NFS Version 2 Protocol and the NFS Version 3 Protocol, automount will choose only a subset of the server locations on the list, so that all entries will be the same protocol. It will choose servers with the NFS Version 3 Protocol so long as an NFS Version 2 Protocol server on a local subnet will not be ignored. See the [《Oracle Solaris 管理：IP 服务》](#) for additional details.

If each *location* in the list shares the same *pathname* then a single *location* may be used with a comma-separated list of hostnames:

```
hostname,hostname . . . : pathname
```

Requests for a server may be weighted, with the weighting factor appended to the server name as an integer in parentheses. Servers without a weighting are assumed to have a value of zero (most likely to be selected). Progressively higher values decrease the chance of being selected. In the example,

```
man -ro alpha,bravo,charlie(1),delta(4) : /usr/man
```

hosts alpha and bravo have the highest priority; host delta has the lowest.

Server proximity takes priority in the selection process. In the example above, if the server delta is on the same network segment as the client, but the others are on different network segments, then delta will be selected; the weighting value is ignored. The weighting has effect only when selecting between servers with the same network proximity. The automounter always selects the localhost over other servers on the same network segment, regardless of weighting.

In cases where each server has a different export point, the weighting can still be applied. For example:

```
man -ro alpha:/usr/man  bravo,charlie(1):/usr/share/man
                        delta(3):/export/man
```

A mapping can be continued across input lines by escaping the NEWLINE with a backslash (\). Comments begin with a number sign (#) and end at the subsequent NEWLINE.

Map Key Substitution The ampersand (&) character is expanded to the value of the key field for the entry in which it occurs. In this case:

```
jane sparserver : /home/&
```

the & expands to jane.

Wildcard Key The asterisk (*) character, when supplied as the key field, is recognized as the catch-all entry. Such an entry will match any key not previously matched. For instance, if the following entry appeared in the indirect map for /config:

```
*                & : /export/config/&
```

this would allow automatic mounts in /config of any remote file system whose location could be specified as:

```
hostname : /export/config/hostname
```

Note that the wildcard key does not work in conjunction with the -browse option.

Variable Substitution Client specific variables can be used within an automount map. For instance, if \$HOST appeared within a map, automount would expand it to its current value for the client's host name. Supported variables are:

ARCH	The output of arch	The architecture name. For example, sun4 on a sun4u machine.
CPU	The output of uname -p	The processor type.

		For example, "sparc"
HOST	The output of <code>uname -n</code>	The host name. For example, myhost.
KARCH	The output of <code>arch -k</code> or <code>uname -m</code>	The kernel architecture name or machine hardware name. For example, sun4u.
OSNAME	The output of <code>uname -s</code>	The OS name. For example, "SunOS"
OSREL	The output of <code>uname -r</code>	The OS release name. For example "5.3"
OSVERS	The output of <code>uname -v</code>	The OS version. For example, "beta1.0"
NATISA	The output of <code>isainfo -n</code>	The native instruction set architecture for the system. For example, "sparcv9"
PLATFORM	The output of <code>uname -i</code>	The platform name. For example, SUNW, Sun-Fire-V240.

If a reference needs to be protected from affixed characters, you can surround the variable name with curly braces ({ }).

Multiple Mounts

A multiple mount entry takes the form:

key [-mount-options] [[mountpoint] [-mount-options] location. . .] . . .

The initial `/[mountpoint]` is optional for the first mount and mandatory for all subsequent mounts. The optional *mountpoint* is taken as a pathname relative to the directory named by key. If *mountpoint* is omitted in the first occurrence, a *mountpoint* of `/` (root) is implied.

Given an entry in the indirect map for `/src`

```
beta      -ro\
/          svr1,svr2:/export/src/beta  \
/1.0      svr1,svr2:/export/src/beta/1.0  \
/1.0/man  svr1,svr2:/export/src/beta/1.0/man
```

All offsets must exist on the server under beta. automount will automatically mount /src/beta, /src/beta/1.0, and /src/beta/1.0/man, as needed, from either svr1 or svr2, whichever host is nearest and responds first.

Other File System Types The automounter assumes NFS mounts as a default file system type. Other file system types can be described using the `fstype` mount option. Other mount options specific to this file system type can be combined with the `fstype` option. The location field must contain information specific to the file system type. If the location field begins with a slash, a colon character must be prepended, for instance, to mount a CD file system:

```
cdrom -fstype=hsfs,ro : /dev/sr0
```

or to perform an autofs mount:

```
src -fstype=autofs auto_src
```

Use this procedure only if you are not using Volume Manager.

See the NOTES section for information on option inheritance.

Indirect Maps An indirect map allows you to specify mappings for the subdirectories you wish to mount under the `directory` indicated on the command line. In an indirect map, each key consists of a simple name that refers to one or more file systems that are to be mounted as needed.

Direct Maps Entries in a direct map are associated directly with `autofs` mount points. Each *key* is the full pathname of an `autofs` mount point. The direct map as a whole is not associated with any single directory.

Direct maps are distinguished from indirect maps by the `/-` key. For example:

```
# Master map for automounter
#
+auto_master
/net      -hosts      -nosuid,nobrowse
/home     auto_home   -nobrowse
/-        auto_direct
```

Included Maps The contents of another map can be included within a map with an entry of the form *+mapname*

If *mapname* begins with a slash, it is assumed to be the pathname of a local file. Otherwise, the location of the map is determined by the policy of the name service switch according to the entry for the automounter in `/etc/nsswitch.conf`, such as

```
automount: files nis
```

If the name service is `files`, then the name is assumed to be that of a local file in `/etc`. If the key being searched for is not found in the included map, the search continues with the next entry.

Special Maps

There are three special maps available: `-hosts`, `-fedfs` and `-null`. The `-hosts` map is used with the `/net` directory and assumes that the map key is the hostname of an NFS server. The `automountd` daemon dynamically constructs a map entry from the server's list of exported file systems. References to a directory under `/net/hermes` will refer to the corresponding directory relative to `hermes` root.

The `-fedfs` map is used with the `/nfs4` directory and assumes that the map key is the DNS domain for which the domain root filesystem is needed. The `automountd` daemon looks up the domain root servers with a query equivalent to:

```
% nslookup -q=srv _nfs4._domainroot._tcp.domain
```

...and mounts `server-list:/.domainroot-domain` at `/nfs4/domain`.

This supports the pending IETF standard documented in:

```
http://datatracker.ietf.org/doc/\  
draft-ietf-nfsv4-federated-fs-dns-srv-namespace/
```

The `-null` map cancels a previous map for the directory indicated. This is most useful in the `/etc/auto_master` for cancelling entries that would otherwise be inherited from the `+auto_master` include entry. To be effective, the `-null` entries must be inserted before the included map entry.

Executable Maps

Local maps that have the execute bit set in their file permissions will be executed by the automounter and provided with a key to be looked up as an argument. The executable map is expected to return the content of an automounter map entry on its stdout or no output if the entry cannot be determined. A direct map cannot be made executable.

Configuration and the auto_master Map

When initiated without arguments, `automount` consults the master map for a list of `autoofs` mount points and their maps. It mounts any `autoofs` mounts that are not already mounted, and unmounts `autoofs` mounts that have been removed from the master map or direct map.

The master map is assumed to be called `auto_master` and its location is determined by the name service switch policy. Normally the master map is located initially as a local file `/etc/auto_master`.

Browsing

The `automount` daemon supports browsability of indirect maps. This allows all of the potential mount points to be visible, whether or not they are mounted. The `-nobrowse` option can be added to any indirect `autoofs` map to disable browsing. For example:

```
/net      -hosts      -nosuid,nobrowse  
/home     auto_home
```

In this case, any *hostnames* would only be visible in `/net` after they are mounted, but all potential mount points would be visible under `/home`. The `-browse` option enables browsability of `autoofs` file systems. This is the default for all indirect maps.

The `-browse` option does not work in conjunction with the wildcard key.

Restricting Mount Maps Options specified for a map are used as the default options for all the entries in that map. They are ignored when map entries specify their own mount options.

In some cases, however, it is desirable to force `nosuid`, `nodevices`, `nosetuid`, or `noexec` for a complete mount map and its submounts. This can be done by specifying the additional mount option, `- restrict`.

```
/home      auto_home      - restrict,nosuid,hard
```

The `- restrict` option forces the inheritance of all the restrictive options `nosuid`, `nodevices`, `nosetuid`, and `noexec` as well as the `restrict` option itself. In this particular example, the `nosuid` and `restrict` option are inherited but the `hard` option is not. The `restrict` option also prevents the execution of “executable maps” and is enforced for auto mounts established by programs with fewer than all privileges available in their zone.

退出状态 The following exit values are returned:

- 0 Successful completion.
- 1 An error occurred.

文件

- `/etc/auto_master` Master automount map.
- `/etc/auto_home` Map to support automounted home directories.
- `/etc/nsswitch.conf` Name service switch configuration file. See [nsswitch.conf\(4\)](#).

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [isainfo\(1\)](#), [ls\(1\)](#), [svcs\(1\)](#), [uname\(1\)](#), [updatehome\(1\)](#), [automountd\(1M\)](#), [mount\(1M\)](#), [mount_nfs\(1M\)](#), [sharectl\(1M\)](#), [svcadm\(1M\)](#), [attributes\(5\)](#), [nfssec\(5\)](#), [smf\(5\)](#)

《Oracle Solaris 管理：网络服务》

附注

`autofs` mount points must not be hierarchically related. `automount` does not allow an `autofs` mount point to be created within another `autofs` mount.

Since each direct map entry results in a new `autofs` mount such maps should be kept short.

Entries in both direct and indirect maps can be modified at any time. The new information is used when `automountd` next uses the map entry to do a mount.

New entries added to a master map or direct map will not be useful until the `automount` command is run to install them as new `autofs` mount points. New entries added to an indirect map may be used immediately.

As of the Solaris 2.6 release, a listing (see [ls\(1\)](#)) of the `autofs` directory associated with an indirect map shows all potential mountable entries. The attributes associated with the potential mountable entries are temporary. The real file system attributes will only be shown once the file system has been mounted.

Default mount options can be assigned to an entire map when specified as an optional third field in the master map. These options apply only to map entries that have no mount options. Note that map entities with options override the default options, as at this time, the options do not concatenate. The concatenation feature is planned for a future release.

When operating on a map that invokes an NFS mount, the default number of retries for the automounter is 0, that is, a single mount attempt, with no retries. Note that this is significantly different from the default (10000) for the [mount_nfs\(1M\)](#) utility.

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same.

引用名	automountd – autofs mount/unmount daemon
用法概要	automountd [-Tvn] [-D <i>name=value</i>]
描述	automountd is an RPC server that answers file system mount and unmount requests from the autofs file system. It uses local files or name service maps to locate file systems to be mounted. These maps are described with the automount(1M) command. If automount finds any non-trivial entries in either the local or distributed automount maps and if the daemon is not running already, the automountd daemon is automatically invoked by automount(1M). automountd enables the svc:/network/nfs/nlockmgr service (lockd(1M)), and the svc:/network/nfs/status service (statd(1M)), if NFS mounts need to be done.
SMF Management	The automountd service is managed by the service management facility, smf(5), under the service identifier: <pre>svc:/system/filesystem/autofs</pre> Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M). The service's status can be queried using svcs(1). If it is disabled, it is enabled by automount(1M) unless the application/auto_enable property is set to false. Values for parameters for automountd are stored in SMF and manipulated using sharectl(1M). The following are the supported parameters. automountd_verbose=TRUE FALSE Verbose mode. Causes status messages to be logged to: <pre>/var/svc/log/system-filesystem-autofs:default.log</pre> See smf(5). The default value is FALSE. Equivalent to the -v option. nobrowse=ON OFF Turn on or off browsing for all autofs mount points. The default value is OFF. Equivalent to the -n option. trace=num Expands each RPC call and logs it to: <pre>/var/svc/log/system-filesystem-autofs:default.log</pre> See smf(5). The default value, 0, turns off such tracing. Starting with 1, with each higher value, the verbosity of trace output increases. This property is equivalent to the -T option. environment=name=value Environment variables. You can specify multiple <i>name=value</i> pairs with a comma separator. If there is more than one value for a name, use the double forward slash to indicate that. See EXAMPLES. There are no environment variable settings supplied. This property is equivalent to the -D option. All the above parameters can be changed using the sharectl(1M) command.

选项	The following options are supported: -D <i>name=value</i> Assign <i>value</i> to the indicated automount map substitution variable. These assignments cannot be used to substitute variables in the master map <code>auto_master</code>. -n Turn off browsing for all autofs mount points. This option overrides the <code>-browse</code> autofs map option on the local host. -T Trace. Expand each RPC call and display it on the standard output. -v Verbose. Log status messages to the console.

用法 See [largefile\(5\)](#) for the description of the behavior of automountd when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).

示例 示例 1 Setting Verbose Mode

The following command turns on verbose mode, causing status messages to be logged to the location specified in “SMF Management,” above. See that section for a description of the properties specified in this and following examples.

```
# sharectl set -p automountd_verbose=true autofs
```

示例 2 Turning on Browsing for Mount Points

The following command turns on browsing for all autofs mount points.

```
# sharectl set -p nobrowse=off autofs
```

示例 3 Specifying and Displaying Environment Variables

The following commands set and display the value for an environment variable named DAY.

```
# sharectl set -p environment=DAY=TUES autofs
% sharectl get -p environment autofs
environment=DAY=TUES
```

The following command sets multiple parameters for an environment variable. The subsequent command displays the result.

```
# sharectl set -p environment=DAY=TUES,TIME=NOON autofs
% sharectl get -p environment autofs
environment=DAY=TUES,TIME=NOON
```

The following command set multiple values for name that is assigned to an environment variable. The subsequent command displays the result.

```
# sharectl set -p environment=DAY=MON\\,TUE,TIME=NOON autofs
% sharectl get -p environment autofs
environment=DAY=MON\\,TUE,TIME=NOON
```

文件 /etc/auto_master Master map for automounter.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [svcs\(1\)](#), [automount\(1M\)](#), [sharectl\(1M\)](#), [svcadm\(1M\)](#), [attributes\(5\)](#), [largefile\(5\)](#), [smf\(5\)](#)

引用名	autopush – configures lists of automatically pushed STREAMS modules
用法概要	<pre>autopush -f <i>filename</i> autopush -g -M <i>major</i> -m <i>minor</i> autopush -r -M <i>major</i> -m <i>minor</i></pre>
描述	The autopush command configures the list of modules to be automatically pushed onto the stream when a device is opened. It can also be used to remove a previous setting or get information on a setting.
选项	The following options are supported: -f <i>filename</i> Sets up the autopush configuration for each driver according to the information stored in <i>filename</i>. An autopush file consists of lines of four or more fields, separated by spaces as shown below: <pre><i>major minor last-minor module1 module2 . . . module8</i></pre> The first field is a string that specifies the <i>major</i> device name, as listed in the <code>/kernel/drv</code> directory. The next two fields are integers that specify the <i>minor</i> device number and <i>last-minor</i> device number. The fields following represent the names of modules. If <i>minor</i> is <code>-1</code>, then all minor devices of a major driver specified by <i>major</i> are configured, and the value for <i>last-minor</i> is ignored. If <i>last-minor</i> is <code>0</code>, then only a single minor device is configured. To configure a range of minor devices for a particular major, <i>minor</i> must be less than <i>last-minor</i>. The remaining fields list the names of modules to be automatically pushed onto the stream when opened, along with the position of an optional anchor. The maximum number of modules that can be pushed is eight. The modules are pushed in the order they are specified. The optional special character sequence <code>[anchor]</code> indicates that a STREAMS anchor should be placed on the stream at the module previously specified in the list; it is an error to specify more than one anchor or to have an anchor first in the list. A nonzero exit status indicates that one or more of the lines in the specified file failed to complete successfully. -g Gets the current configuration setting of a particular <i>major</i> and <i>minor</i> device number specified with the <code>-M</code> and <code>-m</code> options respectively and displays the autopush modules associated with it. It will also return the starting minor device number if the request corresponds to a setting of a range (as described with the <code>-f</code> option). -m <i>minor</i> Specifies the minor device number. -M <i>major</i> Specifies the major device number.

-r
Removes the previous configuration setting of the particular *major* and *minor* device number specified with the -M and -m options respectively. If the values of *major* and *minor* correspond to a previously established setting of a range of minor devices, where *minor* matches the first minor device number in the range, the configuration would be removed for the entire range.

退出状态 The following exit values are returned:

- 0
Successful completion.
- non-zero
An error occurred.

示例 示例 1 Using the autopush command.

The following example gets the current configuration settings for the *major* and *minor* device numbers as indicated and displays the autopush modules associated with them for the character-special device /dev/term/a:

```
example# autopush -g -M 29 -m 0
Major      Minor      Lastminor      Modules
  29         0          1      ldterm ttcompat
```

文件 /etc/iu.ap

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [dladm\(1M\)](#), [ttymon\(1M\)](#), [attributes\(5\)](#), [ldterm\(7M\)](#), [sad\(7D\)](#), [streamio\(7I\)](#), [ttcompat\(7M\)](#)

《STREAMS Programming Guide》

附注 The use of the autopush command is obsolete for networking data-links. The preferred method of configuring a list of automatically pushed STREAMS modules on a given networking data-link interface is the [dladm\(1M\)](#) autopush link property.

Because network devices are self-cloning, the autopush command is inadequate for networking devices. The granularity of the autopush command's configuration is driver-wide, and not per-interface as one might expect. Another reason not to use autopush is that administrators are more familiar with the names of their network interfaces than with device major and minor numbers. The [dladm\(1M\)](#) command allows the configuration using data-link interface names.

引用名

bart – basic audit reporting tool

用法概要

```
/usr/bin/bart create [-n] [-R root_directory]
                    [-r rules_file | -] [-a md5|sha1|sha256|sha384|sha512]

/usr/bin/bart create [-n] [-R root_directory] -I
                    [-a md5|sha1|sha256|sha384|sha512] [file_name]...

/usr/bin/bart compare [-i attribute] [-p]
                     [-r rules_file | -] control-manifest test-manifest
```

描述

bart(1M) is a tool that performs a file-level check of the software contents of a system.

You can also specify the files to track and the types of discrepancies to flag by means of a rules file, `bart_rules`. See `bart_rules(4)`.

The bart utility performs two basic functions:

bart create The manifest generator tool takes a file-level *snapshot* of a system. The output is a catalog of file attributes referred to as a *manifest*. See `bart_manifest(4)`.

You can specify that the list of files be cataloged in three ways. Use `bart create` with no options, specify the files by name on the command line, or create a rules file with directives that specify which the files to monitor. See `bart_rules(4)`.

By default, the manifest generator catalogs all attributes of all files in the root (/) file system. File systems mounted on the root file system are cataloged only if they are of the same type as the root file system.

For example, /, /usr, and /opt are separate UFS file systems. /usr and /opt are mounted on /. Therefore, all three file systems are cataloged. However, /tmp, also mounted on /, is not cataloged because it is a TMPFS file system. Mounted CD-ROMs are not cataloged since they are HSFS file systems.

bart compare The report tool compares two manifests. The output is a list of per-file attribute discrepancies. These discrepancies are the differences between two manifests: a control manifest and a test manifest.

A discrepancy is a change to any attribute for a given file cataloged by both manifests. A new file or a deleted file in a manifest is reported as a discrepancy.

The reporting mechanism provides two types of output: verbose and programmatic. Verbose output is localized and presented on multiple lines, while programmatic output is more easily parsable by other programs. See `OUTPUT`.

By default, the report tool generates verbose output where all discrepancies are reported except for modified directory timestamps (`dirmtime` attribute).

To ensure consistent and accurate comparison results, *control-manifest* and *test-manifest* must be built with the same rules file.

Use the rules file to ignore specified files or subtrees when you generate a manifest or compare two manifests. Users can compare manifests from different perspectives by re-running the `bart compare` command with different rules files.

选项

The following options are supported:

- i *attribute ...* Specify the file attributes to be ignored globally. Specify attributes as a comma separated list.

This option produces the same behavior as supplying the file attributes to a global `IGNORE` keyword in the rules file. See `bart_rules(4)`.
- I [*file_name...*] Specify the input list of files. The file list can be specified at the command line or read from standard input.
- n Prevent computation of content signatures for all regular files in the file list.
- p Display manifest comparison output in “programmatic mode,” which is suitable for programmatic parsing. The output is not localized.
- r *rules_file* Use *rules_file* to specify which files and directories to catalog, and to define which file attribute discrepancies to flag. If *rules_file* is -, then the rules are read from standard input. See `bart_rules(4)` for the definition of the syntax.
- R *root_directory* Specify the root directory for the manifest. All paths specified by the rules, and all paths reported in the manifest, are relative to *root_directory*.

注 – The root file system of any non-global zones must not be referenced with the -R option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See [zones\(5\)](#).
- a [*hash*] Specify the cryptographic digest algorithm to use for the hash of the file contents: md5, sha1, sha256, sha512 are the currently supported values. If -a is not specified, sha256 is used and a Version 1.1 manifest is created that indicates which hash algorithm is used. If md5 is specified then a Version 1.0 manifest is created.

The following operands are supported:

<i>test-manifest</i>	Specify the manifest created by <code>bart create</code> on the test system.

Output

The `bart create` command generates a system manifest. See `bart_manifest(4)`.

When the `bart compare` command compares two system manifests, it generates a list of file differences. By default, the comparison output is localized. However, if the `-p` option is specified, the output is generated in a form that is suitable for programmatic manipulation.

Default Format

filename

```
attribute control:xxxx test:yyyy
```

<i>filename</i>	Name of the file that differs between <i>control-manifest</i> and <i>test-manifest</i> . For file names that contain embedded whitespace or newline characters, see <code>bart.manifest(4)</code> .

attribute The name of the file attribute that differs between the manifests that are compared. *xxxx* is the attribute value from *control-manifest*, and *yyyy* is the attribute value from *test-manifest*. When discrepancies for multiple attributes occur for the same file, each difference is noted on a separate line.

The following attributes are supported:

acl	ACL attributes for the file. For a file with ACL attributes, this field contains the output from <code>acltotext()</code> .

all	All attributes.

contents	Checksum value of the file. This attribute is only specified for regular files. If you turn off context checking or if checksums cannot be computed, the value of this field is -.

dest	Destination of a symbolic link.

devnode	Value of the device node. This attribute is for character device files and block device files only.

dirmtime Modification time in seconds since 00:00:00 UTC, January 1, 1970 for directories.

gid	Numerical group ID of the owner of this entry.

lnmtime	Creation time for links.
mode	Octal number that represents the permissions of the file.
mtime	Modification time in seconds since 00:00:00 UTC, January 1, 1970 for files.
size	File size in bytes.
type	Type of file.
uid	Numerical user ID of the owner of this entry.

The following default output shows the attribute differences for the `/etc/passwd` file. The output indicates that the `size`, `mtime`, and `contents` attributes have changed.

```
/etc/passwd:
  size control:74 test:81
  mtime control:3c165879 test:3c165979
  contents control:daca28ae0de97afd7a6b91fde8d57afa
test:84b2b32c4165887355317207b48a6ec7
```

Programmatic Format	<i>filename attribute control-val test-val [attribute control-val test-val]*</i>
	<i>filename</i> Same as <i>filename</i> in the default format.
	<i>attribute control-val test-val</i> A description of the file attributes that differ between the control and test manifests for each file. Each entry includes the attribute value from each manifest. See <code>bart_manifest(4)</code> for the definition of the attributes.

Each line of the programmatic output describes all attribute differences for a single file.

The following programmatic output shows the attribute differences for the `/etc/passwd` file. The output indicates that the `size`, `mtime`, and `contents` attributes have changed.

```
/etc/passwd size 74 81 mtime 3c165879 3c165979
contents daca28ae0de97afd7a6b91fde8d57afa 84b2b32c4165887355317207b48a6ec7
```

退出状态

Manifest Generator	The manifest generator returns the following exit values:
	0 Success
	1 Non-fatal error when processing files; for example, permission problems
	>1 Fatal error; for example, invalid command-line options
Report Tool	The report tool returns the following exit values:
	0 No discrepancies reported

- 1 Discrepancies found
- >1 Fatal error executing comparison

示例

示例 1 Creating a Default Manifest Without Computing Checksums

The following command line creates a default manifest, which consists of all files in the / file system. The -n option prevents computation of checksums, which causes the manifest to be generated more quickly.

```
bart create -n
```

示例 2 Creating a Manifest for a Specified Subtree

The following command line creates a manifest that contains all files in the /home/nickiso subtree.

```
bart create -R /home/nickiso
```

示例 3 Creating a Manifest by Using Standard Input

The following command line uses output from the find(1) command to generate the list of files to be cataloged. The find output is used as input to the bart create command that specifies the -I option.

```
find /home/nickiso -print | bart create -I
```

示例 4 Creating a Manifest by Using a Rules File

The following command line uses a rules file, rules, to specify the files to be cataloged.

```
bart create -r rules
```

示例 5 Comparing Two Manifests and Generating Programmatic Output

The following command line compares two manifests and produces output suitable for parsing by a program.

```
bart compare -p manifest1 manifest2
```

示例 6 Comparing Two Manifests and Specifying Attributes to Ignore

The following command line compares two manifests. The dirmtime, lnmtime, and mtime attributes are not compared.

```
bart compare -i dirmtime,lnmtime,mtime manifest1 manifest2
```

示例 7 Comparing Two Manifests by Using a Rules File

The following command line uses a rules file, rules, to compare two manifests.

```
bart compare -r rules manifest1 manifest2
```

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	security/bart
Interface Stability	Committed

另请参见 [cksum\(1\)](#), [digest\(1\)](#), [find\(1\)](#), [bart_manifest\(4\)](#), [bart_rules\(4\)](#), [attributes\(5\)](#)

附注 The file attributes of certain system libraries can be temporarily altered by the system as it boots. To avoid triggering false warnings, you should compare manifests only if they were both created with the system in the same state; that is, if both were created in single-user or both in multi-user.

引用名 beadm – 管理 ZFS 引导环境

用法概要

```
beadm create [-a] [-d description]
               [-e non-activeBeName | beName@snapshot]
               [-o property=value] ... [-p zpool] beName

beadm create beName@snapshot

beadm destroy [-fF] beName | beName@snapshot

beadm list [-a | -ds] [-H] [beName]

beadm mount beName mountpoint

beadm unmount [-f] beName

beadm rename beName newBeName

beadm activate beName
```

描述 beadm 命令是管理 ZFS 引导环境 (Boot Environment, BE) 的用户接口。该实用程序由需要在一个系统中管理多个 Oracle Solaris 实例的系统管理员使用。

使用 beadm 可进行如下操作：

- 基于活动引导环境创建新的引导环境。
- 基于非活动引导环境创建新的引导环境。
- 创建现有引导环境的快照。
- 基于现有快照创建新的引导环境。
- 创建新的引导环境并将其复制到另一个 zpool 中。
- 激活现有的非活动引导环境。
- 挂载引导环境。
- 卸载引导环境。
- 销毁引导环境。
- 销毁引导环境的快照。
- 重命名现有的非活动引导环境。
- 显示有关快照和数据集的信息。

子命令 beadm 命令包含以下子命令和选项。下面的“示例”部分对其中的许多子命令和选项的用法进行了说明。

beadm (*no arguments*)
显示命令用法。

```
beadm create [-a] [-d description] [-e non-activeBeName | beName@snapshot] [-o
property=value] ... [-p zpool] beName
```

创建名为 *beName* 的新引导环境。如果不提供 -e 选项，新引导环境将会创建为当前运行的引导环境的克隆。如果提供 -d 选项，相应描述还可作为 x86 系统的 GRUB 菜单或 SPARC 系统的引导菜单中的引导环境的标题。如果不提供 -d 选项，*beName* 可以用作标题。嵌套引导环境不支持使用 -p 选项。此外，不可引导的嵌套引导环境及其快照不能与 -e 选项一起使用。

- a
创建时激活新创建的引导环境。缺省设置为不激活新创建的引导环境。
- d *description*
创建带有相关描述的新引导环境。
- e *non-activeBeName*
基于现有的非活动引导环境创建新引导环境。在嵌套引导环境中，只有可引导的引导环境可以与此选项一起使用。
- e *beName@snapshot*
基于引导环境 *beName* 的现有快照创建新引导环境。在嵌套引导环境中，只有可引导的引导环境快照可以与此选项一起使用。
- o *property=value*
为具有特定 ZFS 属性的新引导环境创建数据集。可以指定多个 -o 选项。有关 -o 选项的更多信息，请参见 [zfs\(1M\)](#)。
- p *zpool*
在指定的 *zpool* 中创建新引导环境。如果不提供该选项，则缺省行为是在与原始引导环境相同的池中创建新引导环境。此选项在嵌套引导环境内不受支持。

beadm create *beName@snapshot*

创建现有引导环境 *beName* 的快照。在嵌套引导环境内，只能为可引导的引导环境创建快照。在嵌套引导环境内，只能销毁可引导的引导环境，以及不可引导、但在重新引导时未标记为活动的引导环境。

beadm destroy [-fF] *beName* | *beName@snapshot*

销毁名为 *beName* 的引导环境或销毁该引导环境下名为 *beName@snapshot*。如果销毁一个引导环境，该引导环境下的所有快照也均被销毁。使用该命令时需慎重。

- f
如果当前引导环境已挂载，强制卸载该引导环境。
- F
强制销毁引导环境，而不提示进行确认。

beadm list [-a | -ds] [-H] [*beName*]

列出有关现有引导环境 *beName* 的信息，或列出有关所有引导环境（如果不提供 *beName*）的信息。Active（活动）字段指示引导环境当前是否为活动状态。N 表示引导环境当前为活动状态；R 表示重新引导时为活动状态；NR 表示引导环境在当前以及重新引导时均为活动状态。嵌套引导环境内的不可引导的引导环境由叹号(!) 表示。

计算机可解析的输出中的每一行将引导环境名称作为第一个字段。Space（空间）字段以字节显示。Created（创建时间）字段以 UTC（国际协调时间）格式显示。如果不与其他选项一起使用，-H 选项在第二个字段中指定引导环境的 UUID（通用唯一标识符）。如果该引导环境不存在 UUID，该字段为空。请参见“示例”部分。在嵌套引导环境内，UUID 字段实际表示与嵌套引导环境相关联的父 ID。

- a
列出有关引导环境的所有可用信息。其中包括从属文件系统和快照。
- d
列出属于该引导环境的所有从属文件系统的信息。
- s
列出有关引导环境的快照的信息。
- H
不列出标题信息。列表信息中的每个字段由分号分隔。

beadm mount *beName mountpoint*

在 *mountpoint* 上挂载名为 *beName* 的引导环境。*mountpoint* 必须是已存在的空目录。

beadm unmount [-f] *beName*

卸载名为 *beName* 的引导环境。

- f
强制卸载引导环境，即使引导环境当前处于忙碌状态也是如此。

beadm rename *beName newBeName*

将名为 *beName* 的引导环境重命名为 *newBeName*。在嵌套引导环境中，只能对可引导的引导环境进行重命名。

beadm activate *beName*

在下次重新引导时使 *beName* 成为活动引导环境。在嵌套引导环境中，只能激活可引导的引导环境。

嵌套引导环境支持

beadm 支持嵌套引导环境概念，具体而言，因为它与非全局区域的引导环境相关。目前，**beadm** 只能从非全局区域内部管理嵌套引导环境。

beadm 函数在非全局区域内与从全局区域运行的情况大致相同，但有少数例外。第一，**beadmcreate** 的 **-p**（备用池）选项在非全局区域内不受支持。第二，在确定任何给定嵌套引导环境（或引导环境快照）是可引导还是不可引导时存在一定的差别。如果嵌套引导环境与当前活动的全局区域引导环境相关联（即，与活动的全局区域引导环境的 UUID 共享同一父 ID），则它是可引导的。否则，它便是不可引导的，且在 **beadm list** 的活动列中标记有 "!"。请注意，虽然非全局区域管理员可以通过 **beadm activate** 将此类引导环境标记为活动，但是重新引导非全局区域并不会导致装入引导环境，因为该引导环境与非活动全局区域引导环境相关联。根据这些条件，**beadm** 于是会限制对不可引导的引导环境执行某些操作：

- 不能销毁不可引导且在重新引导时标记为活动的嵌套引导环境。
- 不能激活不可引导的引导环境。
- 不能为不可引导的引导环境创建快照。
- 不能将不可引导的引导环境或引导环境快照与 **beadm create** 的 **-e** 选项一起使用。
- 不能重命名不可引导的引导环境。

示例

示例1 基于活动的引导环境创建新的引导环境

以下命令通过克隆当前的引导环境来创建新的引导环境 BE1。

```
# beadm create BE1
```

示例2 基于非活动的引导环境创建新的引导环境

以下命令通过克隆现有的名为 BE1 的引导环境创建新的引导环境 BE2。

```
# beadm beadm create -e BE1 BE2
```

示例3 创建现有的引导环境的快照

以下命令为现有的名为 BE1 的引导环境创建名为 now 的快照。

```
# beadm create BE1@now
```

示例4 通过克隆快照创建新的引导环境

以下命令通过克隆名为 BE1 的引导环境的现有快照创建名为 BE3 的新引导环境。

```
# beadm create -e BE1@now BE3
```

示例5 在指定的 zpool 中创建新的引导环境

以下命令基于当前正在运行的引导环境创建名为 BE4 的新引导环境。该命令在名为 rpool2 的 zpool 中创建一个新的引导环境。

```
# beadm create -p rpool2 BE4
```

示例6 在指定的启用了压缩功能的 zpool 中创建新的引导环境

以下命令基于当前正在运行的引导环境创建名为 BE5 的新引导环境。该命令在名为 rpool2 的 zpool 中创建新的引导环境以及启用了压缩功能的数据集。

```
# beadm create -p rpool2 -o compression=on BE5
```

示例7 创建新的引导环境并提供描述

以下命令基于当前正在运行的引导环境创建名为 BE6 的新引导环境并提供描述。

```
# beadm create -d "BE6 used as test environment" BE6
```

示例8 激活引导环境

以下命令激活现有的名为 BE3 的非活动引导环境。

```
# beadm activate BE3
```

示例9 挂载引导环境

以下命令将名为 BE3 的引导环境挂载到 /mnt。

```
# beadm mount BE3 /mnt
```

示例 10 卸载引导环境

以下命令卸载名为 BE3 的引导环境。

```
# beadm unmount BE3
```

示例 11 销毁引导环境

以下命令销毁名为 BE3 的引导环境，而不提示进行确认。

```
# beadm destroy -F BE3
```

示例 12 销毁快照

以下命令销毁 BE1 引导环境的名为 now 的快照。

```
# beadm destroy BE1@now
```

示例 13 重命名引导环境

以下命令将现有的名为 BE1 的非活动引导环境重命名为 BE3。

```
# beadm rename BE1 BE3
```

示例 14 列出所有的引导环境

以下命令列出所有的现有引导环境。

```
# beadm list
BE  Active Mountpoint Space  Policy Created
--  - - - - -
BE2 -      -          72.0K  static 2008-05-21 12:26
BE3 -      -          332.0K static 2008-08-26 10:28
BE4 -      -          15.78M static 2008-09-05 18:20
BE5 NR    /           7.25G  static 2008-09-09 16:53
```

示例 15 列出所有的引导环境及其数据集和快照信息

以下命令列出所有现有的引导环境及其所有数据集和快照信息。

```
# beadm list -d -s
BE/Dataset/Snapshot      Active Mountpoint Space  Policy Created
-----
BE2
  p/ROOT/BE2             -    -          36.0K  static 2008-05-21 12:26
  p/ROOT/BE2/opt          -    -          18.0K  static 2008-05-21 16:26
  p/ROOT/BE2/opt@now      -    -           0      static 2008-09-08 22:43
  p/ROOT/BE2@now          -    -           0      static 2008-09-08 22:43
BE3
  p/ROOT/BE3             -    -          192.0K static 2008-08-26 10:28
  p/ROOT/BE3/opt          -    -           86.0K static 2008-08-26 10:28
  p/ROOT/BE3/opt/local    -    -           36.0K static 2008-08-28 10:58
BE4
```


示例 15 列出所有的引导环境及其数据集和快照信息 (续)

p/ROOT/BE4	-	-	15.78M	static	2008-09-05	18:20
BE5						
p/ROOT/BE5	NR	/	6.10G	static	2008-09-09	16:53
p/ROOT/BE5/opt	-	/opt	24.55M	static	2008-09-09	16:53
p/ROOT/BE5/opt@bar	-	-	18.38M	static	2008-09-10	00:59
p/ROOT/BE5/opt@foo	-	-	18.38M	static	2008-06-10	16:37
p/ROOT/BE5@bar	-	-	139.44M	static	2008-09-10	00:59
p/ROOT/BE5@foo	-	-	912.85M	static	2008-06-10	16:37

示例 16 列出一个引导环境的数据集和快照信息
以下命令列出名为 BE5 的引导环境的所有数据集和快照信息。

```
# beadm list -a BE5
```

BE/Dataset/Snapshot	Active	Mountpoint	Space	Policy	Created
-----	-----	-----	-----	-----	-----
BE5					
p/ROOT/BE5	NR	/	6.10G	static	2008-09-09 16:53
p/ROOT/BE5/opt	-	/opt	24.55M	static	2008-09-09 16:53
p/ROOT/BE5/opt@bar	-	-	18.38M	static	2008-09-10 00:59
p/ROOT/BE5/opt@foo	-	-	18.38M	static	2008-06-10 16:37
p/ROOT/BE5@bar	-	-	139.44M	static	2008-09-10 00:59
p/ROOT/BE5@foo	-	-	912.85M	static	2008-06-10 16:37

示例 17 按计算机可解析的格式列出
以下命令按照计算机可解析的格式列出所有引导环境的信息。

```
# beadm list -H
BE2;;;55296;static;1211397974
BE3;;;339968;static;1219771706
BE4;;;16541696;static;1220664051
BE5;215b8387-4968-627c-d2d0-f4a011414bab;NR;/;7786206208;static;1221004384
```

示例 18 显示不可引导的引导环境
以下命令将列出所有引导环境。在非全局区域内运行时，它会显示可引导的和不可引导的引导环境。不可引导的引导环境在活动列中以叹号 (!) 指明。

```
# beadm list
```

BE	Active	Mountpoint	Space	Policy	Created
--	-----	-----	-----	-----	-----
zbe-0	-	-	29.22M	static	2011-03-04 09:14
zbe-1	NR	/	815.10M	static	2011-03-04 09:28
zbe-2	-	-	35.0K	static	2011-03-04 09:28
zbe-3	-	-	35.0K	static	2011-03-04 09:28
zbe-4	-	-	35.0K	static	2011-03-04 09:28
zbe-5	!	-	35.0K	static	2011-03-04 11:47
zbe-6	!	-	54.0K	static	2011-03-07 14:37

退出状态 0
 成功。

 >0
 失败。

文件 /var/log/beadm/beName/create.log.yyyymmdd_hhmmss
 用于捕获 beadm create 输出的日志。文件名中的时间部分解释如下。

- *yyyyymmdd_hhmmss*—例如，20071130_140558。
- *yyyy*—年， 2007
- *mm*—月， 11
- *dd*—日， 30
- *hh*—时， 14
- *mm*—分， 05
- *ss*—秒， 58

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/boot-environment-utilities
接口稳定性	Committed（已确定）

另请参见 [zfs\(1M\)](#)、[attributes\(5\)](#)

引用名 boot – 启动系统内核或独立程序

用法概要

SPARC boot [OBP names] [file] [-aLV] [-F object] [-D default-file]
 [-Z dataset] [boot-flags] [—] [client-program-args]

x86 kernel\$ /platform/i86pc/kernel/\$ISADIR/unix [boot-args]
 [-B prop=val [,val...]]

描述 引导是装入并执行独立程序的过程。出于此探讨的目的，引导意即装入并执行可引导操作系统的过程。一般而言，独立程序是操作系统内核（请参见 [kernel\(1M\)](#)），但任何独立程序都可以进行引导。在基于 SPARC 的系统上，用于计算机的诊断监视程序就是一个可引导的独立程序而不是操作系统的不错示例。

如果将该独立程序标识为动态链接的可执行文件，则 boot 将按照该可执行文件格式所示装入解释程序（链接程序/装载程序），然后将控制权转移给该解释程序。如果该独立程序是静态链接的，则将直接跳至该独立程序。

装入了内核后，将启动 UNIX 系统，挂载所需的文件系统（请参见 [vfstab\(4\)](#)），并运行 /usr/sbin/init 使系统处于 /etc/inittab 中指定的 "initdefault" 状态。请参见 [inittab\(4\)](#)。

SPARC 引导过程 在基于 SPARC 的系统上，多数计算机上的引导过程包含以下基本阶段。

打开计算机后，系统固件（在 PROM 中）将执行开机自检 (power-on self-test, POST)。自检的形式和范围取决于系统中固件的版本。

自检成功完成后，如果已在由固件使用的非易失性存储区域中设置了相应的标志，则固件将尝试进行自动引导。还可对要装入的文件名以及要从其中装入该文件的设备进行处理。

这些标志和名称可在 shell 中使用 [eeprom\(1M\)](#) 命令进行设置，或在停止系统后在 ok 提示符下使用 PROM 命令进行设置。

二级程序可以是特定于文件系统的引导块（从磁盘引导时），也可以是 inetboot 或 wanboot（通过网络引导时）。

网络引导

网络引导分为两步：客户机首先获取 IP 地址和任何其他允许其装入第二阶段引导程序所必需的参数。第二阶段引导程序进而从引导设备装入引导归档文件。

IP 地址可以通过三种方式获取：RARP、DHCP 或手动配置，具体取决于 PROM 中提供的功能和 PROM 的配置。sun4u 和 sun4v 内核体系结构的计算机都具有支持 DHCP 功能的 PROM。

指定两种网络引导方法的 boot 命令语法是：

```
boot net:rarp
boot net:dhcp
```

此命令：

```
boot net
```

没有 `rarp` 或 `dhcp` 说明符，将通过别名为 `net` 的网络接口调用缺省网络引导方法。

用于使用 `RARP/bootparams` 进行网络引导的事件序列在以下段落中进行了介绍。`DHCP` 的序列在 `RARP/bootparams` 说明之后。

使用 `RARP/bootparams` 通过网络进行引导时，`PROM` 将通过广播反向 `ARP` 请求开始，直至其收到回复。收到回复后，`PROM` 将广播 `TFTP` 请求以获取 `inetboot` 的第一块。后续请求将发送至最初应答第一块请求的服务器。装载后，`inetboot` 还将使用反向 `ARP` 获取其 `IP` 地址，然后广播 `bootparams` `RPC` 调用（请参见 [bootparams\(4\)](#)）以定位配置信息及其根文件系统。之后，`inetboot` 通过 `NFS` 方式装入引导归档文件并将控制权转移给该归档文件。

使用 `DHCP` 通过网络进行引导时，`PROM` 将广播硬件地址和内核体系结构，并请求 `IP` 地址、引导参数和网络配置信息。某个 `DHCP` 服务器响应并被选定后（可能存在多个服务器），该服务器将向客户机发送 `IP` 地址和引导客户机所需的所有其他信息。收到该信息后，客户机 `PROM` 将检查要装入的文件名并采取对应的两种操作方式之一，具体取决于该文件的名称看起来是否是 `HTTP URL`。如果不是，`PROM` 将下载 `inetboot`，将该文件装入至内存中并进行执行。`inetboot` 将装入接管计算机并释放 `inetboot` 的引导归档文件。然后，启动脚本初始化 `DHCP` 代理（请参见 [dhcpage\(1M\)](#)），用于实现进一步的 `DHCP` 活动。

如果要装入的文件是 `HTTP URL`，则 `PROM` 将使用 `HTTP` 装入引用的文件。如果客户机已配置了 `HMAC SHA-1` 密钥，其将在继续执行装入的文件之前检查该文件的完整性。该文件应是 `wanboot` 二进制文件。`WAN Boot` 过程可以配置为使用 `DHCP` 或 `NVRAM` 属性，以搜索安装服务器和路由器以及需要连接至其的代理。当 `wanboot` 开始执行时，将确定提供的信息是否足够供其继续。如果缺少任何必要信息，将退出并报告相应错误，或启动命令解释程序并提示提供更多的配置信息。一旦 `wanboot` 获取了必要的信息，就会通过 `HTTP` 方式将引导装载程序装入至内存。如果客户机上安装了加密密钥，`wanboot` 将验证引导装载程序的签名及其随附的散列。存在加密密钥但没有散列密钥是错误的。

`wanboot` 引导装载程序可以与使用 `HTTP` 或安全 `HTTP` 的客户机通信。如果是前者，并且如果客户机已配置了 `HMAC SHA-1` 密钥，则引导装载程序将检查根文件系统的完整性。在将根文件系统装入至内存后（可能已经执行了完整性检查），引导归档文件将从服务器进行传送。如果提供有 `boot_logger URL`（通过 [wanboot.conf\(4\)](#) 文件方式），`wanboot` 将定期记录其进度。

并非所有 `PROM` 都能够使用 `URL`。可以使用 `list-security-keys` `OBP` 命令确定客户机是否具备此功能（请参见 [monitor\(1M\)](#)）。

`WAN` 引导目前在 `x86` 平台上并未提供。

`wanboot` 命令行

当客户机程序是 wanboot 时，将接受以下格式的 client-program-args：

```
boot ... -o opt1[,opt2[,...]]
```

其中每个选项可能是一个操作：

dhcp

要求 wanboot 通过 DHCP 方式获取配置参数。

prompt

导致 wanboot 进入其命令解释程序。

<cmd>

下面列出了解释程序命令之一。

...或赋值，使用下面列出的解释程序参数名称。

wanboot 命令解释程序

wanboot 命令解释程序通过在引导时提供 "-o prompt" 的 client-program-args 进行调用。输入包括单一的命令或赋值，或以逗号分隔的命令或赋值列表。配置参数为：

host-ip

客户机的 IP 地址（采用点分十进制记法）

router-ip

缺省路由器的 IP 地址（采用点分十进制记法）

subnet-mask

子网掩码（采用点分十进制记法）

client-id

DHCP 客户机标识符（引号括起的 ASCII 字符串或十六进制 ASCII）

hostname

DHCP 事务中请求的主机名 (ASCII)

http-proxy

HTTP 代理服务器规范 (IPADDR[:PORT])

密钥名是：

3des

三重 DES 加密密钥（48 个十六进制 ASCII 字符）

aes

AES 加密密钥（32 个十六进制 ASCII 字符）

sha1

HMAC SHA-1 签名密钥（40 个十六进制 ASCII 字符）

最后，URL 或 WAN Boot CGI 将通过以下方式进行引用：

bootserver

WAN Boot CGI 的 URL（等同于 OBP 的 `file` 参数）

解释程序接受以下命令：

help

输出可用命令的简要说明

var=val

将 `val` 赋值给 `var`，其中 `var` 是配置参数名、密钥名或 `bootserver` 之一。

var=

取消设置参数 `var`。

list

列出所有参数及其值（通过 OBP 检索的密钥值始终不会显示）。

prompt

提示输入取消设置参数的值。输出每个参数的名称及其当前值（如果有），用户可以接受此值（按回车键）或输入新的值。

go

用户在对所有输入的值满意后，可以保持解释程序处于现有状态并继续引导。

exit

退出引导解释程序并返回至 OBP 的 `ok` 提示符。

任何这些赋值或命令都可以在命令行上作为 `-o` 选项的一部分进行传递，遵循引导参数的 128 字节的 OBP 限制。例如，`-o list,go` 将仅列出参数的当前（缺省）值，然后继续引导。

iSCSI 引导

iSCSI 引导在 x86 和 SPARC 上均受支持。

基于 x86 的 iSCSI 引导

对于基于 x86 的 iSCSI 引导，接受引导的主机必须配备有支持 iBFT（iSCSI Boot Firmware Table，iSCSI 引导固件表）的 NIC，或该主机的主板的 BIOS 必须支持 iBFT。iBFT（在高级配置和电源接口 (Advanced Configuration and Power Interface, ACPI) 3.0b 规范中进行定义）指定了一个信息块，其中包含了可用于 iSCSI 引导过程的各种参数。

实现 iBFT 的固件通过建立与 iSCSI 目标的连接，于启动期间在 BIOS 中提供了一个作为可引导设备的 iSCSI 磁盘。iSCSI 引导的其余过程与从本地磁盘引导一样。

要正确配置 iBFT，用户需要参考来自硬件供应商的文档。

基于 SPARC 的 iSCSI 引导

基于 SPARC 的 iSCSI 引导在 OpenBoot level 4.31 及以上版本中受支持，无需特定 NIC。

OpenBoot 中的 `boot` 命令采用 `keyword=value` 格式提取一系列关键字以标识目的地 iSCSI 目标。iSCSI boot 命令的完整格式是：

```
boot net:iscsi-target-ip=t-ip,iscsi-target-name=name
      host-ip=h-ip[,router-ip=r-ip]
      [,subnet-mask=m-ip]
      [,iscsi-port=port]
      [,iscsi-lun=lun]
      [,iscsi-partition=partition]
```

上述关键字的描述如下所述：

host-ip	引导主机的 IP 地址。
router-ip	路由网关的 IP 地址。
subnet-mask	host-ip 的子网掩码。
iscsi-target-ip	存储 OS 的 iSCSI 目标的 IP 地址。
iscsi-target-name	存储 OS 的 iSCSI 目标的名称。
iscsi-partition	包含可引导根的分区。
iscsi-port	目标的 IP 端口。
iscsi-lun	目标上从其进行引导的 LUN。

iscsi-target-ip、route-ip 和 subnet-mask 的值都采用标准的 IPv4 点分十进制格式，例如 255.255.255.0（用于 subnet-mask）。IPv6 在当前 OpenBoot 实现中不受支持。

iscsi-port 的值（一个十进制数）介于 1 到 65535 范围之内。

iscsi-lun 的值采用以短划线相连的十六进制 LUN 格式，即 *fff-fff-fff-fff*。有关详细信息，请参见 RFC 4173 第 5 节。前导零和结尾短划线可以排除在外，因此，例如 3，等同于 *0003-0000-0000-0000*。

iscsi-partition 的值是一个 ASCII 字符，用于指定根分区。多数情况下，其值为 a。

iscsi-target-name 的值采用的格式为字符串，由 RFC 3720 和 RFC 3722 指定。

添加了两个安全密钥以在目标侧提供 CHAP 验证。它们是：

chap-user	CHAP 名称
chap-password	CHAP 密钥

目前，这两个密钥可以通过命令 `set-ascii-security-key` 在 Open Boot PROM (ok) 提示符下进行设置。例如：

```
ok set-ascii-security-key chap-user chap name
ok set-ascii-security-key chap-password chap password
```

目前还不支持双向验证。这两个变量只能在 Open Boot PROM 提示符下进行更改。

支持 RFC 4173，以从 DHCP 服务器检索 iSCSI 引导信息。DHCP 服务器必须为引导客户机指定根路径选项，这样，客户机才可通过简单的命令执行 iSCSI 引导。

```
boot net:dhcp
```

目前，密钥 `boot-device` 用于在 iSCSI 引导期间检索物理引导设备。该密钥在 Solaris 安装期间设置。手动修改的密钥值可能会中断 iSCSI 引导。

从磁盘引导

从磁盘引导时，OpenBoot PROM 固件将从指定为引导设备的分区的块 1 到 15 读取引导块。该独立引导程序通常包含能够读取引导归档文件的特定于文件系统的读取器。

如果到独立程序的路径名是相对的（不是以斜线开头），则二级引导将在依赖于平台的搜索路径中查找独立程序。该路径必须包含 `/platform/platform-name`。许多 SPARC 平台接下来会搜索特定于平台的路径条目 `/platform/hardware-class-name`。请参见 [filesystem\(5\)](#)。如果路径名是绝对的，则 `boot` 将使用指定路径。然后，`boot` 程序将在合适的地址装入独立程序，之后移交控制权。

在引导归档文件从引导设备进行了传输后，Solaris 即可初始化计算机并接管其控制权。下面的“引导归档文件阶段”中将进一步介绍该过程，该过程在所有平台上都是相同的。

如果文件名并未在命令行上给定或，例如，由 `boot-file` NVRAM 变量进行了指定，则 `boot` 将根据系统上安装的软件以及硬件和固件的功能选择要装入的合适缺省文件。

指向内核的路径不得包含任何空格。

从 ZFS 引导

从 ZFS 引导不同于从 UFS 引导，原因是，对于 ZFS，设备说明符标识存储池，而不是单个根文件系统。存储池可能包含多个可引导的数据集（即，根文件系统）。因此，从 ZFS 引导时，只是指定引导设备是不够的。还必须在由引导设备标识的池内标识根文件系统。缺省情况下，选择用来进行引导的数据集是由池的 `bootfs` 属性标识的数据集。通过使用 `-z` 选项指定备用可引导数据集可以覆盖此缺省选择。使用 `-L` 选项列出 ZFS 池内的可引导数据集。

引导归档文件阶段

引导归档文件包含一个使用内存磁盘挂载的文件系统映像。该映像是自描述的，特别是在引导块中包含了一个文件系统读取器。该文件系统读取器挂载并打开 RAM 磁盘映像，然后读取并执行其中包含的内核。缺省情况下，该内核位于：

```
/platform/‘uname -i’/kernel/unix
```

如果从 ZFS 引导，归档文件和内核文件的路径名都在获选用于引导（如上文所述）的根文件系统（即，数据集）中解析。

通过从内存文件系统装入必要驱动程序和模块继续初始化内核，直至 I/O 可以开启且根文件系统装入完毕。挂载了根文件系统后，将不再需要内存文件系统，并将其放弃。

OpenBoot PROM boot 命令行为

OpenBoot boot 命令采用以下格式的参数：

```
ok boot [device-specifier] [arguments]
```

缺省的 boot 命令不带参数：

```
ok boot
```

如果 boot 命令行上未给定 *device-specifier*，OpenBoot 一般会使用 *boot-device* 或 *diag-device* NVRAM 变量。如果该命令上给定了可选 *arguments*，则 OpenBoot 一般会使用 *boot-file* 或 *diag-file* NVRAM 变量作为缺省的 boot 参数。（如果系统处于诊断模式，则使用 *diag-device* 和 *diag-file*，而不是 *boot-device* 和 *boot-file*）。

arguments 可以包含多个字符串。所有 *argument* 字符串都将传递至辅助引导程序；这些字符串不由 OpenBoot 进行解释。

如果在 boot 命令行上指定了任意 *arguments*，则不会使用 *boot-file* 或 *diag-file* NVRAM 变量。NVRAM 变量的内容不与命令行参数进行合并。例如，命令：

```
ok boot -s
```

将忽略 *boot-file* 和 *diag-file* 中的设置；它会将字符串 "-s" 解释为 *arguments*。boot 将不使用 *boot-file* 或 *diag-file* 的内容。

使用较旧的 PROM，命令：

```
ok boot net
```

不带任何参数，将使用 *boot-file* 或 *diag-file* 中的设置（如果有设置）作为要传递至引导程序的缺省文件名和参数。在多数情况下，最好允许 boot 命令根据系统类型、系统硬件和固件以及根文件系统上安装的内容选择适当的缺省值。在某些情况下，更改 *boot-file* 或 *diag-file* 可能会生成意外的结果。

在多数基于 OpenBoot 2.x 和 3.x 的系统中都可以看到此行为。请注意，在一些平台上可能会有所不同。

命令：

```
ok boot cdrom
```

...通常也不带任何参数。相应地，如果将 *boot-file* 设置为 64 位内核文件名并尝试通过 boot cdrom 引导安装 CD 或 DVD，若安装介质仅包含 32 位内核，那么引导将失败。

因为 *boot-file* 或 *diag-file* 的内容可能会根据所用 boot 命令的格式而被忽略，所以，对于多数生产系统而言，不建议过多依赖于 *boot-file*。

从 wanboot 的本地（CD 或 DVD）副本执行 WAN Boot 时，必须使用：

```
ok boot cdrom -F wanboot - install
```

现代 PROM 已增强了网络引导支持软件包，以支持以下将由该软件包处理的参数的语法：

`[protocol,] [key=value,]*`

所有参数均是可选的，显示顺序随意。除非参数位于列表末尾，否则需要使用逗号。如果进行了指定，参数将优先于任何缺省值，或者，如果使用 DHCP 进行引导，则优先于 DHCP 服务器为那些参数提供的配置信息。

上述的 *protocol* 指定要使用的地址搜索协议。

下列配置参数被指定为 *key=value* 属性对。

tftp-server

TFTP 服务器的 IP 地址

file

针对 WAN Boot 要使用 TFTP 或 URL 下载的文件

host-ip

客户机的 IP 地址（采用点分十进制记法）

router-ip

缺省路由器的 IP 地址

subnet-mask

子网掩码（采用点分十进制记法）

client-id

DHCP 客户机标识符

hostname

DHCP 事务中要使用的主机名

http-proxy

HTTP 代理服务器规范 (IPADDR[:PORT])

tftp-retries

TFTP 重试的最大次数

dhcp-retries

DHCP 重试的最大次数

将由网络引导支持软件包处理的参数列表通过以下两种方式之一指定：

- 将参数传递至该软件包的 `open` 方法，或
- 将参数在 NVRAM 变量 `network-boot-arguments` 中列出。

仅当没有参数传递至该软件包的 `open` 方法时，才会处理 `network-boot-arguments` 中指定的参数。

参数值

protocol 指定要使用的地址搜索协议。如果存在，则可能的值是 *rarp* 或 *dhcp*。

如果其他配置参数以本文档指定的新语法和样式指定，则缺少 *protocol* 参数意味着需要手动配置。

如果没有指定其他配置参数，或那些参数以目前支持的位置参数语法指定，则缺少 *protocol* 参数将导致网络引导支持软件包使用特定于平台的缺省地址搜索协议。

手动配置要求向客户机提供其 IP 地址、引导文件名以及提供引导文件映像的服务器地址。根据网络配置，可能还需要指定 *subnet-mask* 和 *router-ip*。

如果未指定 *protocol* 参数，则网络引导支持软件包将使用特定于平台的缺省地址搜索协议。

tftp-server 是提供要下载文件的 TFTP 服务器（如果使用 TFTP）的 IP 地址（采用标准的 IPv4 点分十进制记法）。

使用 DHCP 时，该值（如果指定）将覆盖 DHCP 响应中指定的 TFTP 服务器的值。

如果将 TFTP RRQ 指定为参数或在 DHCP 响应中进行了指定，则其将单点传送到服务器。否则，TFTP RRQ 将进行广播。

file 指定将由 TFTP 从 TFTP 服务器或 URL（如果使用的是 HTTP）装入的文件。如果文件名是 URL，则将触发 HTTP 的使用，即，文件名以 *http:*（不区分大小写）开头。

使用 RARP 和 TFTP 时，缺省文件名是客户机的 IP 地址（以 ASCII 十六进制表示），如本文档前面的部分所述。

使用 DHCP 时，该参数（如果指定）将覆盖 DHCP 响应中指定的根文件名。

使用 DHCP 和 TFTP 时，缺省文件名基于根节点的 *name* 属性进行构造，其中逗号 (,) 替换为句点 (.)。

在命令行上指定时，文件名不得包含斜线 (/)。

RFC 2396 中对 URL 的格式进行了介绍。HTTP 服务器必须指定为 IP 地址（采用标准的 IPv4 点分十进制记法）。可选端口号以十进制方式指定。如果未指定端口，则默认使用端口 80（十进制）。

提供的 URL 必须是“经过安全编码的”，即，软件包不会向提供的 URL 应用转义编码。包含逗号的 URL 必须以引号括起来的字符串形式出现。否则，引号括起来的 URL 为可选。

host-ip 指定要引导系统的客户机 IP 地址（采用标准的 IPv4 点分十进制记法）。如果使用 RARP 作为地址搜索协议，则指定该参数使得 RARP 的使用没有必要。

如果使用 DHCP，指定 *host-ip* 参数将使客户机遵循要求具有“外部配置的网络地址”的客户机遵循的步骤，如 RFC 2131 中所述。

router-ip 是直接连接的网上的路由器的 IP 地址（采用标准的 IPv4 点分十进制记法）。该路由器将用作通信跨越网络的第一个中继站。如果提供了该参数，则此处指定的路由器将优先于 DHCP 响应中指定的首选路由器。

subnet-mask（采用标准的 IPv4 点分十进制记法）是客户机网络上的子网掩码。如果未提供子网掩码（无论是通过该参数还是 DHCP 响应），则会假设一个缺省掩码，该掩码适合分配给引导客户机的地址的网络类型（A、B 或 C 类）。

client-id 为客户机指定唯一标识符。DHCP 客户机标识符派生自该值。客户机标识符可以指定为：

- 采用 ASCII 十六进制表示的标识符，或
- 引号括起来的字符串

因此，**client-id="openboot"** 和 **client-id=6f70656e626f6f74** 均表示 DHCP 客户机标识符为 6F70656E626F6F74。

命令行上指定的标识符不得包含斜线 (/) 或空格。

DHCP 客户机标识符的最大长度是 32 字节，如果使用 ASCII 十六进制格式，则 64 个字符表示 32 字节。如果使用第二种格式，标识符中的字符数必须是偶数。有效的字符是 0-9、a-f 以及 A-F。

为了正确标识客户机，在客户机连接至的子网上使用的客户机标识符中，客户机标识符必须是唯一的。系统管理员负责选择符合此要求的标识符。

在命令行上指定客户机标识符将优先于任何其他指定标识符的 DHCP 机制。

hostname（指定为字符串）指定要在 DHCP 事务中使用的主机名。该名称可以通过本地域名限定，也可以不通过本地域名限定。主机名的最大长度为 255 个字符。

注 - **hostname** 参数可用于要求客户机向 DHCP 服务器提供所需主机名的服务环境中。客户机向 DHCP 服务器提供所需的主机名，然后，该服务器可以通过 DNS 注册分配给该客户机的主机名和 IP 地址。

http-proxy 采用以下标准记法针对主机进行指定：

host [":" *port*]

...其中 **host** 指定为 IP 地址（采用标准的 IPv4 点分十进制记法），可选的 *port* 采用十进制指定。如果未指定端口，则默示使用端口 8080（十进制）。

tftp-retries 是确定 TFTP 进程失败之前重试的最大次数（采用十进制指定）。缺省设置为可进行无限次重试。

dhcp-retries 是确定 DHCP 进程失败之前重试的最大次数（采用十进制指定）。缺省设置为可进行无限次重试。

x86 引导过程

在基于 x86 的系统上，引导过程包含两个概念截然不同的阶段，即内核装入和内核初始化。内核装入在使用系统板上的 BIOS ROM 以及外围板上 ROM 中的 BIOS 扩展的 GRUB (GRand Unified Bootloader) 中实现。BIOS 装入 GRUB，从硬盘、DVD 或 CD 上的第一个物理扇区开始。如果受网络适配器上的 ROM 支持，BIOS 还可从网络引导服务器下载 pxegrub 二进制文件。找到了 GRUB 后，将执行菜单中的命令以装入 unix 内核和包含内核模块与数据的预构造引导归档文件。

如果由 GRUB 标识为引导设备的相应设备包含 ZFS 存储池，则用于创建 GRUB 菜单的 `menu.lst` 文件将位于该池的数据集分层结构的根数据集中。这是一个数据集，其名称与池自身的名称相同。池中正好始终有这样一个数据集存在，因此，此数据集适用于池范围的数据，如 `menu.lst` 文件。引导了系统后，该数据集将挂载在根文件系统下的 `/poolname`。

池中可以有多个可引导数据集（即，根文件系统）。缺省情况下，`menu.lst` 文件中的文件名条目获得解析的文件系统由池的 `bootfs` 属性标识（请参见 [zpool\(1M\)](#)）。但是，`menu.lst` 项可以包含 `bootfs` 命令，该命令可指定池中的备用数据集。这样，`menu.lst` 文件就可以包含池中多个根文件系统的项。

内核初始化在 GRUB 完成装入引导归档文件并将控制权移交至 unix 二进制文件后开始。此时，GRUB 将停止活动，引导设备上不会再出现更多 I/O。Unix 操作系统针对根文件系统进行初始化、在引导归档文件的必要模块中进行链接以及在实际的根设备上挂载。此时，内核重新获取存储 I/O，挂载其他文件系统（请参见 [vfstab\(4\)](#)），并启动各种操作系统服务（请参见 [smf\(5\)](#)）。

启用自动重新引导
(x86)

Solaris 操作系统支持 [smf\(5\)](#) 属性，通过该属性，可使系统自动从当前引导设备进行重新引导，以从诸如过时引导归档文件之类的状态中恢复。

服务 `svc:/system/boot-config:default` 包含布尔属性 `auto-reboot-safe`，该属性缺省设置为 `false`。将其设置为 `true` 表明系统的 BIOS 以及缺省 GRUB 菜单项均设置为从当前引导设备引导。可以使用 [svccfg\(1M\)](#) 和 [svcadm\(1M\)](#) 更改该属性的值。例如，要设置 `auto-reboot-safe` 以启动自动重新引导，请输入诸如以下所示的命令：

```
example# svccfg -s svc:/system/boot-config:default \
      setprop config/auto-reboot-safe = true
```

多数系统均配置为从当前引导设备进行自动重新引导。然而，在某些情况下，自动重新引导至未知操作系统可能会产生意外结果。对于这些情况，`auto-reboot-safe` 属性允许用户指定希望的行为。

故障安全模式

若要从内置于引导归档文件的根文件系统映像引导，而后将根重新挂载至实际的根设备，引导归档文件和根文件系统的内容必须一致。否则，无法保证计算机的正常运行与完整性。

术语“一致”意即根文件系统的所有文件和模块都在引导归档文件中并具有相同的内容。由于引导策略要求首先将引导归档文件读取并挂载为第一阶段根映像，因此，所有可卸载内核模块和初始化（派生自引导归档文件的内容）均要求与实际根文件系统

匹配。如果缺乏一致性，系统运行可能会出现这样的情况：内核模块或参数设置在重新引导之前应用至根设备，但并未在根归档文件中更新。这种不一致性可能导致系统不稳定或数据丢失。

挂载了根文件系统后，在放弃内存文件系统之前，Solaris 将针对两个文件系统执行一致性验证。如果检测到不一致，Solaris 将暂停常规引导序列并回退至故障安全模式。纠正此状态需要管理员采取以下两个步骤之一。推荐的做法是重新引导至故障安全归档文件并重新构建引导归档文件。这将确保引导已知内核并在归档文件重建过程中发挥作用。或者，管理员可以选择清除不一致的引导归档文件服务状态并继续系统初启（如果该不一致性不会影响系统正常运行）。请参见 [svcadm\(1M\)](#)。

如果清除了引导归档文件服务并继续系统初启（上述第二种备选方法），则系统运行时，可卸载内核驱动程序或其他模块可能相对于根文件系统已过时。因而，可能会影响系统正常运行。

为了确保引导归档文件是一致的，标准系统关机过程（由 [reboot\(1M\)](#) 和 [shutdown\(1M\)](#) 启动）将检查更新并将其应用至引导归档文件（在 [umountall\(1M\)](#) 结束后）。

对于任何需要在 [umountall](#) 服务完成后包含在引导归档文件中的内核文件、驱动程序、模块或驱动程序配置文件，对其进行更新都将导致下一次引导时检查引导归档文件一致性失败。为了避免这种情况，建议始终彻底关闭计算机。

如果要求在完成了 [umountall](#) 服务后对内核进行更新，则管理员可以选择通过调用以下内容重新构建归档文件：

```
# bootadm update-archive
```

故障安全引导归档文件

故障安全归档文件可用于随时引导计算机以进行维护或故障排除。故障安全引导归档文件安装在计算机上，源自 [miniroot](#) 归档文件。引导故障安全归档文件将使计算机使用内存文件系统作为根设备进行引导。

SPARC

SPARC 故障安全归档文件是：

```
/platform/'uname -i'/failsafe
```

...可按如下进行引导：

```
ok boot [device-specifier] -F failsafe
```

如果用户希望从特定 ZFS 可引导数据集引导故障安全归档文件，则可使用如下命令来完成：

```
ok boot [device-specifier] -Z dataset -F failsafe
```

x86

x86 故障安全归档文件是：

/boot/x86.miniroot-safe

...可通过从 GRUB 菜单选择 Solaris failsafe（Solaris 故障安全）项进行引导。

选项

SPARC

以下 SPARC 选项受支持：

-a

引导程序将该标志解释为 *ask me*，因此会提示提供独立程序的名称。然后，'-a' 标志将传递至该独立程序。

-D *default-file*

显式指定 *default-file*。在某些系统上，boot 选择动态缺省文件，前提是其他任何地方都没有指定相应文件。该选项允许显式设置 *default-file*，可在引导 [kmdb\(1\)](#) 时使用，因为缺省情况下，kmdb 将在 *default-file* 由 boot 程序导出时对其进行装入。

-F *object*

使用命名对象引导。该对象必须是 ELF 可执行文件或包含引导块的可引导对象。主要用途是引导故障安全或 wanboot 引导归档文件。

-L

列出 ZFS 池内的可引导数据集。可以在该列表中选择一个可引导数据集，其后将显示引导该数据集的详细说明。按照以下说明引导选定的数据集。该选项只有在引导设备包含 ZFS 存储池时才受支持。

-V

显示详细的调试信息。

boot-flags

引导程序将所有 *boot-flags* 传递至 *file*。它们不由 boot 进行解释。有关随缺省独立程序提供的选项的信息，请参见 [kernel\(1M\)](#) 和 [kmdb\(1\)](#) 手册页。

client-program-args

boot 程序将所有 *client-program-args* 传递至 *file*。它们不由 boot 进行解释。

file

要引导的独立程序的名称。如果文件名并未显式指定（无论是在 boot 命令行上还是在 *boot-file* NVRAM 变量中），boot 将选择适当的缺省文件名。

OBP names

指定 Open Boot Prom 的名称。例如，在基于 Desktop SPARC 的系统上，名称 */sbus/esp@0,800000/sd@3,0:a* 意指位于 SCSI 总线上目标 3 lun0 的 SCSI 磁盘 (sd)，esp 主机适配器插入插槽 0。

-Z *dataset*

在指定的 ZFS 数据集中从根文件系统引导。

x86

以下 x86 选项受支持：

-B *prop=val...*

要传递至内核的一个或多个属性-值对。多个属性-值对必须用逗号分开。使用该选项等同于命令：`eeeprom prop=val`。有关可用属性和有效值，请参见 [eeeprom\(1M\)](#)。

如果与此菜单项对应的根文件系统是 ZFS 数据集，则此菜单项需要添加以下选项：

-B *\$ZFS-BOOTFS*

boot-args

引导程序将所有 *boot-args* 传递至 *file*。它们不由 *boot* 进行解释。有关随内核提供的选项的信息，请参见 [kernel\(1M\)](#) 和 [kmdb\(1\)](#)。

除非另有指定，否则 x86 系统都将引导 `/platform/i86pc/kernel/amd64/unix`。

x86 引导序列详细信息

打开与 PC 兼容的计算机后，BIOS ROM 中的系统固件将执行开机自检 (power-on self test, POST)，运行外围板 ROM 中的 BIOS 扩展，并调用软件中断 INT 19h (引导)。INT 19h 处理程序通常执行标准的与 PC 兼容的引导，其中包含尝试从第一个硬盘读取第一个物理扇区。然后，处理器将跳至内存中该扇区映像的第一个字节。

x86 主引导

磁盘介质上的第一个扇区包含主引导块 (GRUB stage1)。stage1 负责装入 GRUB stage2。现在，GRUB 功能齐全了。它将读取并执行菜单文件 `/boot/grub/menu.lst`。类似序列也会在 DVD 或 CD 引导中出现，但主引导块位置和内容由 El Torito 规范确定。El Torito 引导还会导致 `strap.com`，进而装入 `boot.bin`。

硬盘上的第一个扇区包含主引导块，其中包含主引导程序和 FDISK 表，针对维护其的 PC 程序命名。主引导将在 FDISK 表中查找活动分区，装入其第一个扇区 (GRUB stage1)，并跳至其在内存中第一个字节。这将完成标准的与 PC 兼容的硬盘引导序列。如果 GRUB stage1 安装在主引导块上 (请参见 [installgrub\(1M\)](#) 的 `-m` 选项)，则 stage2 将直接从 Solaris FDISK 分区进行装入，与活动分区无关。

针对 Solaris 软件的 x86 FDISK 分区以一个柱面引导分片开始，在第一个扇区中包含 GRUB stage1，在第二个和第三个扇区中包含标准的 Solaris 磁盘标号和卷目录表 (volume table of content, VTOC)，在第五十个和后续扇区中包含 GRUB stage2。从扇区 4 到 49 的区域包含适用于较旧版本的 Solaris 的引导块。这使得在同一 FDISK 上可以同时存在多个 Solaris 版本。当针对 Solaris 软件的 FDISK 分区是活动分区时，主引导程序 (mboot) 会将第一个扇区中的分区引导程序读入到内存中并跳至该位置。然后，将 GRUB stage2 程序读入到内存中并跳至该位置。显示 GRUB 菜单后，用户可以选择在不同的分区、不同的磁盘上或可以从网络引导操作系统。

对于网络引导，受支持的方法是 Intel 的 Preboot eXecution Environment (PXE) 标准。使用 PXE 从网络引导时，系统或网络适配器 BIOS 将使用 DHCP 定位引导服务器上的网络引导程序 (pxegrub)，并使用简单文件传输协议 (Trivial File Transfer Protocol, TFTP) 对其进行读取。BIOS 通过跳至其内存中的第一个字节执行 pxegrub。pxegrub 程序将下载菜单文件并将菜单项显示给用户。

x86 内核启动

内核启动过程独立于内核装入过程。内核启动期间，控制台 I/O 将转至由 `console` 属性指定的设备。

从 UFS 引导时，根设备由 `bootpath` 属性指定，根文件系统类型由 `fstype` 属性指定。这些属性应由 `/boot/solaris/bootenv.rc` 中的 Solaris Install/Upgrade 进程设置，并可由 `-B` 选项覆盖，如上文所述（请参见 [eeprom\(1M\)](#) 手册页）。

从 ZFS 引导时，根设备由引导参数指定，而该参数由 GRUB 菜单项中的 `kernel` 或 `module` 行上的 `-B $ZFS-BOOTFS` 参数指定。该值（与由 `-B` 选项指定的所有参数一样）由 GRUB 传递给内核。

如果未提供控制台属性，则控制台 I/O 缺省为 `screen`（屏幕）和 `keyboard`（键盘）。根设备缺省为 `ramdisk`（RAM 磁盘），文件系统缺省为 `ufs`。

示例**SPARC**

示例 1 在单用户交互模式中引导缺省内核

要在单用户交互模式中引导缺省内核，请使用以下项之一响应 `ok` 提示符：

```
boot -as
```

```
boot disk3 -as
```

示例 2 通过支持 WAN Boot 的 PROM 进行网络引导

为了说明各种引导命令行调用的一些细微反响，假设设置了 `network-boot-arguments`，通过 `devalias` 命令设置了 `net`，如以下命令所示。

在以下命令中，设备别名中的设备参数由该设备驱动程序处理。网络引导支持软件包处理 `network-boot-arguments` 中的参数。

```
boot net
```

以下命令将不会生成任何设备参数。网络引导支持软件包处理 `network-boot-arguments` 中的参数。

```
boot net:
```

以下命令将不会生成任何设备参数。`rarp` 是唯一的网络引导支持软件包参数。`network-boot-arguments` 将被忽略。

```
boot net:rarp
```

在以下命令中，指定的设备参数会被接受。网络引导支持软件包处理 `network-boot-arguments` 中的参数。

```
boot net:speed=100,duplex=full
```

示例 3 结合使用 wanboot 与较旧的 PROM

以下命令将生成要从 DVD 或 CD 装入的 wanboot 二进制文件，此时 wanboot 将执行 DHCP，然后进入其命令解释程序以允许用户输入密钥和任何其他必要配置。

boot cdrom -F wanboot -o dhcp,prompt

x86

示例 4 在单用户交互模式中引导缺省内核

要在单用户交互模式中引导缺省内核，请编辑 GRUB 内核命令行以读取：

kernel /platform/i86pc/kernel/amd64/unix -as

文件

/etc/inittab

其中指定了 initdefault 状态的表

/usr/sbin/init

使系统进入 initdefault 状态的程序

仅 SPARC

/platform/*platform-name*/kernel/sparcv9/unix

要引导系统的缺省程序。

仅限于 x86

/Boot

包含与引导相关的文件的目录。

/rpool/boot/grub/menu.lst

由 GRUB 显示的可引导操作系统菜单。/rpool 是一个通用约定。路径名是可配置的，具体取决于安装程序的功能。

/platform/i86pc/kernel/amd64/unix

要引导系统的缺省程序。

另请参见

[kmdb\(1\)](#)、[uname\(1\)](#)、[bootadm\(1M\)](#)、[eeprom\(1M\)](#)、[init\(1M\)](#)、[installboot\(1M\)](#)、[kernel\(1M\)](#)、[monit](#)

RFC 903 A Reverse Address Resolution Protocol (《反向地址解析协议》)，<http://www.ietf.org/rfc/rfc903.txt>

RFC 2131 Dynamic Host Configuration Protocol (《动态主机配置协议》)，<http://www.ietf.org/rfc/rfc2131.txt>

RFC 2132 DHCP Options and BOOTP Vendor Extensions (《DHCP 选项和 BOOTP 供应商扩展》)，<http://www.ietf.org/rfc/rfc2132.txt>

RFC 2396 Uniform Resource Identifiers (URI): Generic Syntax (《统一资源标识符 (Uniform Resource Identifier, URI)：通用语法》)，<http://www.ietf.org/rfc/rfc2396.txt>

《Oracle Solaris 管理：常见任务》

《Sun 硬件平台指南》

《OpenBoot 命令参考手册》

警告

`boot` 实用程序无法确定可以用作可引导程序的文件。如果请求引导不可引导的文件，`boot` 实用程序将装入并进一步处理该文件。之后会发生什么不可预知。

附注

可以使用 `uname(1)` 的 `-i` 选项查找 *platform-name*。可以使用 `uname(1)` 的 `-m` 选项查找 *hardware-class-name*。

Solaris 操作系统的当前发行版不支持运行 UltraSPARC-I CPU 的计算机。

引用名 bootadm – 管理启用了 GRUB 的操作系统的可引导性

用法概要

```
/usr/sbin/bootadm update-archive [-vn] [-R altroot [-p platform]]
/usr/sbin/bootadm list-archive [-vn] [-R altroot [-p platform]]
x86 only
/usr/sbin/bootadm set-menu [-R altroot [-p platform]] key=value
/usr/sbin/bootadm list-menu [-R altroot [-p platform]]
```

描述

bootadm 命令用于管理引导归档文件以及 x86 引导环境中的 GRUB (GRand Unified Bootloader) 菜单。**update-archive** 选项提供了一种用于用户更新引导归档文件的方法，作为恢复过程的一种预防措施或一部分。**set-menu** 子命令可用于切换 GRUB 菜单中的 **auto-boot** 超时以及缺省的引导项。

list-menu 子命令显示 GRUB 菜单的位置以及当前 GRUB 菜单项。通常，GRUB 菜单在 ZFS 的位置为 `/rpool/boot/grub/menu.lst`。安装方法不同，活动的 GRUB 菜单的位置可能也不同。可以使用 **list-menu** 子命令定位活动的 GRUB 菜单。例如，如果系统是使用 Live Upgrade 进行安装的，GRUB 菜单可能没有在当前引导环境中。有关 **list-menu** 选项的典型输出，请参见“示例”部分。

请注意，基于 OpenBoot PROM (OBP) 的计算机（如 SPARC 系统），不使用 GRUB，也不存在可由 **bootadm** 管理的引导菜单。

bootadm 命令动态确定要管理的映像支持的选项。这样，平台上调用的 **bootadm** 命令可用于管理其他平台类型的无盘客户机。

子命令

bootadm 命令包含以下子命令：

update-archive

根据需要更新当前引导归档文件。适用于 SPARC 和 x86 平台。

list-archive

列出引导归档文件中包含的文件以及目录。适用于 SPARC 和 x86 平台。

set-menu

维护 GRUB 菜单。当前 GRUB 菜单为 `boot/grub/menu.lst`（相对于根）。该位置随时可能变动，因此，不要依赖该位置。仅适用于 x86 平台。

list-menu

列出活动 GRUB 菜单的位置，以及当前的 GRUB 菜单项。包括自动引导超时、缺省的菜单项数量以及每个菜单项的标题。仅适用于 x86 平台。

选项

bootadm 命令包含以下选项：

-v

在 **update-archive** 操作中，过时的文件显示在标准错误中。

-n

在 **update-archive** 操作中，只检查归档文件内容但不更新归档文件内容。

-p *platform*

该客户机的平台或计算机硬件类型。该平台类型只能与 **-R** 一起指定，并且通常只用于管理与服务器不属于同一平台类型的无盘客户机。平台必须为 **i86pc**、**sun4u** 或 **sun4v** 之一。

-R *altroot*

操作适用于一个备用的根路径。

注 – 不得使用 **-R** 选项引用任何非全局区域的根文件系统。否则，可能会损坏全局区域的文件系统，危及全局区域的安全性，或者损坏非全局区域的文件系统。请参见 [zones\(5\)](#)。

key=value

可能的值为：

default=entrynum

GRUB 菜单中的菜单项编号（例如，0、1 或 2），用于指定当计时器超时时要引导的操作系统。

timeout=seconds

由缺省的菜单项编号指定的操作系统进行引导之前等待的时间（以秒为单位）。如果值为 -1，表明禁用自动引导。

示例

示例1 更新当前的引导归档文件

以下命令更新当前的引导归档文件：

```
# bootadm update-archive
```

示例2 更新备用根中的引导归档文件

以下命令更新备用根下的引导归档文件：

```
# bootadm update-archive -R /a
```

示例3 列出安装的操作系统实例

以下命令列出 GRUB 菜单中的安装的操作系统实例：

```
# bootadm list-menu
```

```
default=0
```

```
timeout=10
```

```
(0) Solaris10
```

```
(1) Solaris10 Failsafe
```

```
(2) Linux
```

示例4 切换缺省的引导项

以下命令指向上一示例中显示的菜单。用户选择 **Linux**（菜单项 2）。

```
# bootadm set-menu default=2
```

示例 5 列出 GRUB 菜单项以及 GRUB 菜单的位置

以下命令列出 GRUB 菜单项以及 GRUB 菜单的位置：

```
# bootadm list-menu
The location for the active GRUB menu is: /stubboot/boot/grub/menu.lst
default 0
timeout 10
0 Solaris10
1 Solaris10 failsafe
2 Linux
```

示例 6 显示 GRUB 菜单的位置

以下命令显示 GRUB 菜单的位置：

```
# bootadm list-menu
The location for the active GRUB menu is: /dev/dsk/c0t1d0s0 (not mounted)
The filesystem type of the menu device is <ufs>
default 2
timeout 10
0 c0t1d0s3
1 c0t1d0s3 failsafe
2 Solaris10
3 Solaris10 failsafe
```

在该示例中，活动的 GRUB 菜单位于未挂载的设备上。要访问 GRUB 菜单，请挂载该设备并在 <mountpoint>/boot/grub/menu.lst 中访问 GRUB 菜单。

退出状态

将返回下列退出值：

- 0 命令执行成功。
- 1 命令因运行出错而退出。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）

另请参见

[boot\(1M\)](#)、[installgrub\(1M\)](#)、[attributes\(5\)](#)

登录以下网址查阅 GRUB 主页：

<http://www.gnu.org/>

引用名	bootconfchk – verify the integrity of a network boot configuration file
用法概要	<code>/usr/sbin/bootconfchk [bootconf-file]</code>
描述	The <code>bootconfchk</code> command checks that the file specified is a valid network boot configuration file as described in wanboot.conf(4). Any discrepancies are reported on standard error.
退出状态	0 Successful completion. 1 An error occurred. 2 Usage error.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/boot/wanboot
Interface Stability	Committed

另请参见 [wanboot.conf\(4\)](#), [attributes\(5\)](#)

`-r device-inst`

Read and display all pic values for the specified device

Specify *device-inst* as *device* (name) followed by *instance number*, if specifying an instance number of a device whose counters are to be read and displayed. If all instances of this device are to be read, use *device* (name) without an instance number. All pic values will be sampled when using the `-r` option.

`-w device-inst [,pic0=event] [,picn=event]`

Program (write) the specified devices to count the specified events. Write access to the counters is restricted to root users only. Non-root users can use `-r` option.

Specify *device-inst* as *device* (name) followed by an optional *instance number*. If specifying an instance number of a device to program these events on. If all instances of this device are to be programmed the same, then use *device* without an instance number. Specify an event to be counted for a specified pic by providing a comma separated list of *picn=event* values.

The `-e` option displays all valid event names for each device. Any devices that are programmed will be sampled every interval seconds and repeated count times. It is recommended that the interval specified is small enough to ensure that counter wraparound will be detected. The rate at which counters wraparound varies from device to device. If a user is programming events using the `-w` option and `busstat` detects that another user has changed the events that are being counted, the tool will terminate as the programmed devices are now being controlled by another user. Only one user can be programming a device instance at any one time. Extra devices can be sampled using the `-r` option. Using multiple instances of the `-w` option on the same command line, with the same *device-inst* specifying a different list of events for the pics will give the effect of multiplexing for that device. `busstat` will

switch between the list of events for that device every interval seconds. Event can be a string representing the event name, or even a number representing the bit pattern to be programmed into the Performance Control Register (PCR). This assumes explicit knowledge of the meaning of the control register bits for a device. The number can be specified in hexadecimal, decimal, or octal, using the usual conventions of `strtol(3C)`.

退出状态

The following exit values are returned:

- 0 Successful completion.
- 1 An error occurred.
- 2 Another user is writing to the same devices.

示例

SPARC Only

示例 1 Programming and monitoring the Address Controller counters

In this example, `ac0` refers to the Address Controller instance 0. The counters are programmed to count Memory Bank stalls on an Ultra Enterprise system at 10 second intervals with the values displayed in absolute form instead of deltas.

```
# busstat -a -w ac0,pic0=mem_bank0_stall,pic1=mem_bank1_stall 10
time dev  event0          pic0  event1          pic1
10  ac0  mem_bank0_stall      1234  mem_bank1_stall  5678
20  ac0  mem_bank0_stall      5678  mem_bank1_stall  12345
30  ac0  mem_bank0_stall     12345  mem_bank1_stall  56789
...
```

For a complete list of the supported events for a device, use the `-e` option.

示例 2 Programming and monitoring the counters on all instances of the Address Controller

In this example, `ac` refers to all `ac` instances. This example programs all instances of the Address Controller counters to count `_clock` cycles and `mem_bank0_rds` at 2 second intervals, 100 times, displaying the values as deltas.

```
# busstat -w ac,pic0=clock_cycles,pic1=mem_bank0_rds 2 100
time dev  event0          pic0  event1          pic1
2  ac0  clock_cycles      167242902  mem_bank0_rds    3144
2  ac1  clock_cycles      167254476  mem_bank0_rds    1392
4  ac0  clock_cycles      168025190  mem_bank0_rds    40302
4  ac1  clock_cycles      168024056  mem_bank0_rds    40580
...
```

示例 3 Monitoring the events being counted

This example monitors the events that are being counted on the sbus1 device, 100 times at 1 second intervals. It suggests that a root user has changed the events that sbus1 was counting to be dvma_tlb_misses and interrupts instead of pio_cycles.

```
% busstat -r sbus0 1 100
```

time	dev	event0	pic0	event1	pic1
1	sbus1	pio_cycles	2321	pio_cycles	2321
2	sbus1	pio_cycles	48	pio_cycles	48
3	sbus1	pio_cycles	49	pio_cycles	49
4	sbus1	pio_cycles	2281	pio_cycles	2281
5	sbus1	dvma_tlb_misses	0	interrupts	0
6	sbus1	dvma_tlb_misses	6	interrupts	2
7	sbus1	dvma_tlb_misses	8	interrupts	11
...					

示例 4 Event Multiplexing

This example programs ac0 to alternate between counting (clock cycles, mem_bank0_rds) and (addr_pkts, data_pkts) at 2 second intervals while also monitoring what ac1 is counting :

It shows the expected output of the above busstat command. Another root user on the machine has changed the events that this user had programmed and busstat has detected this and terminates the command with a message.

```
# busstat -w ac0,pic0=clock_cycles,pic1=mem_bank0_rds \
-w ac0,pic0=addr_pkts,pic1=data_pkts \
-r ac1 2
```

time	dev	event0	pic0	event1	pic1
2	ac0	addr_pkts	12866	data_pkts	17015
2	ac1	rio_pkts	385	rio_pkts	385
4	ac0	clock_cycles	168018914	mem_bank0_rds	2865
4	ac1	rio_pkts	506	rio_pkts	506
6	ac0	addr_pkts	144236	data_pkts	149223
6	ac1	rio_pkts	522	rio_pkts	522
8	ac0	clock_cycles	168021245	mem_bank0_rds	2564
8	ac1	rio_pkts	387	rio_pkts	387
10	ac0	addr_pkts	144292	data_pkts	159645
10	ac1	rio_pkts	506	rio_pkts	506
12	ac0	clock_cycles	168020364	mem_bank0_rds	2665
12	ac1	rio_pkts	522	rio_pkts	522

busstat: events changed (possibly by another busstat).
#

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[iostat\(1M\)](#), [mpstat\(1M\)](#), [vmstat\(1M\)](#), [strtol\(3C\)](#), [attributes\(5\)](#)

引用名	captoinfo – convert a termcap description into a terminfo description				
用法概要	captoinfo [-l] [-v]... [-V] [-w width] filename...				
描述	captoinfo looks in <i>filename</i> for termcap descriptions. For each one found, an equivalent terminfo description is written to standard output, along with any comments found. A description which is expressed as relative to another description (as specified in the termcap <code>tc = field</code>) is reduced to the minimum superset before being displayed. If no <i>filename</i> is given, then the environment variable TERMCAP is used for the filename or entry. If TERMCAP is a full pathname to a file, only the terminal whose name is specified in the environment variable TERM is extracted from that file. If the environment variable TERMCAP is not set, then the file <code>/usr/share/lib/termcap</code> is read.				
选项	–l Display the fields one to a line. Otherwise, the fields are printed several to a line, with a maximum width of 60 characters. –v Display tracing information on the standard error as the program runs. Specifying additional –v options displays more detailed information. –V Display the version of the program in use on the standard error and then exit. –w width Change the output to <i>width</i> characters.				
文件	<code>/usr/share/lib/terminfo/??/*</code> compiled terminal description database <code>/usr/share/lib/termcap</code>				
属性	See attributes(5) for descriptions of the following attributes:				
	<table><tr><th>ATTRIBUTETYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>system/core-os</td></tr></table>	ATTRIBUTETYPE	ATTRIBUTE VALUE	Availability	system/core-os
ATTRIBUTETYPE	ATTRIBUTE VALUE				
Availability	system/core-os				
另请参见	infocmp(1M) , curses(3CURSES) , terminfo(4) , attributes(5)				
附注	captoinfo should be used to convert termcap entries to terminfo entries because the termcap database may not be supplied in future releases.				

引用名

catman – create the formatted files for the reference manual

用法概要

```
/usr/bin/catman [-c] [-n] [-p] [-t] [-w] [-M directory]  
                [-T macro-package] [sections]  
  
/usr/bin/catman [-M directory] -w
```

描述

The catman utility creates the preformatted versions of the on-line manual from the [nroff\(1\)](#) or [sgml\(5\)](#) input files. This feature allows easy distribution of the preformatted manual pages among a group of associated machines, since it makes the directories of preformatted manual pages self-contained and independent of the unformatted entries.

With the -w option, catman also creates index files, in the directories specified by the MANPATH or the -M option. If there is no MANPATH or -M option specified, unless -n is specified, catman creates index files at the /usr/share/man/ and /usr/gnu/share/man/ directories by default. When any specified or default directory is read-only, catman fails and displays an error message to stderr, indicating that writing is not allowed to the directory.

Each manual page is examined and those whose preformatted versions are missing or out of date are recreated. If any changes are made, catman recreates the index files.

If a manual page is a *shadow* page, that is, it sources another manual page for its contents, a symbolic link is made in the catx or fmtx directory to the appropriate preformatted manual page.

Shadow files in an unformatted nroff source file are identified by the first line being of the form `.so manx/yyy.x`.

Shadow files in the SGML sources are identified by the string SHADOW_PAGE. The file entity declared in the shadow file identifies the file to be sourced.

选项

The following options are supported:

-c

Create unformatted nroff source files in the appropriate man subdirectories from the SGML sources. This option will overwrite any existing file in the man directory of the same name as the SGML file.

-n

Do not create (or recreate) the index files. If the -n option is specified, the index files are not created and the [apropos\(1\)](#) and [whatis\(1\)](#) commands might run more slowly than otherwise.

-p

Dry—run option. That is, display what would be done instead of doing it.

-t

Create troffed entries in the appropriate fmt subdirectories instead of nroffing into the cat subdirectories.

-w

Create the index files that are used by [apropos\(1\)](#), [whatis\(1\)](#) and the [man\(1\)](#) -f, -k, and -K options, in the directories specified by the MANPATH environment variable or the -M option. If no MANPATH or -M option is specified, index files are created in the /usr/share/man/ and /usr/gnu/share/man/ directories by default. No manual reformatting is done.

-M *directory*

Update manual pages located in the specified *directory*, (/usr/share/man by default). If the -M option is specified, the directory argument must not contain a ',' (comma), since a comma is used to delineate section numbers. See [man\(1\)](#).

-T *macro-package*

Use *macro-package* in place of the standard manual page macros, ([man\(5\)](#) by default).

操作数

The following operand is supported:

sections

If there is one parameter not starting with a '-', it is taken to be a space separated list of manual sections to be processed by catman. If this operand is specified, only the manual sections in the list will be processed. For example,

```
catman 1 2 3
```

only updates manual sections 1, 2, and 3. If specific sections are not listed, all sections in the man directory specified by the environment variable MANPATH are processed.

环境变量

TROFF

The name of the formatter to use when the -t flag is given. If not set, [troff\(1\)](#) is used.

MANPATH

A colon-separated list of directories that are processed by catman and [man\(1\)](#). Each directory can be followed by a comma-separated list of sections. If set, its value overrides /usr/share/man as the default directory search path, and the man.cf file as the default section search path. The -M and -s flags, in turn, override these values.

示例

示例1 Creating an Index File

The following command creates an index file in the /usr/local/share/man directory.

```
# catman -M /usr/local/share/man -w
```

文件

```
/usr/share/man
    default manual directory location

/usr/share/man/man*/*.*
    raw nroff input files

/usr/share/man/sman*/*.*
    raw SGML input files

/usr/share/man/cat*/*.*
    preformatted nroffed manual pages
```

`/usr/share/man/fmt*/*.*`
preformatted troffed manual pages

`/usr/share/lib/tmac/an`
default macro package

`/usr/share/man/man_index/*`
`/usr/share/man/man_index/man.idx`
`/usr/share/man/man_index/man.dic`
`/usr/share/man/man_index/man.frq`
`/usr/share/man/man_index/man.pos`
index file for -K query

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	text/doctools
CSI	Enabled
Interface Stability	Committed

另请参见 [apropos\(1\)](#), [man\(1\)](#), [nroff\(1\)](#), [rm\(1\)](#), [troff\(1\)](#), [whatis\(1\)](#), [attributes\(5\)](#), [man\(5\)](#), [sgml\(5\)](#)

诊断 `man?/xxx.? (.so'ed from man?/yyy.?): No such file or directory`
The file outside the parentheses is missing, and is referred to by the file inside them.

`target of .so in man?/xxx.? must be relative to /usr/man`
catman only allows references to filenames that are relative to the directory `/usr/man`.

`opendir:man?: No such file or directory`
A harmless warning message indicating that one of the directories catman normally looks for is missing.

`*.*: No such file or directory`
A harmless warning message indicating catman came across an empty directory.

警告 If a user, who has previously run catman to install the cat* directories, upgrades the operating system, the entire cat* directory structure should be removed prior to running catman. See [rm\(1\)](#).

Do not re-run catman to rebuild the index files unless the complete set of man* directories is present. catman builds the index files based on the man* directories.

附注 The windex database has been replaced by index files. Unlike the case with windex, index file generation does not pose any specific restrictions or prerequisites on what can be indexed.

引用名	cfgadm – 配置管理
用法概要	<pre> /usr/sbin/cfgadm [-f] [-y -n] [-v] [-o hardware_options] -c function ap_id... /usr/sbin/cfgadm [-f] [-y -n] [-v] [-o hardware_options] -x hardware_function ap_id... /usr/sbin/cfgadm [-v] [-a] [-s listing_options] [-o hardware_options] [-l [ap_id ap_type]] /usr/sbin/cfgadm [-v] [-o hardware_options] -t ap_id... /usr/sbin/cfgadm [-v] [-o hardware_options] -h [ap_id ap_type] </pre>

描述 `cfgadm` 命令提供了针对动态可重新配置的硬件资源的配置管理操作。这些操作包括显示状态 (-l)、启动测试 (-t)、调用配置状态更改 (-c)、调用特定于硬件的函数 (-x) 以及获取配置管理帮助消息 (-h)。配置管理操作在**挂接点**上执行，这些位置的系统软件支持在继续运行 Solaris 的过程中动态重新配置硬件资源。

配置管理会区分计算机上实际存在的硬件资源和对 Solaris 可见的已配置硬件资源。配置管理函数本性是特定于硬件的，可通过调用特定于硬件的库来执行。

配置管理在**挂接点**上运行。位于挂接点处的硬件资源在系统运行期间可能是物理上可替换的，也可能是物理上不可替换的，但可通过配置管理界面动态重新配置。

挂接点定义了两个唯一的元素，这些元素不同于位于挂接点之外的硬件资源。挂接点的两个元素是**插口**和**插卡**。物理插入或移除硬件资源通常在挂接点上执行，这些操作会导致获得插口或失去插卡。配置管理支持在挂接点上执行物理插入和移除操作以及其他配置管理函数。

挂接点具有相关联的状态和状况信息。配置管理界面可用于控制挂接点状态的转换。插口可按以下三种状态之一的形式存在：`empty`（空）、`disconnected`（断开）或 `connected`（已连接），而插卡则按以下两种状态之一的形式存在：`configured`（已配置）或 `unconfigured`（未配置）。

插口可提供 `empty`（空）状态，这是挂接点不包含任何插卡时插口的常规状态。如果插口具有将其插卡与常规系统访问隔离的功能，则还可以提供 `disconnected`（断开）状态。通常情况下，此状态用于在系统将插卡资源完全投入使用之前进行各种特定于硬件的测试，或作为准备物理移除或重新配置插卡的步骤之一。处于 `disconnected` 状态的插口会尽量在硬件允许的范围内将其插卡与系统隔离开，但也会提供针对测试和设置的访问。插口必须提供 `connected`（已连接）状态，这样才能允许对插卡内包含的硬件资源进行常规访问。对于包含插卡且当前未执行配置管理操作的插口，`connected`（已连接）状态是其常规状态。

处于 `unconfigured`（未配置）状态的插卡上所包含的硬件资源不会用常规 Solaris 数据结构表示，因此不能供 Solaris 使用。未配置插卡上允许执行的操作仅限于配置管理操作。处于 `configured`（已配置）状态的插卡的硬件资源用常规 Solaris 数据结构表

示，因此其中部分或全部硬件资源均可供 Solaris 使用。所有插卡都提供 **configured**（已配置）和 **unconfigured**（未配置）两种状态。

挂接点可处于以下五种状况之一：**unknown**（未知）、**ok**（正常）、**failing**（即将失败）、**failed**（失败）或 **unusable**（不可用）。挂接点可在任何状况下进入系统，具体取决于加电自检和非易失性记录保存的结果。

包含处于 **configured**（已配置）状态的插卡的挂接点处于以下四种状况之一：**unknown**（未知）、**ok**（正常）、**failing**（即将失败）或 **failed**（失败）。如果状况不为 **failing**（即将失败）或 **failed**（失败），挂接点会在因超出与硬件有关的可恢复错误阈值而执行的操作期间更改为 **failing**（即将失败）。如果状况不为 **failed**（失败），挂接点可在因不可恢复错误而导致的操作期间更改为 **failed**（失败）。

具有处于 **unconfigured**（未配置）状态的插卡的挂接点可处于已定义的任何一种状况。在与计算机有关的时间阈值过后，具有 **unconfigured**（未配置）插卡的挂接点的状况会从 **ok**（正常）衰退为 **unknown**（未知）。启动测试函数会将挂接点的状况更改为 **ok**（正常）、**failing**（即将失败）或 **failed**（失败），具体更改为哪种状况取决于测试结果。对于未提供测试函数的挂接点，可将其置于 **unknown**（未知）状况。如果测试中断，可将挂接点状况设置为上一个状况，即 **unknown**（未知）或 **failed**（失败）。可对处于 **unknown**（未知）、**ok**（正常）、**failing**（即将失败）或 **failed**（失败）状况的挂接点进行重新测试。

挂接点可能会因为多种原因而处于 **unusable**（不可用）状况，例如插口的电力不足或冷却不够、插卡无法识别、不受支持或配置错误等。处于 **unusable**（不可用）状况的挂接点绝对无法供系统使用。挂接点通常会始终处于这种状况，直到物理方面的问题被修复。

挂接点还会保持指示正在执行状态更改或重新评估状况的忙信息。

挂接点可使用特定于硬件的标识符 (*ap_id*) 来引用，这些标识符与挂接点在系统设备分层结构中的类型和位置相关。*ap_id* 不能有歧义，必须能够确定单个挂接点。支持两种类型的 *ap_id* 规范：物理和逻辑。物理 *ap_id* 包含完全指定的路径名，逻辑 *ap_id* 包含以更用户友好的方式来标识挂接点的缩略表示法。

例如，代表系统底板槽号 7 的挂接点包含的物理 *ap_id* 为

`/devices/central/fhc/sysctrl:slot7`，而逻辑 *ap_id* 为 `system:slot7`。再如，系统的第二个 PCI I/O 总线上的第三个插口可能包含的逻辑 *ap_id* 为 `pci2:plug3`。

挂接点也可以动态创建。动态挂接点是相对于系统中包含的基本挂接点而命名的。动态挂接点的 *ap_id* 由基本组件后接两个冒号 (:) 和动态组件组成。基本组件即基本挂接点 *ap_id*。动态组件是特定于硬件的，由相应的特定于硬件的库生成。

例如，考虑一个用 SCSI HBA 表示的基本挂接点，此挂接点包含物理 *ap_id*

`/devices/sbus@1f,0/SUNW,fas@e,88000000:scsi` 和逻辑 *ap_id* `c0`。连接至此 SCSI HBA 的磁盘可用包含逻辑 *ap_id* `c0::dsk/c0t0d0` 的动态挂接点表示，其中 `c0` 是基本组件，`dsk/c0t0d0` 是特定于硬件的动态组件。与之类似，此动态挂接点的物理 *ap_id* 应该是：`/devices/sbus@1f,0/SUNW,fas@e,88000000:scsi::dsk/c0t0d0`。

ap_type 是逻辑 *ap_id* 的部分形式，可能存在歧义，无法指定特定的挂接点。*ap_type* 最多只是逻辑 *ap_id* 部分的子字符串，但不包含冒号 (:) 分隔符。例如，*ap_type* 为 *pci* 应该会显示所有具有以 *pci* 开头的逻辑 *ap_id* 的挂接点。

建议不要使用 *ap_type*。-s 选项的新的 select 子选项为选择挂接点提供了更为通用和灵活的机制。请参见“选项”部分。

cfgadm 命令主要与特定于硬件的库中包含的与硬件有关的函数进行交互，因此其行为也与硬件有关。

每个配置管理操作可能都需要服务中断。完成请求的函数需要对交互用户执行明显的服务中断，在标准输入上确认的标准错误输出中将会出现提示，然后才会启动函数。始终使用 -y 或 -n 选项来分别回答是或否，即可覆盖确认。特定于硬件的选项，例如测试级别，将所作为使用 -o 选项的子选项提供。

更改系统配置状态的操作由系统日志守护进程 [syslogd\(1M\)](#) 进行审计。

此命令的参数符合 [getopt\(3C\)](#) 和 [getsubopt\(3C\)](#) 语法约定。

选项

所支持的选项如下：

- a

指定 -l 选项也必须列出动态挂接点。
- cfunction

在 *ap_id* 指定的挂接点上执行状态更改 *function*。

将 *function* 指定为
insert、*remove*、*disconnect*、*connect*、*configure* 或 *unconfigure*。这些函数可通过调用特定于硬件的库例程来在挂接点上执行状态转换，以下列表对其进行了定义。
- insert

执行相应的操作以允许用户手动插入插卡，或激活可执行物理插入的硬件提供机制。*insert* 可能产生特定于硬件的副作用，即会临时暂停部分系统活动。在这种情况下，特定于硬件的库会生成相应的警告消息，并告知用户专门针对相应硬件的特殊注意事项或操作过程。各种特定于硬件的错误都可导致此函数运行失败，并将插口状况设置为 *unusable*（不可用）。
- remove

执行相应的操作以允许用户手动移除插卡，或激活可执行物理移除的硬件提供机制。*remove* 可能产生特定于硬件的副作用，即会临时暂停部分系统活动。在这种情况下，特定于硬件的库会生成相应的警告消息，并告知用户专门针对相应硬件的特殊注意事项或操作过程。各种特定于硬

	件的错误都可导致此函数运行失败，并将插口状况设置为 unusable （不可用）。
disconnect	执行特定于硬件的操作，以将插口置于 disconnected （断开）状态，这样可以避免插卡通过此插口在正常情况下运行。
connect	执行特定于硬件的操作，以将插口置于 connected （已连接）状态，这样可以允许插卡通过此插口在正常情况下运行。
configure	执行特定于硬件的操作，以允许 Solaris 使用插卡的硬件资源。已配置的插卡是系统配置的组成部分，可通过 Solaris 设备处理维护命令（例如 psradm(1M) 、 mount(1M) 、 ifconfig(1M) ）来进行处理。
unconfigure	执行特定于硬件的操作，以从系统中逻辑移除插卡的硬件资源。插卡必须现在配置，且其硬件资源当前不能由 Solaris 使用。

状态转换函数可能因挂接点的状况或其他与硬件有关的注意事项而失败。当挂接点处于 **ok**（正常）或 **unknown**（未知）状况时，所有旨在添加资源的状态更改 *function*（**insert**、**connect** 和 **configure**）均会传送到特定于硬件的库。所有其他状况都需要使用强制选项，才能够将这些 *function* 传送到特定于硬件的库。挂接点状况不会阻止调用特定于硬件的库以从系统中执行与移除（**remove**、**disconnect** 和 **unconfigure**）硬件资源相关的操作。如果挂接点处于 **unknown**（未知）状况，特定于硬件的库可以拒绝状态更改 *function*。

挂接点的状况不一定是状态更改函数更改的，但如果状态更改操作期间出错，就会更改挂接点的状况。通过指定强制选项（**-f**），即可尝试覆盖状况以及强制执行可能会失败的状态更改操作。特定于硬件的安全性和完整性检查可以避免强制选项产生任何影响。

- f** 强制执行指定的操作。通常情况下，这会覆盖与硬件有关的安全功能。强制执行状态更改操作，即可根据任何与硬件有关的安全性检查，使用未处于 **ok**（正常）或 **unknown**（未知）状况的插卡的硬件资源。
- h [ap_id | ap_type ...]** 输出帮助消息文本。如果已指定 **ap_id** 或 **ap_type**，则会为用此参数表示的挂接点调用特定于硬件的库的帮助例程。

- `-l [ap_id | ap_type ...]` 列出指定挂接点的状态和状况。可使用 `-s` 选项和 `select` 子选项来过滤挂接点。调用不含任何一个操作选项的 `cfgadm` 等同于调用不含参数的 `-l`。列表的显示格式由 `-v` 和 `-s` 选项控制。如果指定了 `-a` 选项，挂接点便会动态扩展。
- `-n` 抑制任何交互式确认，并假定回答结果为 *no*（否）。如果未指定 `-n` 或 `-y`，则会通过标准错误输出和标准输入来获取交互式确认。如果上述任何标准通道均未响应终端（由 `isatty(3C)` 确定），则会假定使用 `-n` 选项。
- `-ohardware_options` 向主命令选项提供特定于硬件的选项。此硬件选项字符串的格式和内容完全是特定于硬件的。选项字符串 `hardware_options` 符合 `getsubopt(3C)` 语约定。
- `-slisting_options` 向列表 (`-l`) 命令提供列表选项。`listing_options` 符合 `getsubopt(3C)` 语约定。此子选项用于指定挂接点选择条件 (`select=select_string`)、所需的匹配类型 (`match=match_type`)、列表顺序 (`sort=field_spec`)、显示的数据 (`cols=field_spec` 和 `cols2=field_spec`)、列分界符 (`delim=string`) 以及是否抑制列标题 (`noheadings`)。

如果指定了 `select` 子选项，则只会列出符合指定条件的挂接点。`select` 子选项具有以下语法：

```
cfgadm -s select=attr1(value1):attr2(value2)...
```

其中 `attr` 是 `ap_id`、`class` 或 `type` 其中的一个。`ap_id` 是指逻辑 `ap_id` 字段，`class` 是指挂接点类，`type` 是指类型字段。`value1`、`value2` 等是要匹配的相应值。匹配类型可由 `match` 子选项按以下方式指定：

```
cfgadm -s match=match_type,select=attr1(value1)...
```

其中，`match_type` 可以是 `exact` 或 `partial`。缺省值为 `exact`。

`select` 子选项的参数可用引号引起来，以免被 `shell` 调用。

`field_spec` 是用冒号 (:) 串联起来的一个或多个 *data-field*，例如 `data-field:data-field:data-field`。`data-field` 是 `ap_id`、`physid`、`r_state`、`o_state`、`condition`、`type`、`busy`、`status` 和 `info` 中的其中一个。`ap_id` 字段的输出是挂接点的逻辑名称，`physid` 字段则包含物理名称。`r_state` 字段可以是 `empty`（空）、`disconnected`（断开）或 `connected`（已连接）。`o_state` 字段可以是 `configured`（已配置）或 `unconfigured`（未配置）。如果挂接点正忙，`busy`（忙）字段为 `y`；如果不忙，则为 `n`。`type`（类型）和 `info`（信息）字段均是特定于硬件的。`status_time` 字段会提供

`r_state`、`o_state` 或挂接点状况上次更改的时间。`status_time_p` 字段是 `status_time` 字段的可解析版本。如果挂接点具有关联的类，`class`（类）字段会列出类名。如果挂接点没有关联的类，`class`（类）字段会列出 `none`（无）。

`field_spec` 中的字段顺序非常重要：对于 `sort` 子选项，提供的第一个字段是主排序键。对于 `cols` 和 `cols2` 子选项，会按照请求的顺序输出这些字段。通过在 `sort` 子选项的 `field_spec` 内 `data-field` 名称之前添加减号 (-)，即可保留 `data-field` 上的排序顺序。`sort` 的缺省值为 `ap_id`。`cols` 和 `cols2` 的缺省值取决于是否提供 `-v` 选项，如果未提供，`cols` 的值为 `ap_id:r_state:o_state:condition`，`cols2` 未设置任何值。如果提供 `-v`，`cols` 的值为 `ap_id:r_state:o_state:condition:info`，`cols2` 的值为 `status_time:type:busy:physid`。`delim` 的缺省值是单个空格。`delim` 的值可以是任意长度的字符串。分界符不能包含逗号 (,) 字符，请参见 [getsubopt\(3C\)](#)。这些列表选项可用于创建可解析的输出。请参见“说明”部分。

-t

对一个或多个挂接点执行测试。测试函数用于重新评估挂接点的状况。如果 `hardware_options` 不含测试级别说明符，则会使用可确定硬故障的最快测试。

更全面的测试均是特定于硬件的，并且使用 `hardware_options` 进行选择。

测试结果用于更新指定插卡的状况，如果找不到故障，则更新为 `ok`（正常）；如果找到了可恢复的故障，则更新为 `failing`（即将失败）；如果找到了不可恢复的故障，则更新为 `failed`（失败）。

如果测试被中断，挂接点的状况可恢复为上一个值，如果未找到任何错误，可设置为 `unknown`（未知）；如果只找到可恢复的错误，可设置为 `failing`（即将失败）；如果找到了任何不可恢复的错误，则设置为 `failed`（失败）。如果测试正常完成且没有出现任何错误，应只将挂接点设置为 `ok`（正常）。

-v

在详细模式下执行。对于 `-c`、`-t` 和 `-x` 选项，将会输出提供每个尝试操作结果的消息。对于 `-h` 选项，将会输出详细的帮助信息。对于 `-l` 选项，将会输出每个挂接点的详细信息。

-xhardware_function

执行特定于硬件的函数。特定于硬件的专用函数可以更改插口或插卡的状态。挂接点状况会因为在特定于硬件的专

用函数运行期间出现错误而更改。*hardware_function* 字符串的格式和内容完全是特定于硬件的。选项字符串 *hardware_function* 符合 [getsubopt\(3C\)](#) 语法约定。

-y 抑制任何交互式确认，并假定回答结果为 yes（是）。

用法 使用此命令所需的特权取决于硬件。通常情况下，缺省系统配置会限制除超级用户的列出选项之外的所有选项。

示例 **示例1** 在设备树中列出挂接点
以下示例列出了除动态挂接点之外的所有挂接点。

example# cfgadm

Ap_Id	Type	Receptacle	Occupant	Cond
system:slot0	cpu/mem	connected	configured	ok
system:slot1	sbus-upa	connected	configured	ok
system:slot2	cpu/mem	connected	configured	ok
system:slot3	unknown	connected	unconfigured	unknown
system:slot4	dual-sbus	connected	configured	failing
system:slot5	cpu/mem	connected	configured	ok
system:slot6	unknown	disconnected	unconfigured	unusable
system:slot7	unknown	empty	unconfigured	ok
c0	scsi-bus	connected	configured	unknown
c1	scsi-bus	connected	configured	unknown

示例2 列出所有可配置的硬件信息
以下示例列出了当前所有可配置的硬件信息，包括由动态挂接点表示的硬件信息：

example# cfgadm -al

Ap_Id	Type	Receptacle	Occupant	Cond
system:slot0	cpu/mem	connected	configured	ok
system:slot1	sbus-upa	connected	configured	ok
system:slot2	cpu/mem	connected	configured	ok
system:slot3	unknown	connected	unconfigured	unknown
system:slot4	dual-sbus	connected	configured	failing
system:slot5	cpu/mem	connected	configured	ok
system:slot6	unknown	disconnected	unconfigured	unusable
system:slot7	unknown	empty	unconfigured	ok
c0	scsi-bus	connected	configured	unknown
c0::disk/c0t14d0	disk	connected	configured	unknown
c0::disk/c0t11d0	disk	connected	configured	unknown
c0::disk/c0t8d0	disk	connected	configured	unknown
c0::rmt/0	tape	connected	configured	unknown
c1	scsi-bus	connected	configured	unknown

示例3 根据挂接点属性有选择性地列出

以下示例列出了包含以 `scsi` 开头的类、以 `c` 开头的 `ap_id` 和以 `scsi` 开头的 `type`（类型）字段的所有挂接点。`-s` 选项的参数可用引号引起来，以免被 `shell` 调用。

```
example# cfgadm -s "match=partial,select=class(scsi):ap_id(c):type(scsi)"
```

Ap_Id	Type	Receptacle	Occupant	Cond
c0	scsi-bus	connected	configured	unknown
c1	scsi-bus	connected	configured	unknown

示例4 在详细模式下列出当前可配置的硬件信息

以下示例为详细模式下的 *ap-type system* 列出了当前可配置的硬件信息。

```
example# cfgadm -v -l system
Ap_Id          Receptacle Occupant  Condition Information
When          Type      Busy    Phys_Id
system:slot1   connected  configured ok
Apr  4 23:50 sbus-upa n      /devices/central/fhc/sysctrl:slot1
system:slot3   connected  configured ok      non-detachable
Apr 17 11:20 cpu/mem n      /devices/central/fhc/sysctrl:slot3
system:slot5   connected  configured ok
Apr  4 23:50 cpu/mem n      /devices/central/fhc/sysctrl:slot5
system:slot7   connected  configured ok
Apr  4 23:50 dual-sbus n      /devices/central/fhc/sysctrl:slot7
```

When 列表示 `status_time` 字段。

示例5 使用特定于硬件的 `Extended` 测试来测试两个插卡

以下示例使用特定于硬件的 `extended` 测试来测试两个插卡：

```
example# cfgadm -v -o extended -t system:slot3 system:slot5
Testing attachment point system:slot3 ... ok
Testing attachment point system:slot5 ... ok
```

示例6 使用强制选项配置插卡

以下示例使用强制选项将处于 `failing`（即将失败）状态的插卡配置到系统：

```
example# cfgadm -f -c configure system:slot3
```

示例7 从系统中取消配置插卡

以下示例从系统中取消配置了插卡：

```
example# cfgadm -c unconfigure system:slot4
```


示例8 在挂接点上配置插卡

以下示例配置了插卡：

```
example# cfgadm -c configure c0::dsk/c0t0d0
```

环境变量

请参见 [environ\(5\)](#)，了解影响 `cfgadm` 执行情况的以下环境变量的描述：`LC_TIME`、`LC_MESSAGES`、`NLS_PATH` 和 `TZ`。

- `LC_MESSAGES` 确定 `cfgadm` 如何显示列标题和错误消息。输出数据的列出方式不受此变量的设置影响。
- `LC_TIME` 确定 `cfgadm` 如何显示可读的状态更改时间 (`status_time`)。
- `TZ` 指定转换状态更改的时间时所使用的时区。这也适用于可读 (`status_time`) 和可解析 (`status_time_p`) 格式。

退出状态

将返回下列退出值：

- 0 成功完成。
- 1 出现错误。
- 2 指定目标上不支持配置管理。
- 3 用法错误。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见

[cfgadm_fp\(1M\)](#)、[cfgadm_ib\(1M\)](#)、[cfgadm_pci\(1M\)](#)、[cfgadm_sata\(1M\)](#)、[cfgadm_sbd\(1M\)](#)、[cfgadm_vhci\(1M\)](#)

诊断

诊断消息会出现在标准错误输出中。以下内容不是选项和用法错误，而是此实用程序生成的诊断消息：

```
cfgadm: Configuration administration not supported on ap_id
cfgadm: No library found for ap_id
cfgadm: ap_id is ambiguous
cfgadm: operation: Insufficient privileges
cfgadm: Attachment point is busy, try again
cfgadm: No attachment points with specified attributes found
cfgadm: System is busy, try again
cfgadm: operation: Operation requires a service interruption
cfgadm: operation: Data error: error_text
```

cfgadm: operation: Hardware specific failure: error_text

有关错误消息的其他详细信息，请参见 [config_admin\(3CFGADM\)](#)。

附注

硬件资源以特定于硬件的方式进入取消配置的池。这种情况会出现多次，例如系统初始化或执行取消配置操作之后。处于 `unconfigured`（未配置）状态的插卡不能供系统使用，除非进行特别干预。这种干预可表现为操作者发出命令或通过自动配置机制的形式干预。

`cfgadm` 命令的列表选项可用于为其他命令提供可解析输入，例如在 `shell` 脚本内。对于可解析输出，`-s` 选项必须用于选择所需的字段。`-s` 选项还可用于抑制列标题。下列字段会始终生成可解析的输出

出：`ap_id`、`physid`、`r_state`、`o_state`、`condition`、`busy`、`status_time_p`、`class` 和 `type`。可解析的输出决不能在字段值中插入空白字符。

以下 `shell` 脚本片段会查找第一个运行良好的 CPU 类型的 `unconfigured`（未配置）插卡。

```
found=
cfgadm -l -s "noheadings,cols=ap_id:r_state:condition:type" | \
while read ap_id r_state cond type
do
    if [ "$r_state" = unconfigured -a "$cond" = ok -a "$type" = CPU ]
    then
        if [ -z "$found" ]
        then
            found=$ap_id
        fi
    fi
done
if [ -n "$found" ]
then
    echo "Found CPU $found"
fi
```

可解析时间字段 (`status_time_p`) 的格式为 `YYYYMMDDhhmmss`，以适于字符串比较的格式提供年、月、日、时、分和秒。

若要了解系统配置管理支持的详细信息，应该参考特定于硬件的文档。

引用名	cfgadm_ac – EXX00 内存系统管理		
用法概要	<pre> /usr/sbin/cfgadm [-c configure] [-f] [-o disable-at-boot enable-at-boot] ac#:bank# ... /usr/sbin/cfgadm [-c unconfigure] [-o disable-at-bootp enable-at-boot] ac#:bank# ... /usr/sbin/cfgadm [-v] [-o quick normal extended, [max_errors=#]] -t ac#:bank#... /usr/sbin/cfgadm -x relocate-test ac#:bank# ... /usr/sbin/cfgadm [-l] -o disable-at-boot enable-at-boot ac#:bank# ... </pre>		
描述	特定于 ac 硬件的库 <code>/usr/platform/sun4u/lib/cfgadm/cfgadm_ac.so.1</code> 通过 cfgadm_sysctrl(1M) 提供在 E6X00、E5X00、E4X00 以及 E3X00 系统中配置和取消配置内存区的功能。该功能是 CPU/内存板动态重新配置的一部分。 内存区显示为设备树的挂接点。对于每个 CPU/内存板，发布两个挂接点（为板上的每个内存区提供一个）：<code>bank0</code> 和 <code>bank1</code>。如果内存区未存储任何内容，则插口状态为 <code>empty</code>（空）。如果内存区存储了内容，则插口状态为 <code>connected</code>（已连接）。内存区的插口状态永远不会为 <code>disconnected</code>（断开）。已连接的内存区的插卡状态可以为 <code>configured</code>（已配置）或 <code>unconfigured</code>（未配置）。如果插卡状态为 <code>configured</code>（已配置），表明 Solaris 正在使用内存；如果插卡状态为 <code>unconfigured</code>（未配置），表明 Solaris 没有使用内存。		
选项	有关命令选项的完整说明，请参见 cfgadm(1M)。 所支持的选项如下： <table> <tr> <td>-c configure unconfigure</td><td> 更改插卡状态。<code>configure</code> 参数确保内存已经初始化并将内存添加到 Solaris 内存池中。<code>unconfigure</code> 参数取消 Solaris 对内存的使用。如果 CPU/内存板从系统中移除，内存的两个区必须为未配置状态。 如果板上的内存标记为 <code>disabled-at-boot</code>（请参见 <code>info</code> 字段），<code>cfgadm</code> 将拒绝 <code>configure</code> 操作，除非指定 <code>-f</code>（强制）选项或“引导时启用”标志（<code>-o enable-at-boot</code>）。配置操作在很短的时间内完成并且执行时间与需初始化的内存大小成正比。 如果系统中没有足够的未确定内存（<code>VM viability</code>（VM 可行性）错误）或未配置的内存区中仍存在不能删除的内存（<code>non-relocatable pages</code> 错误），<code>cfgadm</code> 将拒绝 <code>unconfigure</code> 操作。不可重定位的页面存在与否由 <code>info</code> 列表字段中的词 <code>permanent</code> 指示。由于各种因素（如系统 </td></tr> </table>	-c configure unconfigure	更改插卡状态。<code>configure</code> 参数确保内存已经初始化并将内存添加到 Solaris 内存池中。<code>unconfigure</code> 参数取消 Solaris 对内存的使用。如果 CPU/内存板从系统中移除，内存的两个区必须为未配置状态。 如果板上的内存标记为 <code>disabled-at-boot</code>（请参见 <code>info</code> 字段），<code>cfgadm</code> 将拒绝 <code>configure</code> 操作，除非指定 <code>-f</code>（强制）选项或“引导时启用”标志（<code>-o enable-at-boot</code>）。配置操作在很短的时间内完成并且执行时间与需初始化的内存大小成正比。 如果系统中没有足够的未确定内存（<code>VM viability</code>（VM 可行性）错误）或未配置的内存区中仍存在不能删除的内存（<code>non-relocatable pages</code> 错误），<code>cfgadm</code> 将拒绝 <code>unconfigure</code> 操作。不可重定位的页面存在与否由 <code>info</code> 列表字段中的词 <code>permanent</code> 指示。由于各种因素（如系统
-c configure unconfigure	更改插卡状态。<code>configure</code> 参数确保内存已经初始化并将内存添加到 Solaris 内存池中。<code>unconfigure</code> 参数取消 Solaris 对内存的使用。如果 CPU/内存板从系统中移除，内存的两个区必须为未配置状态。 如果板上的内存标记为 <code>disabled-at-boot</code>（请参见 <code>info</code> 字段），<code>cfgadm</code> 将拒绝 <code>configure</code> 操作，除非指定 <code>-f</code>（强制）选项或“引导时启用”标志（<code>-o enable-at-boot</code>）。配置操作在很短的时间内完成并且执行时间与需初始化的内存大小成正比。 如果系统中没有足够的未确定内存（<code>VM viability</code>（VM 可行性）错误）或未配置的内存区中仍存在不能删除的内存（<code>non-relocatable pages</code> 错误），<code>cfgadm</code> 将拒绝 <code>unconfigure</code> 操作。不可重定位的页面存在与否由 <code>info</code> 列表字段中的词 <code>permanent</code> 指示。由于各种因素（如系统		

	负荷以及辅助存储器所需分页的数量），取消 Solaris 对内存的使用可能需要大量时间。<code>unconfigure</code> 操作可随时取消并且可通过发出信号中断该命令的调用使内存恢复满配置状态。如果在超时期限内未能删除内存，<code>unconfigure</code> 操作会自行取消。缺省的超时期限为 60 秒，可通过 <code>-o timeout=#</code> 选项更改。如果值设置为 0，将禁止超时。
-f	强制选项。在非易失的 <code>disabled-memory-list</code> 变量中配置标记为“引导时禁用”的内存区时，使用该选项覆盖数据块。请参见《<i>Platform Notes: Sun Enterprise 6x00/5x00/4x00/3x00 Systems</i>》。
-l	列表选项。支持此选项，如 <code>cfgadm(1M)</code> 中所述。 <code>type</code>（类型）字段始终为 <code>memory</code>（内存）。 对于空的内存区，<code>info</code>（信息）字段包含以下信息： slot# empty slot# 表示 CPU/内存板所插入的系统插槽。例如，如果为 slot11，则与 <code>cfgadm</code> 一起用于处理相关板的挂接点为 <code>sysctrl0:slot11</code>。对于连接的内存区，<code>info</code>（信息）字段包含以下信息： <pre>slot# sizeMb sizeGb [(sizeMb sizeGb used)] base 0x### [interleaved #-way] [disabled at boot] [permanent]</pre> 根据实际需要，内存区大小可以 Mb 或 Gb 为单位表示。如果内存没有完全使用，将报告所使用的内存大小。物理基址可用十六进制表示。如果该内存区与其他某个内存区相互交错，将报告交错因子。如果板上的内存通过非易失 <code>disabled-memory-list</code> 变量在引导时禁用，将报告这种情况。如果内存区中存在无法删除的内存，将报告该内存为永驻内存。
-o disable-at-boot enable-at-boot	这些选项用于修改非易失的 <code>disabled-memory-list</code> 变量的状态。如果不需要执行命令，这些选项可以与 <code>-c</code> 选项或显式或隐式列表命令 <code>-l</code> 一起使用。使用 <code>-o enable-at-boot</code> 与 <code>configure</code> 命令通过在禁用内存列表中配置板上的内存来覆盖数据块。
-o extended normal quick	该选项与 <code>-t</code> 选项一起指定测试级别。

	normal 测试级别确保每个内存单元存储 0 和 1 并且检查所有单元是否可单独寻址。quick 测试级别只进行 0 与 1 的测试，通常会跳过地址线问题。extended 测试使用模式来检测相邻单元干扰问题。缺省的测试级别为 normal。请参见 -t 选项。
-o max_errors=#	该选项与 -t 选项一起使用指定所允许的错误的最大数量。如果不指定该选项，则采用缺省值 32。
-o timeout=#	该选项与 unconfigure 命令一起使用来设置自取消超时。缺省值为 60，并且单位为秒。如果该选项的值为 0，表明无超时。
-t	测试未配置的内存区。使用 -o quick normal extended 选项指定测试级别。 如果该测试能够在内存区上运行，cfgadm 命令返回 0（表明运行成功）并退出。如果存在挂点，可获取测试结果。
-v	详细选项。将该选项与 -t 选项一起使用来显示详细的测试进度和结果。
-x relocate-test	对于指定的内存区中使用的所有内存页，尝试执行 unconfigure 命令中所使用的重定位操作。该操作的成功并不能保证内存区为未配置的。如果该操作失败，表明内存区可能不是未配置的。该选项只用于测试目的。

操作数

支持以下操作数：	
ac#:bank#	内存区的挂接点由地址控制器 (ac) 驱动程序 (ac#) 的实例发布。将会为每个系统板创建一个 ac 驱动程序实例，但只有与 CPU/内存板关联的实例会发布两个内存区挂接点：bank0 和 bank1。 这种形式符合 cfgadm(1M) 中给出的逻辑 ap_id 规范。对应的物理 ap_id 已在“文件”部分中列出。 ac 驱动程序实例编号与对应的板的槽号无关。完整的物理挂接点标识符中所包含的槽号是紧跟在 fhc@ 部分后面的十六进制槽号的两倍。

文件

/devices/fhc@*,f8800000/ac@0,1000000:bank?	挂接点
/usr/platform/sun4u/lib/cfgadm/cfgadm_ac.so.1	特定于硬件的库文件

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/library/platform

另请参见 [cfgadm\(1M\)](#)、[cfgadm_sysctrl\(1M\)](#)、[config_admin\(3CFGADM\)](#)、[attributes\(5\)](#)

《*Sun Enterprise 6x00, 5x00, 4x00 and 3x00 Systems Dynamic Reconfiguration User's Guide*》

《*Platform Notes: Sun Enterprise 6x00/5x00/4x00/3x00 Systems*》

附注 有关 EXX00 系统 CPU/内存板动态重新配置的更多详细信息，请参见《*Sun Enterprise 6x00, 5x00, 4x00 and 3x00 Systems Dynamic Reconfiguration User's Guide*》。

引用名 `cfgadm_cardbus` – 特定于 cardbus 硬件的 `cfgadm` 命令

用法概要

```
/usr/sbin/cfgadm [-f ] [-y | -n ] [-v]
                [-o hardware_options] -c function ap_id [ap_id]

/usr/sbin/cfgadm [-f ] [-y | -n ] [-v]
                [-o hardware_options] -x hardware_function ap_id
                [ap_id]

/usr/sbin/cfgadm [-v] [-s listing_options]
                [-o hardware_options] [-l [ap_id | ap_type]]

/usr/sbin/cfgadm [-v] [-o hardware_options] -t ap_id [ap_id]

/usr/sbin/cfgadm [-v] [-o hardware_function] -h
                [ap_id| ap_type]
```

描述

Solaris 中的 CardBus 插槽为热插拔插槽。热插拔能力由特定于 PCI 硬件的库 `/usr/lib/cfgadm/pci.so.1` 通过 `cfgadm` 命令支持（请参见 [cfgadm\(1M\)](#)）。

CardBus、PCI、CompactPCI 以及 PCI Express 之间的热插拔管理模型以同样的方式运行。有关使用信息，请参见 [cfgadm_pci\(1M\)](#)。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/library

另请参见

[cfgadm\(1M\)](#)、[config_admin\(3CFGADM\)](#)、[libcfgadm\(3LIB\)](#)、[attributes\(5\)](#)

《Oracle Solaris 管理：常见任务》

引用名 cfgadm_fp – 驱动程序专用 cfgadm 命令

用法概要

```
/usr/sbin/cfgadm [-f] [-n | -y] [-v] [-o hardware_options]
                 -c function ap_id [ap_id]

/usr/sbin/cfgadm [-v] [-a] [-s listing_options]
                 [-o hardware_options] [-l [ap_id]]

/usr/sbin/cfgadm [-v] [-o hardware_options] -h [ap_id]
```

描述

fp 端口驱动程序插件 /usr/lib/cfgadm/fp.so.1 提供通过 [cfgadm\(1M\)](#) 管理光纤通道光纤网络设备节点的功能。cfgadm 在挂接点上运行。挂接点是可动态重新配置硬件资源的系统位置。有关挂接点的其他详细信息，请参见 [cfgadm\(1M\)](#)。

对于光纤通道光纤网络设备节点管理，每个 fp 端口节点采用设备树中的一个挂接点表示。此外，每个光纤通道设备还采用一个动态挂接点表示。挂接点通过 ap_id 命名。定义下列两种类型的 ap_id：逻辑和物理。物理 ap_id 基于物理路径名。逻辑 ap_id 是更短、更加用户友好的名称。对于 fp 端口节点，逻辑 ap_id 为对应的磁盘控制器号。例如，c0 是一个典型的逻辑 ap_id。

光纤通道设备使用端口全球名称 (World Wide Name, WWN) 命名。如果一个磁盘设备连接到控制器 c0，其 ap_id 可以是：

c0::50020f2300006077

其中 50020f2300006077 用于标识特定光纤通道设备的端口 WWN。

缺省情况下，对光纤通道专用环路端口、光纤网络端口或公共环路端口上的每个设备进行探测，并使其可供 Solaris 使用。通过启动应用程序或最终用户操作，可使连接到光纤通道光纤网络端口或公共环路端口的设备不可供 Solaris 使用。此操作类似于通过管理用户接口的方法热拔除设备。应用程序或用户可以使用 /usr/lib/cfgadm/fp.so.1 库来启用 libcfgadm 提供接口以完成此任务。

当前连接的光纤网络设备的列表以挂接点的形式生成。

此系统中挂接点的简单列表包括 fp 端口节点（非光纤通道设备）中的挂接点。以下示例使用列表选项 (-l) 的 -a 标志来列出光纤通道设备：

```
# cfgadm -l
Ap_Id      Type      Receptacle  Occupant  Condition
c0          fc-fabric connected   configured unknown
c1          fc-private connected   configured unknown
c2          fc-pt_to_pt connected   configured unknown
c3          fc       connected   unconfigured unknown
sysctrl0:slot0  cpu/mem  connected   configured ok
sysctrl0:slot1  sbus-upa connected   configured ok
```

以下示例列出了连接到 fp 端口的光纤通道设备。

```
# cfgadm -al
Ap_Id      Type      Receptacle  Occupant  Condition
c0          fc-fabric  connected   configured unknown
```



```

c0::50020f2300006077 disk          connected  configured  unknown
c0::50020f23000063a9 disk          connected  configured  unknown
c0::50020f2300005f24 disk          connected  configured  unknown
c0::50020f2300006107 disk          connected  configured  unknown
c1          fc-private             connected  configured  unknown
c1::220000203708b69c disk          connected  configured  unknown
c1::220000203708ba7d disk          connected  configured  unknown
c1::220000203708b8d4 disk          connected  configured  unknown
c1::220000203708b9b2 disk          connected  configured  unknown
c2          fc-pt_to_pt            connected  configured  unknown
c2::500104f000937528 tape          connected  configured  unknown
c3          fc                     connected  unconfigured unknown
sysctrl0:slot0      cpu/mem        connected  configured  ok
sysctrl0:slot1      sbus-upa        connected  configured  ok

```

在此示例中，fc-fabric 类型的 `ap_id c0` 指示 fp 端口与光纤网络相连。对于具有光纤网络相关类型（例如 fc-fabric 和 fc-public）的 fp 端口，缺省情况下在引导时创建设备节点并可通过 `cfgadm` 配置与取消配置操作进行管理。fc-private 类型的 `ap_id c1` 指示 fp 端口连接到专用环路且缺省情况下也创建设备节点。fc-pt_to_pt 类型的 `ap_id c2` 指示 fp 端口直接连接到另一个 N_port 且缺省情况下也创建设备节点。fc 类型的 `ap_id c3` 指示 fp 端口 c2 未与任何对象相连。光纤通道设备 `ap_id` 的 "Type"（类型）字段显示了设备中 LUN 0 的 SCSI 设备类型。

在 Solaris 中配置了具有多个 FCP SCSI LUN 的光纤通道设备，其中每一个 FCP SCSI LUN 均可作为一个 Solaris 设备使用。假定 `ap_id c0::50020f2300006077` 和 `c0::50020f23000063a9` 表示具有多个 FCP SCSI LUN 的光纤通道设备。

以下示例显示了如何列出具有 FCP SCSI LUN 的 `ap_id` 的信息：

```

# cfgadm -al -o show SCSI_LUN

```

Ap_Id	Type	Receptacle	Occupant	Condition
c0	fc-fabric	connected	configured	unknown
c0::50020f2300006077,0	disk	connected	configured	unknown
c0::50020f2300006077,1	disk	connected	configured	unknown
c0::50020f2300006077,2	disk	connected	configured	unknown
c0::50020f2300006077,3	disk	connected	configured	unknown
c0::50020f23000063a9,0	disk	connected	configured	unknown
c0::50020f23000063a9,1	disk	connected	configured	unknown
c0::50020f23000063a9,2	disk	connected	configured	unknown
c0::50020f23000063a9,3	disk	connected	configured	unknown
c0::50020f2300005f24,0	disk	connected	unconfigured	unknown
c0::50020f2300005f24,1	disk	connected	unconfigured	unknown
c0::50020f2300006107,0	disk	connected	unconfigured	unknown
c0::50020f2300006107,1	disk	connected	unconfigured	unknown
c1	fc-private	connected	configured	unknown
c1::220000203708b69c,0	disk	connected	configured	unknown
c1::220000203708ba7d,0	disk	connected	configured	unknown
c1::220000203708b8d4,0	disk	connected	configured	unknown

c1::220000203708b9b2,0	disk	connected	configured	unknown
c2	fc-pt_to_pt	connected	configured	unknown
c2::500104f000937528,0	tape	connected	configured	unknown
c3	fc	connected	unconfigured	unknown

在此示例中，ap_id c0::50020f2300006077,0 标识光纤通道设备的 FCP SCSI LUN 0，采用端口 WWN 50020f2300006077 表示。根据报告，光纤通道设备具有 4 个 FCP SCSI LUN，并且这些 FCP SCSI LUN 均已配置。列出了与端口 WWN 50020f2300006077 相关联的 4 个 FCP SCSI LUN 级别 ap_id。此列表还显示了未配置的光纤通道设备的 FCP SCSI LUN。根据报告，由 c0::50020f2300005f24 表示的光纤通道设备具有两个 FCP SCSI LUN。c0::50020f2300005f24 上的配置操作会创建两个 Solaris 设备。FCP SCSI LUN 级别 ap_id 的 Type（类型）字段显示每个 LUN 的 SCSI 设备类型。如果光纤通道设备具有不同的设备类型 LUN，Type（类型）字段会反映出来。

fp 端口上挂接点的插口和插卡状态具有下列含义：

- configured（已配置）
在 fp 端口上配置了一个或多个设备
- connected（已连接）
fp 端口处于活动状态
- 断开
fp 端口停止（IO 活动暂停）
- empty（空）
不适用
- unconfigured（未配置）
在 fp 端口上未配置任何设备

fp 端口上各光纤通道设备的状态：

- 配置网络
设备已在 Solaris 中配置并可供使用
- connected（已连接）
设备连接的 fp 端口处于活动状态
- 断开
设备连接的 fp 端口处于停止状态
- unconfigured（未配置）
可配置设备

fp 端口上挂接点的 condition（状态）字段具有下列含义：

- failed（失败）
一个错误状态已阻止 fp 端口检测光纤通道连接的存在情况或类型的功能。

fp 端口上各光纤通道设备的 condition（状态）字段具有下列含义：

failed (失败)

在光纤网络上探测设备时发生错误。

failing (即将失败)

在主机上配置了一个设备且其状态在 Solaris 上显示为正常 (即, 联机), 但是该设备目前不存在或在光纤网络上不显示, 或者因为用于配置该设备的本地端口上存在错误状态, 该设备的存在情况无法验证。

unusable (不可用)

已在主机上配置设备, 但其目前处于脱机状态或存在故障。

condition (状态) 为 **unknown** 指示在光纤网络上探测设备已完成且没有错误, 并且如果之前配置了该设备, Solaris 主机中的设备状态为正常。无法保证设备的内部状态。

选项

除列出 (-l) 外, **cfgadm** 还定义了一些类型的操作。这些操作包括调用配置状态更改以及获取配置管理帮助消息 (-h)。

所支持的选项如下:

-c function

为 **fp** 传输特定库定义了下列通用命令:

对于 **fc-fabric** 类型 **fp** 端口挂接点上的光纤通道设备挂接点, 支持下列配置状态更改操作:

configure

配置连接到主机的光纤通道光纤网络设备。如果在列表操作的输出中将光纤通道设备列为 **unknown** (未知) 类型, 可能无法配置该设备。未尝试配置类型为 **unknown** (未知) 的设备。可使用强制选项 (-f) 来强制 **fp** 端口驱动程序插件尝试配置任意设备。报告进程中的任何错误。缺省情况下, 在光纤通道光纤网络设备上发现的每个 FCP SCSI LUN 都进行了配置。但是, 在 **fp.conf** 文件的 "pwwn-lun-blacklist" 属性中指定的 FCP SCSI LUN 仍将处于 **unconfigured** (未配置) 状态。FCP SCSI LUN 级别列表反映了这些 FCP SCSI LUN 的状态。重新引导或在连接的控制器上 Solaris 动态重新配置之后, 它们仍会处于 **unconfigured** (未配置) 状态。有关 "pwwn-lun-blacklist" 属性的更多详细信息, 请参见 [fp\(7d\)](#)。

unconfigure

在主机上取消配置光纤通道光纤网络设备。在下次重新引导或在连接该设备的控制器上 Solaris 动态重新配置之前, 该设备保持处于 **unconfigured** (未配置) 状态, 而届时所有光纤网络设备都会自动枚举。通过使用 **fp.conf** 文件中的 "manual_configuration_only" 属性可以更改缺省的行为。如果设置了该属性, 设备会在重新引导后仍处于 **unconfigured** (未配置) 状态。有关 "manual_configuration_only" 属性的更多详细信息, 请参见 [fp\(7d\)](#)。

对于光纤通道专用环路设备和 **N_Port** 点到点设备, **configure** 命令返回成功, 但不会执行任何操作。**unconfigure** 命令在专用环路设备和 **N_Port** 点到点设备上不受支持。缺省情况下, 专用环路设备和 **N_Port** 点到点设备由 Solaris 光纤通道驱动程序

配置，且不通过最终用户或应用程序启动的操作进行管理。`fp.conf` 文件中的 `pwwn-lun-blacklist` 属性以与应用于光纤网络设备相同的方式应用于专用环路设备和 `N_Port` 点到点设备。

-f

强制实施 `configure` 更改状态操作，无论 `condition`（状态）或 `type`（类型）为何。请参见上述有关 `configure` 更改状态操作的说明。

-h *ap_id*

获取 `fp` 传输特定帮助。指定任意 `fp` 挂接点。

-o *hardware_options*

支持下列硬件选项。

show_SCSI_LUN

如果使用 `list` 选项 `-al` 指定，列出与发现的光纤通道设备的每个 FCP SCSI LUN 相关联的 `ap_id`。请参见上述 FCP SCSI LUN 级别列表的说明和示例。在 FCP SCSI LUN 级别上不支持设备节点创建。请参见“说明”部分。

缺省情况下，所有光纤通道设备对 Solaris 都可用。通过设置 `/kernel/drv/fp.conf` 文件中的属性 `"manual_configuration_only"`，可完成在缺省情况下仅启用对 Solaris 可用的光纤网络设备的子集。设置了 `fp.conf` 中的 `"manual_configuration_only"` 时，除非应用程序或最终用户之前已请求在 Solaris 中配置该设备，否则所有光纤网络设备对 Solaris 均不可用。`configure` 状态更改命令使该设备对 Solaris 可用。在光纤网络设备上成功执行 `configure` 操作之后，将相关链接添加到 `/dev` 名称空间。`unconfigure` 状态更改命令使设备对 Solaris 不可用。

使用 `-c configure` 操作在主机上成功配置了光纤通道光纤网络设备时，其物理 `ap_id` 存储在系统信息库中。使用 `-c unconfigure` 操作取消光纤通道光纤网络设备配置时，从该系统信息库中删除物理 `ap_id`。缺省情况下，自动枚举所有光纤网络设备，且仅当设置了 `fp.conf` 中的 `"manual_configuration_only"` 属性时才使用该系统信息库。有关 `"manual_configuration_only"` 属性的更多详细信息，请参见 [fp\(7d\)](#)。

可以使用 `-c` 选项指定下列命令以控制系统信息库的更新行为：

force_update

对于 `configure`，无条件地将挂接点添加到系统信息库，对于 `unconfigure`，无条件地删除挂接点。

no_update

无论实施任何操作，均不更新系统信息库。

对于正常的 `configure` 和 `unconfigure` 操作，不应使用这些选项。请参见“警告”部分。

如果为一个光纤通道设备配置了多个 FCP SCSI LUN 且与其 FCP SCSI LUN 相关联的任意 Solaris 处于不可用状态，将整个光纤通道设备报告为不可用。在使用 `-c unconfigure` 命令的情况下的以下选项仅移除光纤通道设备的处于不可用状态的 Solaris 设备。

`unusable_SCSI_LUN`

对于 `unconfigure` 操作，移除目标设备的任何脱机的设备节点。

`-s listing_options`

有关用法信息，请参见 `cfgadm(1M)`。

`-t ap_id`

目前没有可用的测试命令。

`-x hardware_function`

目前没有可用的硬件专用函数。

所有其他选项的含义与 `cfgadm(1M)` 手册页中定义的含义相同。

示例

示例1 取消配置磁盘

以下命令可用于取消配置磁盘：

```
# cfgadm -c unconfigure c0::210000203708b606
```

示例2 取消配置单个挂接点下的所有配置的磁盘

以下命令取消配置挂接点 `c0` 下的所有配置的磁盘。

```
# cfgadm -c unconfigure c0
```

示例3 配置磁盘

以下命令配置磁盘：

```
# cfgadm -c configure c0::210000203708b606
```

示例4 配置单个挂接点下的所有未配置的磁盘

以下命令配置挂接点 `c0` 下的所有未配置的磁盘。

```
# cfgadm -c configure c0
```

示例5 从系统信息库中移除光纤通道光纤网络设备挂接点

以下命令从光纤网络设备系统信息库中无条件地移除光纤通道光纤网络设备挂接点。

```
# cfgadm -c unconfigure -o force_update c0::210000203708b606
```

示例6 从目标设备中移除脱机的 Solaris 设备节点

以下命令从目标设备中移除脱机的 Solaris 设备节点：

```
# cfgadm -c unconfigure -o unusable_SCSI_LUN c0::210000203708b606
```

文件

`/usr/lib/cfgadm/fp.so.1`
光纤通道光纤网络设备节点管理的硬件专用库。

`/etc/cfg/fp/fabric_WWN_map`
当前配置的光纤网络设备的物理 `ap_id` 的系统信息库。仅用于在引导时重新配置这些光纤网络设备。仅当设置了 `/kernel/drv/fp.conf` 文件中的 `"manual_configuration_only"` 时才使用该系统信息库。

`/etc/rcS.d/fdevattach`
重新配置引导时物理 `ap_id` 在 `/etc/cfg/fp/fabric_WWN_map` 中列出的光纤网络设备。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/storage/fibre-channel/fc-fabric、service/storage/fibre-channel

另请参见 [svcs\(1\)](#)、[cfgadm\(1M\)](#)、[svcadm\(1M\)](#)、[config_admin\(3CFGADM\)](#)、[libcfcgadm\(3LIB\)](#)、[attributes\(5\)](#)、

警告 对于正常 `configure/unconfigure` 操作下的系统信息库更新，请勿使用硬件专用选项。如果当引导时光纤网络设备的节点创建失败且错误状态被认为是永久性的，希望使用硬件专用选项。带 `force_update` 硬件专用选项的 `unconfigure` 命令从系统信息库中无条件地移除存在故障的光纤网络设备的挂接点。

附注 对于具有未知或无 SCSI 设备类型的设备（例如，光纤通道主机总线适配器），`configure` 操作可能不适用。

`configure` 和 `unconfigure` 命令在由端口 WWN `ap_id` 表示的光纤通道设备级别上运行。如果为一个光纤通道设备配置了多个 FCP SCSI LUN，关联端口 WWN `ap_id` 上的 `configure` 命令用于为每个 FCP SCSI LUN 创建一个 Solaris 设备，除非使用 `fp.conf` 文件中的 `"pwwn-lun-blacklist"` 属性进行了指定。`unconfigure` 命令移除所有与端口 WWN `ap_id` 相关联的 Solaris 设备。FCP SCSI LUN 级别 `ap_id` 对于 `configure` 和 `unconfigure` 命令无效。

过时的 `show_FCP_dev` 选项已被新的 `show_SCSI_LUN` 选项代替，且过时的 `unusable_FCP_dev` 选项已被新的 `unusable_SCSI_LUN` 选项代替。

`cfgadm_fp` 服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

```
svc:/system/device/fc-fabric:default
```

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

引导时启动此服务后，对于光纤网络设备配置，无需实施此服务的管理操作。

引用名	cfgadm_ib – InfiniBand 硬件专用 cfgadm 命令
用法概要	<pre> /usr/sbin/cfgadm -f [-y -n] [-v] -c function ap_id... /usr/sbin/cfgadm [-f] [-y -n] [-v] [-o hardware_options] -x hardware_function ap_id... /usr/sbin/cfgadm -v [-a] [-s listing_option] [-] [ap_id ap_type...] /usr/sbin/cfgadm -v -h [ap_id]... </pre>
描述	InfiniBand 硬件专用库 <code>/usr/lib/cfgadm/ib.so.1</code> 提供通过 cfgadm(1M) 实用程序管理光纤网络的功能。cfgadm 在挂接点上运行。请参见 cfgadm(1M)。 基于 IB 设备管理器 (IB Device Manager, IBDM) 提供的服务，InfiniBand (IB) 设备由 IB nexus 驱动程序 ib(7D) 进行枚举。 IB nexus 驱动程序创建并初始化五种类型的子设备节点。 ■ IB 端口设备 ■ IB HCA 服务 (HCA_SVC) 设备 ■ IB 虚拟物理挂接点 (Virtual Physical Point of Attachment, VPPA) 设备 ■ I/O 控制器 (I/O Controller, IOC) ■ IB 伪设备 有关 IB 端口、IB VPPA 以及 IB HCA_SVC 设备的枚举的详细信息，请参见 ib(7D)。有关 IBDM 的其他信息，请参见 ibdm(7D)。有关 IB 伪设备的详细信息，请参见 ib(4)。 对于 IB 管理，为光纤网络管理创建两种类型的静态挂接点，如在给定主机中所显示的。一种静态挂接点是 <code>ib</code>，且该光纤网络中的所有 IB 设备（IOC、端口、VPPA、HCA_SVC 或伪设备）表示为基于该静态挂接点的动态挂接点。另一种静态挂接点是基于其节点的全局唯一标识符 (Globally Unique Identifier, GUID) 值为主机中的每个主机通道适配器 (Host Channel Adapter, HCA) 创建的。 挂接点通过 <code>ap_id</code> 命名。<code>ap_id</code> 包括逻辑 <code>ap_id</code> 和物理 <code>ap_id</code>。物理 <code>ap_id</code> 基于物理路径名。对于 IB 光纤网络，物理 <code>ap_id</code> 为 <code>/devices/ib:fabric</code>。逻辑 <code>ap_id</code> 为更短、更加用户友好的名称。 IB 光纤网络的静态 <code>ap_id</code> 为 <code>ib</code>。IB 设备为动态挂接点，没有物理 <code>ap_id</code>。IOC 的逻辑 <code>ap_id</code> 包含其 GUID，<code>ib::IOC-GUID</code>。<code>ib::80020123456789a</code> 就是 IOC <code>ap_id</code> 的一个示例。伪设备的逻辑 <code>ap_id</code> 的格式为 <code>ib::driver_name,unit-address</code>。有关详细情况，请参见 ib(4)。<code>ib::sdp,0</code> 就是伪 <code>ap_id</code> 的一个示例，其中 "sdp" 是驱动程序的名称，"0" 是相应的 <code>unit-address</code> 属性。端口、VPPA 以及 HCA_SVC 设备的逻辑 <code>ap_id</code> 包含其分区键 (<code>P_Key</code>)、端口 GUID/节点 GUID 以及通信服务名称。<code>ap_id</code> 的格式如下： 端口设备 <pre>ib::PORT_GUID,0,service-name</pre> VPPA 设备 <pre>ib::PORT_GUID,P_Key,service-name</pre>

HCA_SVC 设备
ib::HCA_GUID,0,servicename

端口和 HCA_SVC 设备的分区键 (P_Key) 为 0。P_Key 有助于确定 VPPA 设备节点的端口所属的分区。一个端口可能有多个 P_Key。ib::80245678,ffff,ipib 就是 VPPA 设备的逻辑 ap_id 点的一个示例。在该示例中，port-GUID 为 80245678，P_Key 为 0xffff，服务名称为 ipib。服务名称信息可从 /kernel/drv/ib.conf 文件中获取，该文件包含服务名称字符串。HCA 的逻辑 ap_id 包含其节点的 GUID 值：hca:HCA-GUID。例如，hca:21346543210a987。

IB 挂接点列表包含有关该光纤网络中的所有 IB 设备（包含在 IBDM 以及 IB nexus 驱动程序上显示的 IOC 设备、VPPA 设备、HCA_SVC 设备、伪设备以及端口设备）的信息，即使这些 IB 设备在主机上未显示且未配置用于使用，也如此。

以下显示包含五个 IB 设备（两个 IOC 设备、一个 VPPA 设备、一个端口设备以及一个 HCA_SVC 设备）的列表和一个 HCA。

```
example# cfgadm -al
Ap_Id                                Type      Receptacle  Occupant    Condition
hca:21346543210a987                 IB-HCA    connected   configured  ok
ib                                    IB-FABRIC connected   configured  ok
ib::80020123456789a                 IB-IOC    connected   configured  ok
ib::802abc9876543                   IB-IOC    connected   unconfigured unknown
ib::80245678,ffff,ipib               IB-VPPA   connected   configured  ok
ib::12245678,0,nfs                   IB-PORT   connected   configured  ok
ib::21346543,0,hnfs                  IB-HCA_SVC connected   configured  ok
ib::sdp,0                             IB-PSEUDO connected   configured  ok
```

ap_id ib::802abc9876543 表示没有被主机配置为用于使用的 IOC 设备或者之前通过显式

```
cfgadm -c unconfigure
```

操作脱机的 IOC 设备。Condition（状态）列中显示的信息可以区分这两种情况。如果 IB 设备的 P_Key 为 0 且具备 HCA GUID，该 IB 设备为 HCA_SVC 设备。有关挂接点列表的更多信息，请参见 [cfgadm\(1M\)](#)。

挂接点的插口状态的含义如下：

connected（已连接）
对于 IOC 设备、VPPA 设备、端口设备、伪设备以及 HCA_SVC 设备，connected（已连接）表示设备在主机上显示。相应设备可能还未配置为由 Solaris 使用。

对于 HCA 挂接点，connected（已连接）表示该挂接点已配置并处于使用状态。

所有的 IB ap_id 总是显示为 connected（已连接）状态。

挂接点的插卡状态的含义如下：

configured (已配置)

IB 设备以及 HCA ap_id 已配置并且可供 Solaris 使用。

unconfigured

如果 IB 设备的 ap_id 通过 `cfgadm -c unconfigure` 显式脱机，该设备将不能成功配置。这种情况是可能的，因为该 IB 设备并没有成功配置为由 Solaris（没有驱动器或设备有问题）使用或者该设备没有配置为由 IB nexus 驱动程序使用。

unconfigured (未配置) 操作不支持 HCA 挂接点。如果系统中不存在 IB 硬件，则 IB 静态 ap_id ib 显示为 unconfigured (未配置)。

挂接点的状态含义如下：

failed (失败)

未使用。

FAILED

未使用。

ok

正常状态。可供使用。

unknown (未知)

该状态只对那些已经由 IBDM 进行探测但尚未被 Solaris 配置为使用的 IB 设备有效。通过 `cfgadm -c unconfigure` 操作显式脱机的设备也显示为 unknown (未知) 状态。该状态不适用于 HCA 挂接点。

unusable (不可用)

未使用。

选项

所支持的选项如下：

`-c function`

IB 硬件专用库支持两个通用命令 (*function*)。静态挂接点 (即 HCA ap_id 和 IB 静态 ib ap_id) 不支持这两个命令。

支持以下两个通用命令：

configure (配置)

配置 IB 设备由 Solaris 使用。

unconfigure

取消配置 IB 设备。如果成功取消配置 IB 设备，`cfgadm` 报告 ap_id 的 condition (状态) 为 unknown (未知)。

`-f`

不支持。

`-h ap_id`

为 IB 挂接点获取 IB 专用帮助。

-l

列出 IB 挂接点的状态和状况。按照 [cfgadm\(1M\)](#) 中的描述使用 -l 选项。

如果该选项与 -a 选项一起使用，也显示动态挂接点（包含 IOC 设备、VPPA 设备、端口设备、伪设备以及 HCA_SVC 设备）。

如果该选项与 -v 选项一起使用，显示有关 ap_id 的详细数据。对于一个 IOC，输出中的 Info（信息）字段

```
cfgadm -avl
```

显示 IOC 控制器配置文件如下信息：供应商 ID、IOC 设备 ID、设备版本、子系统供应商 ID、子系统 ID、类、子类、协议、协议版本以及 ID 字符串。如果 ID 字符串没有提供，该字段不会显示任何信息。这些字段在《InfiniBand Specification Volume 1》(<http://www.infinibandta.org>) 中进行了定义。

对于 VPPA 设备、端口设备或 HCA_SVC 设备，`cfgadm -lav` 中的 Info（信息）字段显示相应设备绑定的服务名称信息。如果相应设备没有绑定服务名称，该字段不显示任何信息。

对于伪设备，`cfgadm -alv` 显示驱动程序名称及其 unit-address 信息。对于 HCA，详细列表显示该 HCA 的供应商 ID、产品 ID、端口数量以及端口 PortGUID 值。请参见“示例”部分。

-o *hardware_option*

该选项当前未定义。

-s *listing_option*

可以通过 select 子选项列出 ib 类的挂接点。有关更多信息，请参见 [cfgadm\(1M\)](#) 手册页。

-x *hardware_function*

执行硬件专用函数。请注意，*name* 长度不能超过 4 个字符。

所支持的硬件专用函数如下：

```
add_service -ocomm=[port|vppa|hca_svc],service=name
```

静态 IB 挂接点支持该硬件专用函数。该函数可用于在 `/kernel/drv/ib.conf` 文件中添加一个新服务并且更新 [ib\(7D\)](#) 驱动程序。

必须使用 `service=name` 选项标识要添加的新服务。必须使用 `comm=[port|vppa|hca_svc]` 选项将 `name` 服务添加到 `port-svc-list` 或 `/kernel/drv/ib.conf` 文件中的 `hca-svc-list`。请参见“示例”部分。

```
delete_service -ocomm=[port|vppa|hca_svc],service=name
```

只有静态 IB 挂接点才支持该硬件专用函数。该函数用于将现有服务从 `/kernel/drv/ib.conf` 文件以及 [ib\(7D\)](#) 驱动程序的数据库中删除。必须使用 `service=name` 选项标识要删除的服务。必须使用 `comm=[port|vppa|hca_svc]` 选项将该服务从 `port-svc-list`、`vppa-svc-list` 或 `/kernel/drv/ib.conf` 文件中的 `vppa-svc-list` 中删除。请参见“示例”部分。

list_clients
HCA 挂接点支持该函数。显示所有使用该 HCA 的内核 IB 客户机。如果内核 IB 客户机已经打开了备用的 HCA 设备，该选项也显示这些内核 IB 客户机的相应 ap_id。请参见“示例”部分。

.

如果一个给定的内核 IB 客户机没有一个有效的 ap_id，该列显示 -。

list_services
只有静态 IB 挂接点才支持该硬件专用函数。该函数列出从 /kernel/drv/ib.conf 文件中读取的所有端口以及 VPPA 服务。请参见“示例”部分。

unconfig_clients
只有静态 HCA 挂接点才支持该硬件专用函数。该函数用于取消配置给定的 HCA 的所有 IB 内核客户机。只有不具备备用 HCA 的 IB 内核客户机才被取消配置。请参见“示例”部分。

update_ioc_config
静态 ib 挂接点以及 IOC 挂接点支持该硬件专用函数。对于 ib APID，该函数更新所有 IOC 设备节点的属性。对于 IOC APID，该函数更新指定的 IOC 设备节点的属性。该命令更新 port-list、port-entries、service-id 以及 service-name IOC 节点属性。

请参见 [ib\(7D\)](#)。

update_pkey_tbls
静态 ib 挂接点支持该函数。该函数更新 IBTL 中的 PKEY 信息。IBTL 重新读取主机上显示的每个 HCA 上的所有端口的 P_Key 表。

请参见 [ibt1\(7D\)](#)。

示例

示例 1 列出 IB 设备的状态和状况
以下命令列出系统中的 IB 设备的状态和状况。该命令只显示了静态挂接点。

```
example# cfgadm
hca:21346543210a987      IB-HCA      connected    configured    ok
ib                        IB-FABRIC    connected    configured    ok
```

-a 选项列出所有挂接点。以下示例使用 -a 选项列出所有挂接点：

```
example# cfgadm -a
hca:21346543210a987      IB-HCA      connected    configured    ok
ib                        IB-FABRIC    connected    configured    ok
ib::80020123456789a      IB-IOC      connected    unconfigured  ok
ib::80245678,ffff,ipib    IB-VPPA     connected    configured    ok
ib::21346543,0,hnfs       IB-HCA_SVC  connected    configured    ok
ib::12245678,0,nfs        IB-PORT     connected    configured    ok
ib::sdp,0                 IB-PSEUDO   connected    configured    ok
```

示例2 列出IB VPPA 设备的详细状态

以下命令列出 IB VPPA 设备的详细状态：

```
example# cfgadm -alv ib::80245678,ffff,ipib
Ap_Id      Receptacle Occupant  Condition Information
When      Type      Busy  Phys_Id
ib::80245678,ffff,ipib  connected  configured  ok      ipib
unavailable IB-VPPA  n      /devices/ib:fabric::80245678,ffff,ipib
```

IOC 的详细列表显示其他信息。以下命令显示详细列表：

```
example# cfgadm -alv ib::80020123456789a
Ap_Id      Receptacle  Occupant  Condition Information
When      Type      Busy  Phys_Id
ib::80020123456789a  connected  configured  ok      VID: 0xeaea
DEVID: 0xeaea VER: 0x5 SUBSYS_VID: 0x0 SUBSYS_ID: 0x0 CLASS: 0xffff
SUBCLASS: 0xff PROTO: 0xff PROTOVER: 0x1 ID_STRING: Sample Host Adapter
unavailable IB-IOC  n      /devices/ib:fabric::80020123456789a
```

伪设备的详细列表显示如下信息：

```
example# cfgadm -alv ib::sdp,0
Ap_Id      Receptacle Occupant  Condition Information
When      Type      Busy  Phys_Id
ib::sdp,0  connected  configured  ok      Driver = "sd
p" Unit-address = "0"
unavailable IB-PSEUDO n /devices/ib:fabric::sdp,0
```

HCA 的详细列表显示如下信息：

```
example# cfgadm -alv hca:21346543210a987
Ap_Id      Receptacle  Occupant  Condition Information
When      Type      Busy  Phys_Id
hca:21346543210a987  connected  configured  ok      VID: 0x15b3,
PID: 0x5a44, #ports: 0x2, port1 GUID: 0x80245678, port2 GUID: 0x80245679
unavailable IB-HCA  n /devices/ib:21346543210a987
```

如果指定以下 `cfgadm` 类以及字段选择选项，可以获取更多的用户友好的输出：`-s "select=class(ib),cols=ap_id:info"`。

以下命令只显示 IB 的 `ap_id`。该命令的输出只包含 `ap_id` 和 `Information`（信息）字段。

```
# cfgadm -al -s "cols=ap_id:info" ib::80245678,ffff,ipib
Ap_Id      Information
ib::80245678,ffff,ipib  ipib
```

示例3 取消配置现有的IBIOC

以下命令取消配置 `ib::80020123456789a` 连接的 IB IOC 并显示 `ap_id` 的状态。

示例3 取消配置现有的 IBIOC （续）

```
# cfgadm -c unconfigure ib::80020123456789a
Unconfigure the device: /devices/ib:fabric::80020123456789a
This operation will suspend activity on the IB device
Continue (yes/no)?
```

输入 y

```
IB device unconfigured successfully.
# cfgadm -al ib::80020123456789a
Ap_Id          Type      Receptacle  Occupant    Condition
ib::80020123456789  IB-IOC   connected   unconfigured unknown
#
```

如果 condition（状态）显示为 unknown（未知），表示该设备节点不存在并且只有 IB 设备管理器知道该 IB 设备的存在。

示例4 配置 IBIOC

以下一系列命令配置 ib::80020123456789a 连接的 IB 设备。

```
# cfgadm -yc configure ib::80020123456789a
# cfgadm -al ib::80020123456789a
Ap_Id          Type      Receptacle  Occupant    Condition
ib::80020123456789a  IB-IOC   connected   configured  ok
```

示例5 列出 HCA 的所有内核 IB 客户机

以下命令列出 hca:21346543210a987 连接的 HCA 中的所有内核 IB 客户机。

```
# cfgadm -x list_clients hca:21346543210a987
Attachment Point  Clients                      Alternate HCA
ib::80020123456789a  iocl                        Yes
ib::80245678,ffff,ipib  ipib                        No
ib::21346543,0,hnfs    hnfs                        No
-                    ibdm                        No
-                    ibmf                        No
```

示例6 添加端口服务

以下命令添加一个名为 srp 的新端口服务：

```
# cfgadm -o comm=port,service=srp -x add_service ib
```

示例7 删除 VPPA 服务

以下命令删除 ibd VPPA 服务 ibd：

```
# cfgadm -o comm=vppa,service=ipib -x delete_service ib
```

示例 8 列出端口服务、VPPA 服务以及 HCA_SVC 服务
以下命令列出所有的端口服务、VPPA 服务以及 HCA_SVC 服务：

```
# cfgadm -x list_services ib
Port communication services:
    srp

VPPA communication services:
    ipib
    nfs

HCA_SVC communication services:
    hnfs
```

示例 9 重新探测 IOC 设备
以下命令重新探测所有的 IOC 设备节点：

```
# cfgadm -x update_ioc_config ib
This operation can update properties of IOC devices.
Continue (yes/no)?

Enter: y

#
```

示例 10 取消配置 HCA 的所有内核客户机
以下命令取消配置 HCA 的所有内核客户机：

```
# cfgadm -x unconfig_clients hca:21346543
This operation will unconfigure clients of this HCA.
Continue (yes/no)?

Enter: y
```

文件 /usr/lib/cfgadm/ib.so.1
通用 InfiniBand 设备管理的硬件专用库

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/library

另请参见 [cfgadm\(1M\)](#)、[config_admin\(3CFGADM\)](#)、[libcfgadm\(3LIB\)](#)、[ib\(4\)](#)、[attributes\(5\)](#)、[ib\(7D\)](#)、[ibdm\(7D\)](#)
《InfiniBand Specification Volume 1》(<http://www.infinibandta.org>)

附注

除列出（`cfgadm -l` 或 `cfgadm -x list_clients`）外，只有超级用户才能在挂接点执行任意函数。

引用名 `cfgadm_pci` – PCI、CompactPCI 以及 PCI Express 热插拔硬件专用 `cfgadm` 命令

用法概要

```
/usr/sbin/cfgadm [-f] [-y | -n] [-v]
                [-o hardware_options] -c function ap_id [ap_id]

/usr/sbin/cfgadm [-f] [-y | -n] [-v]
                [-o hardware_options] -x hardware_function ap_id
                [ap_id]

/usr/sbin/cfgadm [-v] [-s listing_options]
                [-o hardware_options] [-l [ap_id | ap_type]]

/usr/sbin/cfgadm [-v] [-o hardware_options] -t ap_id [ap_id]

/usr/sbin/cfgadm [-v] [-o hardware_function] -h
                [ap_id| ap_type]
```

描述

PCI 硬件专用库 `/usr/lib/cfgadm/pci.so.1` 支持通过 `cfgadm` 命令（请参见 [cfgadm\(1M\)](#)）在支持热插拔的系统中相应可热插拔的插槽中热插拔 PCI 和 CompactPCI 适配卡的功能。该库不支持 PCI Express 热插拔适配卡或标准的 PCI 热插拔适配卡，这些适配卡由另外一个库支持（请参见 [cfgadm_shp\(1M\)](#)）。除了该文档中标注的部分，PCI 与 CompactPCI 之间的热插拔管理模式保持不变。

对于 PCI 热插拔，每个特定的 PCI 总线上的热插拔插槽由该 PCI 总线上的挂接点表示。

一个挂接点由一个插口和一个插卡组成。PCI 热插拔中的 `receptacle`（插口）通常指物理的热插拔插槽，`occupant`（插卡）指插入到插槽中的 PCI 适配卡。

挂接点通过 `ap_id` 命名。`ap_id` 包括逻辑 `ap_id` 和物理 `ap_id`。物理 `ap_id` 基于物理路径名，即 `/devices/pci@1/hpc0_slot3`，而逻辑 `ap_id` 是更短、更加用户友好的名称。对于 PCI 热插拔插槽，逻辑 `ap_id` 的名称通常为对应的热插拔控制器驱动程序名称加上逻辑槽号，即 `pci0:hpc0slot1`，也就是说，PCI nexus 驱动程序、名为 `hpc` 的热插拔控制器驱动程序以及槽号 1。PCI 热插拔的 `ap_type` 为 `pci`。

请注意，`ap_type` 与 `Type`（类型）字段中的信息不同。

有关热插拔过程的详细说明，请参见《[Oracle Solaris 管理：常见任务](#)》。

选项

所支持的选项如下：

```
-c function
    以下 function 支持 PCI 热插拔插槽：

    configure
        将 PCI 设备配置在 Solaris 使用的插槽中。

    connect
        将插槽与 PCI 总线相连。

    disconnect
        断开插槽与 PCI 总线之间的连接。
```


- `insert`
不支持。
- `remove`
不支持。
- `unconfigure`
将 PCI 设备的资源从系统中逻辑移除。
- `-f`
不支持。
- `-h ap_id | ap_type`
列出 PCI 热插拔专用帮助消息。
- `-l list`
列出 PCI 热插拔插槽的值。
- `-o hardware_options`
当前没有定义任何硬件专用选项。
- `-s listing_options`
与通用的 `cfgadm(1M)` 相同。
- `-t ap_id`
该命令只适用于支持插槽测试功能的平台。
- `-v`
在详细模式下执行。

如果 `-v` 选项与 `-l` 选项一起使用，`cfgadm` 命令输出有关挂接点的信息。有关位于 PCI Express 分层结构中的 PCI 热插拔挂接点的详细信息，请参见 `cfgadm_shp(1M)`。如果 PCI 热插拔挂接点不在 PCI Express 分层结构中，`Information`（信息）字段为该插槽的系统标签（如果有）。该字符串可从该插槽的总线节点的 `slot-name` 属性中获取。`Type`（类型）字段中的信息可使用 `v` 选项列出，也可不使用该选项列出。插卡的 `Type`（类型）字段描述有关插槽的信息。有以下两个可能值：

`unknown`（未知）

该插槽为空。如果该插槽中安装了卡，表明该卡未配置或卡上的设备未安装驱动程序。

`subclass/board`

该插槽中的卡是一个单功能或多功能设备。

`subclass` 是表示该设备的子类代码的字符串，如 `SCSI`、`ethernet`、`pci-isa` 等。如果该卡是一个多功能设备，则列出 `MULT`。

`board` 是表示设备的板类型的字符串。例如，`hp` 是表示 PCI 热插拔适配器的字符串，`hs` 是表示热交换板的字符串，`nhs` 是表示非热交换 cPCI 板的字符串，`bhs` 是表示基本热交换 cPCI 板的字符串，`fhs` 是表示完全热交换 cPCI 板的字符串。

大多数配有多个设备的 PCI 卡并不是多功能设备，但这些 PCI 卡是作为其后面的任意设备的 PCI 桥实现的。在这些情况下，显示的子类为 PCI 桥的子类。最常见的 PCI 桥为通用 PCI-PCI 桥 `pci-pci` 或半透明的 PCI 桥 `stpci`。

`-x hardware_function`

执行硬件专用函数。通常，这些硬件专用函数不应更改插口或插卡的状态。

所支持的 `hardware_function` 如下：

`enable_slot | disable_slot`

更改插槽状态并在重新引导后仍保留插槽状态。只有选定的平台在重新引导后才能保留状态。

`enable_slot` 允许在引导时将硬件添加到热插拔插槽。

`disable_slot` 禁止在引导时将硬件添加到热插拔插槽。如果插槽被禁用，该插槽的状态显示为不可用。

`enable_autoconfig | disable_autoconfig`

更改插槽上插卡的自动配置的功能。只有支持自动配置的平台才支持该功能。

`enable_autoconfig` 启用插槽自动配置的功能。

`disable_autoconfig` 禁用插槽自动配置的功能。

利用 PCI Express 平台上的警示按钮以及 CompactPCI 平台上的卡入/弹出锁存器可以完成自动配置。如果自动配置被禁用，警示按钮或锁存器装置不能用于配置插槽上的插卡。

`led=[led_sub_arg],mode=[mode_sub_arg]`

如果不带子参数，列出当前 LED 指示灯设置的列表。如果带有子参数，则设置插槽的专用 LED 指示灯的模式。

将 `led_sub_arg` 指定为 `fault`、`power`、`attn` 或 `active`。

将 `mode_sub_arg` 指定为 `on`、`off` 或 `blink`。

更改 LED 指示灯的状态不会更改插口或插卡的状态。通常，LED 指示灯由热插拔控制器控制，无需用户干预。使用该命令用于测试目的。

注意：更改 LED 指示灯的状态可能歪曲插卡或插口的状态。

以下命令列出 LED 指示灯的值：

```
example#  cfgadm -x led pci0:hpc0_slot1
Ap_Id          Led
pci0:hpc0_slot1  power=on,fault=off,active=off,attn=off
```

以下命令开启故障 LED 指示灯：

```
example#  cfgadm -x led=fault,mode=on pci0:hpc0_slot1
```

以下命令关闭电源 LED 指示灯：

```
example# cfgadm -x led=power,mode=off pci0:hpc0_slot0
```

以下命令将**活动的** LED 指示灯设置为闪烁状态以指示该插槽的位置：

```
example# cfgadm -x led=active,mode=on pci0:hpc0_slot3
```

示例

示例1 列出每个插槽的值

以下命令列出每个插槽的值：

```
example# cfgadm -l
```

Ap_Id	Type	Receptacle	Occupant	Condition
c0	scsi-bus	connected	configured	unknown
c1	scsi-bus	connected	unconfigured	unknown
c2	scsi-bus	connected	unconfigured	unknown
cpci_slot1	stpci/fhs	connected	configured	ok
cpci_slot2	unknown	empty	unconfigured	unknown
cpci_slot4	stpci/fhs	connected	configured	ok
cpci_slot5	stpci/fhs	connected	configured	ok

示例2 替换卡

以下命令列出所有支持 DR 的挂接点：

```
example# cfgadm
```

Type	Receptacle	Occupant	Condition
c0	scsi-bus	connected	configured unknown
c1	scsi-bus	connected	unconfigured unknown
c2	scsi-bus	connected	unconfigured unknown
cpci_slot1	stpci/fhs	connected	configured ok
cpci_slot2	unknown	empty	unconfigured unknown
cpci_slot4	stpci/fhs	connected	configured ok
cpci_slot5	stpci/fhs	connected	configured ok

以下命令取消配置卡并使卡断电：

```
example# cfgadm -c disconnect cpci_slot4
```

该更改可通过输入以下命令验证：

```
example# cfgadm cpci_slot4
```

Ap_Id	Type	Receptacle	Occupant	Condition
cpci_slot4	unknown	disconnected	unconfigured	unknown

现在可以交换卡了。以下命令使卡接电并配置卡：

```
example# cfgadm -c configure cpci_slot4
```

示例2 替换卡 (续)

该更改可通过输入以下命令验证：

example# **cfgadm cpci_slot4**

Ap_Id	Type	Receptacle	Occupant	Condition
cpci_slot4	stpcipci/fhs	connected	configured	ok

文件

/usr/lib/cfgadm/pci.so.1
PCI 热插拔硬件专用库。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/library

另请参见

[cfgadm\(1M\)](#)、[cfgadm_shp\(1M\)](#)、[config_admin\(3CFGADM\)](#)、[libcfgadm\(3LIB\)](#)、[attributes\(5\)](#)
《Oracle Solaris 管理：常见任务》

引用名	cfgadm_sata – SATA 硬件专用 cfgadm 命令										
用法概要	<pre>/usr/sbin/cfgadm [-f] [-y -n] [-v] [-o hardware_options] -c function ap_id... /usr/sbin/cfgadm [-f] [-y -n] [-v] [-o hardware_options] -x hardware_function ap_id... /usr/sbin/cfgadm [-v] [-a] [-s listing_options] [-o hardware_options] [-l [ap_id ap_type]...] /usr/sbin/cfgadm [-v] [-o hardware_options] -t ap_id... /usr/sbin/cfgadm [-v] [-o hardware_options] -h [ap_id]...</pre>										
描述	SATA 硬件专用库 <code>/usr/lib/cfgadm/sata.so.1</code> 提供通过 <code>cfgadm</code> 命令热插拔 SATA 的功能。<code>cfgadm</code> 在挂接点上运行，这些挂接点是可动态重新配置硬件资源的系统位置。有关挂接点的信息，请参见 cfgadm(1M)。 每个 SATA 控制器以及端口复用器的设备端口由设备树上的挂接点表示。系统中连接并配置的 SATA 设备显示为挂接点名称的扩展名。术语“挂接点”和“SATA 端口”可在下列描述中互换使用。 挂接点可通过 <code>ap_id</code> 命名。所有的 SATA 挂接点 <code>ap_id</code> 由以下格式的字符串组成： <pre>sataX/P[.M][::dsk/cXtYd0]</pre> 其中 <table><tr><td>X</td><td>表示 SATA 控制器号</td></tr><tr><td>P</td><td>表示 SATA 控制器的设备端口号（0 到 31）</td></tr><tr><td>M</td><td>表示端口复用器的设备端口号（0 到 14）。如果该项的值为 15，表示端口复用器的主机端口号。只有端口复用器连接到 SATA 控制器的设备端口时，该项才使用。</td></tr><tr><td>dev/cXtYd0</td><td>标识连接的 SATA 设备</td></tr><tr><td>Y</td><td>表示目标号</td></tr></table> 通常，设备标识符源自 <code>/dev</code> 中的设备的相应逻辑链接。由于 SATA 设备只支持一个 LUN (LUN 0)，因此，设备字符串中的“d”字符总是设置为 0（零）。 例如，SATA 控制器 2 的设备端口 5 所连接的端口复用器的设备端口 4 的逻辑 <code>ap_id</code> 可能为： <pre>sata2/5.4</pre> 如果 SATA 磁盘或 CD/DVD 设备连接到该挂接点并且该设备已配置，<code>ap_id</code> 可能为： <pre>sata2/5.4::dsk/c2t645d0</pre> 标识设备的 <code>cXtYd0</code> 字符串与设备挂接点之间存在一对一的对应关系。	X	表示 SATA 控制器号	P	表示 SATA 控制器的设备端口号（0 到 31）	M	表示端口复用器的设备端口号（0 到 14）。如果该项的值为 15，表示端口复用器的主机端口号。只有端口复用器连接到 SATA 控制器的设备端口时，该项才使用。	dev/cXtYd0	标识连接的 SATA 设备	Y	表示目标号
X	表示 SATA 控制器号										
P	表示 SATA 控制器的设备端口号（0 到 31）										
M	表示端口复用器的设备端口号（0 到 14）。如果该项的值为 15，表示端口复用器的主机端口号。只有端口复用器连接到 SATA 控制器的设备端口时，该项才使用。										
dev/cXtYd0	标识连接的 SATA 设备										
Y	表示目标号										

系统挂接点的简单列表包含所有的 SATA 设备端口和连接的设备。例如：

```
#cfgadm -l
Ap_Id                                Type      Receptacle  Occupant    Condition
sata0/0::dev/c0t0d0                 disk      connected   configured  ok
sata0/1::dev/c0t1d0                 disk      connected   configured  ok
sata0/2::dev/c0t2d0                 cd-dvd    connected   configured  ok
sata0/3                             sata-port empty       unconfigured ok
sata1/0                             sata-port disconnected unconfigured unknown
sata1/1                             sata port  disconnected unconfigured unknown
sata1/2                             sata port  empty       unconfigured ok
sata1/3.15                          sata-pmult connected   configured  ok
sata1/3.0::dev/c0t512d0             disk      connected   configured  ok
sata1/3.1                           sata-port  empty       unconfigured ok
sata1/3.2                           sata-port  empty       unconfigured ok
sata1/3.3                           sata-port  empty       unconfigured ok
usb0/1                              unknown    empty       unconfigured ok
usb0/2                              unknown    empty       unconfigured ok
```

有关挂接点列表的更多信息，请参见 [cfgadm\(1M\)](#)。

SATA 端口上的挂接点的插口状态的含义如下：

- empty（空） SATA 端口已加电并启用。在该端口未检测到任何设备。
- disconnected（断开） SATA 端口未启用或检测到 SATA 设备但该设备的通信连接未建立，或者该端口故障。
- connected（已连接） 在与设备建立通信的端口上检测到 SATA 设备。

插卡（SATA 端口连接的设备）状态的含义如下：

- configured（已配置） 所连接的 SATA 设备已配置并且随时可由操作系统使用。
- unconfigured（未配置） 未连接任何设备或 SATA 端口连接的 SATA 设备尚未配置。如果要配置该设备，请运行 "cfgadm -c configure ap_id" 命令。

挂接点（SATA 端口）的状态的含义如下：

- ok（正常） SATA 端口已加电并启用，随时可用。
- failed（失败） SATA 端口故障。该端口被系统禁用和/或断电。这种情况下，端口不可用且处于未知状态。这可能是由插入的设备导致的。
- unknown（未知） SATA 端口被禁用且处于 unknown（未知）状态。

“状态表”包含挂接点的插口状态、插卡状态以及挂接点（SATA 端口）状态。有效的状态如下：

empty/unconfigured/ok	SATA 端口已启用且处于活动状态。未检测到任何设备。
disconnected/unconfigured/ok	SATA 端口已启用并且检测到一个设备，但该设备的通信连接未建立。
disconnected/unconfigured/unknown	SATA 端口被禁用且处于未知状态。
disconnected/unconfigured/failed	SATA 端口被禁用且不可用。由于被检测的系统出现故障，该端口被系统禁用。
connected/unconfigured/ok	SATA 端口已启用且处于活动状态。检测到一个设备并且该设备的通信连接已建立。该设备没有配置为由操作系统使用。
connected/configured/ok	该设备存在且已配置，随时可由操作系统使用。

选项

除列出 (-l)， `cfgadm` 外， 还定义了几种类型的操作。这些操作包括测试 (-t)、调用配置状态更改 (-c)、调用硬件专用函数 (-x) 以及获取配置管理帮助消息 (-h)。

-c *function*

以下通用的 *function* 是为 SATA 硬件专用库定义的。对于 SATA 端口挂接点， 支持以下配置状态更改操作：

connect

启用（激活）SATA 端口并建立与所连接的设备之间的通信。该操作意味着给端口加电（如有必要）。

disconnect

如果尚未取消配置该设备， 则取消配置所连接的设备并禁用（取消激活）SATA 端口。后面的 "connect" 命令启用 SATA 端口操作但并不会使设备处于 "configured"（已配置）状态。

对于 SATA 端口连接的 SATA 设备， 支持以下状态更改操作：

configure 配置新设备以便由操作系统使用（如果尚未配置）。该命令意味着执行连接操作（如有必要）。

unconfigure 如果尚未取消配置 SATA 端口连接的设备， 则取消配置该设备。

configure 与 **unconfigure** 操作不能用于端口复用器连接的挂接点。端口复用器由系统自动配置以及取消配置。但是， 配置与取消配置操作适用于端口复用器的设备端口连接的所有 SATA 设备。

-f

不支持。

-h *ap_id*

SATA 专用帮助可通过使用帮助选项以及任何 SATA 挂接点获取。

-l [-v]

按照 [cfgadm\(1M\)](#) 中的描述使用 **-l** 选项。如果该选项与 **-v** 选项一起使用，"Information"（信息）字段包含以下 SATA 专用信息：

- Mfg：生产商字符串
- Product：产品字符串
- No：产品序列号

-o hardware_options

当前没有定义任何硬件专用选项。

-s listing_options

可使用 **select** 子选项列出 SATA 类的挂接点。请参见 [cfgadm\(1M\)](#)。

-t ap_id

如果 SATA 控制器支持自检操作，执行 SATA 端口自检。如果 SATA 控制器不支持端口自检操作，将发送错误消息。

-x hardware_function

执行硬件专用函数。

SATA 端口或 SATA 控制器上使用的下列某些命令可能影响所连接的 SATA 设备，请参见相应说明。**ap_id** 指 SATA 端口或者整个 SATA 控制器，请参见相应说明。如果该操作意味着取消配置设备但该设备不能被取消配置（即该设备包含一个挂载的文件系统），将发送错误消息并且不执行该操作。如果 SATA 控制器不支持指定的操作，也将发送错误消息。

sata_reset_device ap_id

重置 **ap_id** SATA 端口连接的 SATA 设备。SATA 端口状态不会更改。

sata_reset_port ap_id

重置由 **ap_id** 指定的 SATA 端口。如果 SATA 设备与端口相连接，该设备也会重置。该操作也可能在端口复用器连接的端口上执行。如果端口复用器连接到 SATA 控制器端口，端口复用器连接的 SATA 设备可能不会重置。

sata_reset_all ap_id

重置 **ap_id** 中的控制器号指定的 SATA 控制器以及所有连接的设备并重新枚举所有连接的设备，包括端口复用器以及端口复用器的设备端口所连接的设备。

该操作意味着执行该操作前取消配置所有连接的设备。任何新枚举的设备将处于未配置状态。

sata_port_deactivate ap_id

如果所有方法都无效，强制取消端口激活。该操作为紧急步骤，因此使用时需谨慎。

sata_port_activate ap_id

强制激活端口。该操作用于在紧急情况下激活端口（之前被取消激活以便从错误中恢复）。

sata_port_self_test ap_id

在 SATA 控制器上执行自检操作。该操作意味着取消配置所有设备并重置 SATA 控制器。

-v

在详细模式下执行。

以下转换表报告由 -c 操作以及热插拔操作导致的状态转换：

current state	operation	possible new state
-----	-----	-----
empty/ unconfigured/ok	device plug-in	connected/unconfigured/ok, or disconnected/unconfigured/ok, or disconnected/unconfigured/failed
empty/ unconfigured/ok	-c unconfigure	error message, no state change
empty/ unconfigured/ok	-c configure	error message, no state change
empty/ unconfigured/ok	-c connect	error message, no state change
empty/ unconfigured/ok	-c disconnect	disconnected/unconfigured/unknown, or disconnected/unconfigured/failed
disconnected/ unconfigured/ok	device unplug	no state change
disconnected/ unconfigured/ok	-c unconfigure	error message, no state change
disconnected/ unconfigured/ok	-c configure	error message, no state change
disconnected/ unconfigured/ok	-c connect	error message, no state change
disconnected/ unconfigured/ok	-c disconnect	error message, no state change
disconnected/ unconfigured/ unknown (no disk plugged)	-c configure	error message, state change to empty/unconfigured/ok, or

		disconnected/unconfigured/failed
disconnected/ unconfigured/ unknown (disk plugged)	-c configure	state change to connected/configured/ok or, connected/unconfigured/ok, or disconnected/unconfigured/failed and possible error message
disconnected/ unconfigured/ unknown	-c connect	empty/unconfigured/ok, or connected/unconfigured/ok, or disconnected/unconfigured/ok, or disconnected/unconfigured/unknown, or disconnected/unconfigured/failed
disconnected/ unconfigured/ unknown	-c disconnect	error message, no state change
disconnected/ unconfigured/ failed	any command other than -x commands	error message, no state change
connected/ unconfigured/ok	disk unplug	error message and state: empty/unconfigured/ok, or disconnected/unconfigured/failed
connected/ unconfigured/ok	-c configure	connected/unconfigured/ok, or connected/configured/ok, or disconnected/unconfigured/ok, or disconnected/unconfigured/failed
connected/ unconfigured/ok	-c unconfigure	error message, no state change
connected/ unconfigured/ok	-c connect	error message, no state change
connected/ unconfigured/ok	-c disconnect	disconnected/unconfigured/unknown, or disconnected/unconfigured/failed

connected/ configured/ok	disk unplug	error message and state: empty/unconfigured/ok, or disconnected/unconfigured/failed
connected/ configured/ok	-c configure	error message, no state change
connected/ configured/ok	-c unconfigure	error message, if device cannot be unconfigured, no state change, or connected/unconfigured/ok, or disconnected/unconfigured/ok, or disconnected/unconfigured/failed
connected/ configured/ok	-c connect	error message, no state change
connected/ configured/ok	-c disconnect	error message, if device cannot be unconfigured, no state change, or disconnected/unconfigured/unknown, or disconnected/unconfigured/failed

示例

示例1 配置磁盘

以下命令配置 SATA 控制器 0 上的端口 0 连接的磁盘：

```
example# cfgadm -c configure sata0/0
```

当设备连接到 SATA 端口时，发出该命令。

示例2 取消配置磁盘

以下命令取消配置 SATA 控制器 0 上的端口 3 连接的磁盘：

```
example# cfgadm -c unconfigure sata0/3
```

当挂接点的插口状态为 "connected"（已连接）并且插卡状态为 "configured"（已配置）时，显示设备标识字符串。

示例3 取消配置磁盘时遇到已挂载的文件系统

以下命令表明在取消配置磁盘时遇到已挂载的文件系统：

```
example# cfgadm -c unconfigure sata1/5
```

系统作出如下响应：

```
cfgadm: Component system is busy, try again: failed to offline:  
/devices/pci@0,0/pci8086,244e@1e/pci1095,3124@1/sd@5,0
```

示例 3 取消配置磁盘时遇到已挂载的文件系统 (续)

文件属性

Resource	Information
/dev/dsk/c1t5d0s0	mounted filesystem "/mnt"
/usr/lib/cfgadm/sata.so.1	通用 SATA 热插拔硬件专用库。

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/library

另请参见

[cfgadm\(1M\)](#)、[config_admin\(3CFGADM\)](#)、[libcfgadm\(3LIB\)](#)、[attributes\(5\)](#)

附注

如果端口所连接的磁盘包含关键分区，如根目录 (/)、/usr、swap 或者 /var，则该端口不支持 "sata_port_deactivate" 紧急操作。不应在这些端口尝试取消激活操作。不正确的用法可能导致系统挂起且需要重新引导。

并非所有的 SATA 控制器都支持热插拔操作。

如果 SATA 连接器为可热插拔类型并且 SATA 控制器支持热插拔，则 SATA 设备在任何时候都可进行热插拔。系统检测事件并建立与设备之间的通信。该设备需要通过 "cfgadm -c configure ap_id" 显式命令进行配置。

如果 SATA 连接器为可热插拔类型并且 SATA 控制器支持热插拔，拔下设备前需取消配置设备，否则，可能导致系统挂起或数据丢失。如果设备是未配置的但插口不处于断开状态，从 SATA 端口拔下设备将导致错误消息。

警告

某些 SATA 设备上的连接器不符合 SATA 热插拔规范。在这些设备上执行热插拔操作可能损坏 SATA 控制器和/或 SATA 设备。

引用名	cfgadm_sbd – cfgadm 命令，用于系统板管理
用法概要	<pre>cfgadm -l [-a] [-o parsable] ap_id... cfgadm -c function [-f] [-y -n] [-o unassign nopoweroff] [-v] ap_id... cfgadm -t [-v] ap_id... cfgadm -x [-f] [-v] function ap_id...</pre>
描述	cfgadm_sbd 插件提供动态重新配置功能，适用于 sbd 类系统板的连接、配置、取消配置和断开连接操作。此外，您还可以通过它连接系统板或断开其与正在运行的系统的连接而不必重新引导系统。 cfgadm 命令驻留在 /usr/sbin 中。请参见 cfgadm(1M)。cfgadm_sbd 插件驻留在 /usr/platform/sun4u/lib/cfgadm 中。 每个板槽都会在设备树中显示为单个挂接点。每个组件都会显示为一个动态挂接点。您可以使用 -a 选项查看每个组件的类型、状态和状况，以及每个板槽的状态和状况。 cfgadm 选项的执行情况取决于具体的平台。此外，挂接点的形式在不同的平台也各不相同。有关更多信息，请参见“平台说明”部分。
组件状况	以下为组件状况的名称和描述： failed（失败） 组件测试失败。 ok（正常） 组件运行正常。 unknown（未知） 组件未经测试。
组件状态	以下为组件插口状态的名称和描述： connected（连接） 组件已连接到板槽。 以下为组件插卡状态的名称和描述： configured（已配置） 组件可供 Solaris 操作环境使用。 unconfigured（未配置） 组件无法供 Solaris 操作环境使用。

板状况 以下为板状况的名称和描述：

failed（失败）
板测试失败。

ok（正常）
板运行正常。

unknown（未知）
板未经测试。

unusable（不可用）
板槽不可用。

板状态 插入板时，插口状态将从 **empty**（空）更改为 **disconnected**（断开）。移除板时，插口状态将从 **disconnected**（断开）更改为 **empty**（空）。

注意：移除处于 **connected**（已连接）状态或虽已通电但仍处于 **disconnected**（断开）状态的板会导致操作系统崩溃，并可能造成系统的永久损坏。

以下为板的插口状态的名称和描述：

connected（已连接）
板已通电并连接到系统总线。您只能在板进入 **connected**（已连接）状态以后才能查看其上的组件。

disconnected（断开）
板已断开与系统总线的连接。板在不断电的情况下，也可以处于 **disconnected**（断开）状态。不过，在将板从插槽中移除之前，该板必须已断电并处于 **disconnected**（断开）状态。

empty（空）
板不存在。

断开的板的插卡状态始终为 **unconfigured**（未配置）。下表包含板的插卡状态的名称和描述：

configured（已配置）
板上至少有一个组件已配置。

unconfigured（未配置）
板上的所有组件均未配置。

动态系统域 基于动态系统域 (dynamic system domain, DSD)（在本文档中简称域）的平台可以将机箱中的插槽分隔成以电子方式隔离的硬件分区（即 DSD）。不基于 DSD 的平台可将所有插槽永久性地分配给系统。

插槽可以为空，也可以处于填充状态，还可以分配或提供给任何数目的域使用。可供给定域使用的插槽数由在系统控制器上进行维护的可用组件列表 (ACL) 控制。ACL 不是指由 Solaris 操作环境提供的访问控制列表。

插槽仅在位于域的 ACL 中且尚未分配给其他域的情况下才对该域可见。未分配的插槽对于其 ACL 中包含该插槽的所有域均可见。将插槽分配给某个域以后，该插槽对于其他任何域都不再可见。

如果插槽对于某个域可见但尚未分配，则必须先将该插槽分配给该域，然后才能应用其他任何状态更改命令。分配操作可以通过 `-x assign` 显式完成，也可以在连接过程中隐式完成。必须先取消分配给某个域的插槽，然后才能允许其他域使用该插槽。取消分配始终是显式的，可以直接使用 `-x unassign` 取消分配，也可以使用 `-o unassign` 以选项方式断开连接。

状态更改函数

更改板槽状态或板上组件状态的函数可以针对任何挂接点并行发出。在某个给定时间，只允许一个状态更改操作。状态更改信息中 "Busy"（忙）字段中的 Y 表示某个操作正在进行中。

以下列表包含可以进行状态更改的函数：

- `configure`
- `unconfigure`
- `connect`
- `disconnect`

可用性更改函数

更改板可用性的命令可以针对任何挂接点并行发出。在某个给定时间，只允许一个可用性更改操作。这些函数也更改 `cfgadm -l` 输出中的信息字符串。"Busy"（忙）字段中的 Y 表示某个操作正在进行中。

以下列表包含可以进行可用性更改的函数：

- `assign`
- `unassign`

状况更改函数

更改板槽状况或板上组件状况的函数可以针对任何挂接点并行发出。在某个给定时间，只允许一个状况更改操作。这些函数也更改 `cfgadm -l` 输出中的信息字符串。"Busy"（忙）字段中的 Y 表示某个操作正在进行中。

以下列表包含可以进行状况更改的函数：

- `poweron`
- `poweroff`
- `test`

取消配置过程

本节包含取消配置过程的描述，并说明了在移动常驻内存过程中的不同阶段源板和目标板的状态。

在以下代码示例中，板 0 上的常驻内存必须移至域中的其他板。因此，板 0 为源板，板 1 为目标板。

当系统将板的状态标记为忙时，不能在该板上启动状态更改操作。为了简便起见，已删除代理示例中的 CPU 信息。

此过程通过以下命令开始：

```
# cfgadm -c unconfigure -y SB0::memory &
```

首先，必须删除板 1 上与板 0 上的常驻内存处于同一地址范围内的内存。在此阶段，源板、目标板和内存挂接点均标记为忙。可以使用以下命令来显示状态：

```
# cfgadm -a -s cols=ap_id:type:r_state:o_state:busy SB0 SB1
```

Ap_Id	Type	Receptacle	Occupant	Busy
SB0	CPU	connected	configured	y
SB0::memory	memory	connected	configured	y
SB1	CPU	connected	configured	y
SB1::memory	memory	connected	configured	y

删除板 1 上的相应内存以后，该内存将标记为 **unconfigured**（未配置）。板 0 上的内存始终为 **configured**（已配置），但仍会标记为忙，如下例所示。

Ap_Id	Type	Receptacle	Occupant	Busy
SB0	CPU	connected	configured	y
SB0::memory	memory	connected	configured	y
SB1	CPU	connected	configured	y
SB1::memory	memory	connected	unconfigured	n

然后，板 0 上的内存将复制到板 1。复制完成以后，将切换内存的插卡状态。板 0 上的内存变为 **unconfigured**（未配置），而板 1 上的内存变为 **configured**（已配置）。在该过程的这一阶段，仅板 0 的状态仍为忙，如下例所示。

Ap_Id	Type	Receptacle	Occupant	Busy
SB0	CPU	connected	configured	y
SB0::memory	memory	connected	unconfigured	n
SB1	CPU	connected	configured	n
SB1::memory	memory	connected	configured	n

整个过程完成以后，板 0 上的内存仍为 **unconfigured**（未配置），但挂接点的状态不为忙，如下例所示。

Ap_Id	Type	Receptacle	Occupant	Busy
SB0	CPU	connected	configured	n
SB0::memory	memory	connected	unconfigured	n
SB1	CPU	connected	configured	n
SB1::memory	memory	connected	configured	n

常驻内存已移动，板 0 上的内存已取消配置。此时，您可以在任意一个板上启动新的状态更改操作。

特定于平台的选项 您可以紧随由系统板插件进行解释的选项之后指定特定于平台的选项。所有特定于平台的选项必须以 `platform` 关键字开头。以下示例包含了带特定于平台的选项的命令的通用格式：

```
command -o sbd_options,platform=platform_options
```

选项 此手册页不包括 `cfgadm` 命令的 `-v`、`-a`、`-s` 或 `-h` 选项。有关这些选项的描述，请参见 `cfgadm(1M)`。以下选项受 `cfgadm_sbd` 插件支持：

-c *function*

执行状态更改函数。您可以使用以下函数：

unconfigure

将插卡状态更改为 `unconfigured`（未配置）。此函数适用于系统板槽以及系统板上的所有组件。

`unconfigure` 函数可从 CPU 列表中移除 CPU，并从系统内存池中删除物理内存。如果有设备仍处于使用状态，`cfgadm` 命令会失败并将失败情况报告给用户。只要设备不再处于忙的状态，您就可以立即重试该命令。如果 CPU 仍在使使用，则在继续操作之前，必须确保该 CPU 已关闭。请参见 [pbind\(1M\)](#)、[psradm\(1M\)](#) 和 [psrinfo\(1M\)](#)。

`unconfigure` 函数会将物理内存移至其他系统板，然后才会将该内存从要取消配置的板上删除。如果无法在其他板上找到足够的内存，或者无法找到合适的物理内存范围，命令将失败，具体取决于要移动的内存的类型。

对于常驻内存，必须在移动内存以及对内存控制器进行重新编程时暂停（即停止）操作系统。如果必须暂停操作系统，系统会提示您执行该操作。您始终可以使用 `-y` 或 `-n` 选项来回答（分别代表“是”或“否”）。

移动内存可能需要数分钟才能完成，具体取决于内存量和系统负载。您可以通过对内存挂接点发出状态命令来监视操作进度。您也可以通过停止 `cfgadm` 命令来中断内存操作。删除的内存将返回给系统内存池。

disconnect

将插口状态更改为 `disconnected`（断开）。此函数仅适用于系统板槽。

如果插卡状态为 `configured`（已配置），`disconnect` 函数会尝试取消对该插卡的配置，然后才会关闭系统板的电源。此时可以将板从插槽中移除。

应用此函数后，板在支持动态系统域的平台上将处于 `assigned`（已分配）状态。

如果指定 `-o nopoweroff`，则 `disconnect` 函数将使板处于通电状态。如果指定 `-o unassign`，则 `disconnect` 函数将从域中取消对板的分配。

如果从域中取消对板的分配，可将其分配给其他域。不过，如果将其分配给其他域，该板将无法提供给取消分配操作所发生的域使用。

configure

将插卡状态更改为 **configured**（已配置）。此函数适用于系统板槽以及系统板上的任意组件。

如果插口状态为 **disconnected**（断开），**configure** 函数会尝试连接该插口。然后，它会遍历由 **connect** 函数创建的设备树，并在必要情况下附加这些设备。运行该函数可以对板上的所有组件进行配置，已经配置的组件除外。

对于 CPU，**configure** 函数可将 CPU 添加到 CPU 列表。对于内存，**configure** 函数会在确保内存已初始化的基础上将其添加到系统内存池。CPU 和内存存在 **configure** 函数成功完成后即可使用。

对于 I/O 设备，必须在使用设备之前先使用 **mount** 和 **ifconfig** 命令。请参见 [ifconfig\(1M\)](#) 和 [mount\(1M\)](#)。

connect

将插口状态更改为 **connected**（已连接）。此函数仅适用于系统板槽。

如果板槽尚未分配给域，**connect** 函数会尝试将该板槽分配给域。接下来，该函数会对板执行通电和测试操作，然后将板以电子方式连接到系统总线并对组件进行探测。

成功完成 **connect** 函数以后，可使用 **-a** 选项来查看板上组件的状态。**connect** 函数可以让所有组件都处于 **unconfigured**（未配置）状态。

赋值步骤仅适用于支持动态系统域的平台。

-f

覆盖软件状态更改约束。

-f 选项从不覆盖硬件和操作系统的根本安全性和可用性约束。

-l

列出按 **-s**、**-v** 和 **-a** 选项（详见 [cfgadm\(1M\)](#)）控制格式指定的挂接点的状态和状况。**cfgadm_sbd** 插件在 **info**（信息）字段中提供了具体的信息，如下所述。该信息的格式可通过 **-o parsable** 选项进行更改。

可解析的 **info**（信息）字段由以下内容组成：

cpu

cpu 类型显示以下信息：

cpuid=#[,#...]

其中 **#** 为数字，表示 CPU 的 ID。如果存在多个 **#**，则表示该 CPU 具有多个活动的虚拟处理器。

speed=#

其中 **#** 为数字，表示 CPU 的速度（单位：MHz）。

ecache=#

其中 # 为数字，表示 **ecache** 的大小（单位：MB）。如果 CPU 具有多个活动的虚拟处理器，则 **ecache** 既可在这些虚拟处理器之间共享，也可在其间进行划分。

存储区

memory 类型会根据情况显示以下信息：

address=#

其中 # 为数字，表示基本物理地址。

size=#

其中 # 为数字，表示内存大小（单位：KB）。

永久的=#

其中 # 为数字，表示常驻内存的大小（单位：KB）。

unconfigurable

一种操作系统设置，可防止取消对内存的配置。

inter-board-interleave

该板与其他板交错在一起。

source=*ap_id*

表示源挂接点。

target=*ap_id*

表示目标挂接点。

deleted=#

其中 # 为数字，表示已删除的内存量（单位：KB）。

remaining=#

其中 # 为数字，表示要删除的内存量（单位：KB）。

io

io 类型显示以下信息：

device=*path*

表示 I/O 组件的物理路径。

引用的

引用了 I/O 组件。

板

board 类型显示以下布尔名称。如果这些名称不存在，则使用相反的名称。

assigned

板已分配给域。

powered-on

板已通电。

如果未指定 `-o parsable` 选项，同样的项会以更具可读性的格式显示在 `info`（信息）字段中。

`-o parsable`

以布尔 *name* 或一组由空格字符分隔的 `name=value` 对的方式返回 `info`（信息）字段中的信息。

`-o parsable` 选项可与 `-s` 选项结合使用。有关 `-s` 选项的更多信息，请参见 [cfgadm\(1M\)](#) 手册页。

`-t`

对板进行测试。

连接板之前，该板必须通过相应级别的测试。

使用此选项时，将始终尝试对板进行测试，即使该板已通过相应级别的测试也是如此。发出 `-c connect` 状态更改函数时，也会进行测试，此时，如果该板显示已完成相应级别的测试，则可跳过测试步骤。因此，可以使用 `-t` 选项来明确要求对板进行测试。

`-x function`

执行 `sbd-class` 函数。您可以使用以下函数：

`assign`

将板分配给域。

插口状态必须为 `disconnected`（断开）或 `empty`（空）。板也必须列在域可用组件列表中。请参见“动态系统域”。

`unassign`

从域中取消对板的分配。

插口状态必须为 `disconnected`（断开）或 `empty`（空）。板也必须列在域可用组件列表中。请参见“动态系统域”。

`poweron`

为系统板通电。

插口状态必须为 `disconnected`（断开）。

`poweroff`

关闭系统板的电源。

插口状态必须为 `disconnected`（断开）。

操作数

支持下列操作数：

`Receptacle ap_id`

对于 Sun Fire 15K 之类的 Sun Fire 高端系统来说，插口挂接点 ID 采用的是 SBX 或 IOX 格式，其中 X 即槽号。

具体格式取决于平台，通常与计算机上的物理标签相对应。请参见“说明”部分中特定于平台的信息。

Component *ap_id*
组件挂接点 ID 采用 *component_typeX* 格式，其中 *component_type* 即“组件类型”中描述的组件类型之一，*X* 即组件号。组件号是与板相关的设备号。

上述约定不适用于内存组件。内存挂接点上的任何 DR 操作都会影响系统板上的所有内存。

示例

以下示例显示了 Sun Fire 15K 系统上的用户输入和系统输出。用户输入（尤其是对挂接点的引用）和系统输出在诸如 Sun Fire 中型系统（如 6800）之类的其他 Sun Fire 系统上可能会有所不同。有关如何在非 Sun Fire 高端机型上使用 `cfgadm_sbd` 插件的具体信息，请参见“平台说明”。

示例1 列出所有系统板

```
# cfgadm -a -s "select=class(sbd)"

Ap_Id      Type      Receptacle  Occupant    Condition
SB0        CPU       connected   configured  ok
SB0::cpu0   cpu       connected   configured  ok
SB0::memory memory    connected   configured  ok
IO1        HPCI      connected   configured  ok
IO1::pci0   io        connected   configured  ok
IO1::pci1   io        connected   configured  ok
SB2        CPU       disconnected  unconfigured failed
SB3        CPU       disconnected  unconfigured unusable
SB4        unknown   empty        unconfigured unknown
```

本示例演示以下状况的映射：

- 插槽 2 中的板测试失败。
- 插槽 3 不可用，因此无法将板以热插拔方式插入该插槽。

示例2 列出系统板上的所有 CPU

```
# cfgadm -a -s "select=class(sbd):type(cpu)"

Ap_Id      Type      Receptacle  Occupant    Condition
SB0::cpu0   cpu       connected   configured  ok
SB0::cpu1   cpu       connected   configured  ok
SB0::cpu2   cpu       connected   configured  ok
SB0::cpu3   cpu       connected   configured  ok
```

示例3 显示 CPU 信息字段

```
# cfgadm -l -s noheadings,cols=info SB0::cpu0

cpuid 16, speed 400 MHz, ecache 8 Mbytes
```

示例4 显示可解析格式的CPU信息字段

```
# cfgadm -l -s noheadings,cols=info -o parsable SB0::cpu0

cpuid=16 speed=400 ecache=8
```

示例5 显示 I/O 板上的设备

```
# cfgadm -a -s noheadings,cols=ap_id:info -o parsable I01

I01      powered-on assigned
I01::pci0 device=/devices/saf@0/pci@0,2000 referenced
I01::pci1 device=/devices/saf@0/pci@1,2000 referenced
```

示例6 监视取消配置操作

在以下示例中，显示的内存大小以 KB 为单位。

```
# cfgadm -c unconfigure -y SB0::memory &
# cfgadm -l -s noheadings,cols=info -o parsable SB0::memory SB1::memory

address=0x0 size=2097152 permanent=752592 target=SB1::memory
      deleted=1273680 remaining=823472
address=0x1000000 size=2097152 source=SB0::memory
```

示例7 将插槽分配给域

```
# cfgadm -x assign SB2
```

示例8 从域中取消对插槽的分配

```
# cfgadm -x unassign SB3
```

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/library/platform
稳定性	请参见下文。

接口稳定性在提升。输出稳定性并不稳定。

另请参见

[cfgadm\(1M\)](#)、[devfsadm\(1M\)](#)、[ifconfig\(1M\)](#)、[mount\(1M\)](#)、[pbind\(1M\)](#)、[psradm\(1M\)](#)、[psrinfo\(1M\)](#)、

附注

本节包含有关如何监视内存删除操作进度的信息，以及特定于平台的信息。

内存删除监视

可以使用以下 shell 脚本来监视内存删除操作的进度。

```
# cfgadm -c unconfigure -y SB0::memory &
# watch_memdel SB0
```

```

#!/bin/sh
# This is the watch_memdel script.

if [ -z "$1" ]; then
    printf "usage:  %s board_id\n" `basename $0`
    exit 1
fi

board_id=$1

cfgadm_info='cfgadm -s noheadings,cols=info -o parsable'

eval `cfgadm_info $board_id::memory`

if [ -z "$remaining" ]; then
    echo no memory delete in progress involving $board_id
    exit 0
fi

echo deleting target $target

while true
do
    eval `cfgadm_info $board_id::memory`

    if [ -n "$remaining" -a "$remaining" -ne 0 ]
    then
        echo $deleted KBytes deleted, $remaining KBytes remaining
        remaining=
    else
        echo memory delete is done
        exit 0
    fi
    sleep 1
done
exit 0

```

Sun Enterprise 10000
平台说明

以下语法用于引用 Sun Enterprise 10000 系统上的挂接点：

board::component

...其中 *board* 是指系统板，*component* 是指单个组件。系统板的编号范围是 SB0（零）到 SB15。最多可提供十六个系统板。

在 Sun Enterprise 10000 域上运行的 DR 3.0 型号只支持 `cfgadm_sbd` 插件所提供功能中很有限的一部分。唯一支持的操作是查看域中挂接点的状态。这对应于 `-l` 选项及其所有关联的选项。

尝试在域中执行其他任何操作都会导致错误，指出操作不受支持。所有添加或移除系统板的操作都必须从系统服务处理器启动。

Sun Fire 高端系统平台说明

以下语法用于引用 Sun Fire 高端系统上的挂接点：

board::component

其中 *board* 是指系统板或 I/O 板，*component* 是指单个组件。

系统板的编号范围从 SB0（零）到 SB17，I/O 板的编号范围从 IO0（IO 零）到 IO17，具体取决于系统的配置。（最多可提供十八个系统板和 I/O 板。）

-t 和 -x 选项在 Sun Fire 高端系统平台上的行为并不相同。以下列表对其行为进行了描述：

-t

系统控制器可通过运行 LPOST，按 hpost 命令排序使用 CPU 来测试系统板。测试 I/O 板时，驱动程序会按 -t 选项的要求开始测试，且测试不需用户干预即可自动运行。驱动程序将取消对某个 CPU 及某段连续物理内存的配置，然后向系统控制器发送进行板测试的命令。系统控制器将使用 CPU 和内存从事务/错误箱内部对 I/O 板进行测试。只能使用系统板（而非 MCPU 板）上的 CPU 来测试 I/O 板。

-x assign | unassign

在 Sun Fire 高端系统管理模型中，平台管理员可以使用每个域的可用组件列表来控制平台硬件。此信息保留在系统控制器上。只有平台管理员可以修改域的可用组件列表。

仅当该域的可用组件列表中包含某个板的情况下，域管理员才能对该板执行分配或取消分配操作。平台管理员没有这种限制，因此即使在域的可用组件列表中不含某个板，也可以对该板执行分配或取消分配操作。

Sun Fire 15K 组件类型 以下为组件类型的名称和描述：

cpu

CPU

io

I/O 设备

memory（内存）

内存

注意：针对某个内存组件的操作会影响板上的所有内存组件。

Sun Fire 中型系统平台说明

Sun Fire 中端服务器（如 6800、4810、4800 及 3800 系统）上的挂接点引用与 Sun Fire 高端系统上的挂接点引用略有不同。以下语法用于引用除 Sun Fire 15K 之外的 Sun Fire 系统上的挂接点：

N#.board::component

其中 *N#* 是指节点，*board* 是指系统板或 I/O 板，*component* 是指单个组件。

系统板的编号范围从 SB0 到 SB5，I/O 板的编号范围从 IB6 到 IB9，具体取决于系统的配置。（最多可提供六个系统板和四个 I/O 板。）

Sun Fire 中型系统组
组件类型

以下为组件类型的名称和描述：

cpu

CPU

pci

I/O 设备

memory (内存)

内存

注意：针对某个内存组件的操作会影响板上的所有内存组件。

引用名

cfgadm_scsi – 特定于 SCSI 硬件的 cfgadm 命令

用法概要

```
/usr/sbin/cfgadm [-f] [-y | -n] [-v] [-o hardware_option]
                  -c function ap_id...

/usr/sbin/cfgadm [-f] [-y | -n] [-v] [-o hardware_option]
                  -x hardware_function ap_id...

/usr/sbin/cfgadm [-v] [-a] [-s listing_option] [-o hardware_option]
                  [-l [ap_id | ap_type ... ]]

/usr/sbin/cfgadm [-v] [-o hardware_option] -t ap_id...

/usr/sbin/cfgadm [-v] [-o hardware_option] -h [ap_id]...
```

描述

特定于 SCSI 硬件的库 `/usr/lib/cfgadm/scsi.so.1` 通过 [cfgadm\(1M\)](#) 命令提供 SCSI 热插拔功能。cfgadm 在挂接点上运行，这些挂接点是可动态重新配置硬件资源的系统位置。有关挂接点的信息，请参见 [cfgadm\(1M\)](#)。

对于 SCSI 热插拔，每个 SCSI 控制器由设备树中的挂接点表示。此外，每个 SCSI 设备由动态挂接点表示。挂接点通过 `ap_id` 命名。定义下列两种类型的 `ap_id`：逻辑和物理。物理 `ap_id` 基于物理路径名称，而逻辑 `ap_id` 的名称较短且更加用户友好。对于 SCSI 控制器，逻辑 `ap_id` 通常是对应的磁盘控制器号。例如，一个典型的逻辑 `ap_id` 为 `c0`。

SCSI 设备的命名与控制器的 `ap_id` 相关。因此，如果磁盘设备附加到控制器 `c0`，其 `ap_id` 可以是：

```
c0::dsk/c0t0d0
```

其中，`dsk/c0t0d0` 标识特定的设备。通常，设备标识符源自 `/dev` 中的设备的对应逻辑链接。例如，SCSI 磁带机的逻辑 `ap_id` 可能为 `c0::rmt/0`。其中，`c0` 是 SCSI 控制器的逻辑 `ap_id`，`rmt/0` 源自 `/dev/rmt` 中的磁带机的逻辑链接。如果标识符无法从 `/dev` 中的链接获取，将为其分配唯一的标识符。例如，如果 `/dev` 中不存在磁带设备链接，可分配 `c0::st3` 格式的 `ap_id`，其中，`st3` 是内部生成的唯一标识符。

如果启用了支持 Solaris I/O 多路径功能（以前称为 "MPxIO" 或 "Sun StorEdge Traffic Manager [STMS]"）的控制器，该控制器可提供到 [scsi_vhci\(7D\)](#) 多路径设备的路径。如果 [scsi_vhci\(7D\)](#) 支持附加到此类控制器的设备，该控制器的 `ap_id` 可为：

```
c0::0,0
```

...其中，`0,0` 唯一标识目标和逻辑单元信息。这些 `ap_id` 的路径的 `Type`（类型）字段表明其是否表示到 [scsi_vhci\(7D\)](#) 多路径设备的路径，以及通过该路径连接到的设备的类型。

系统中挂接点的简单列表包含 SCSI 控制器的挂接点，不包含 SCSI 设备的挂接点。在 `(-l)` 列表选项中使用 `-a` 标志也可列出 SCSI 设备。例如：

```
# cfgadm -l
Ap_Id          Type          Receptacle    Occupant      Condition
c0              scsi-bus      connected     configured    unknown
```

```
sysctrl0:slot0    cpu/mem    connected   configured  ok
sysctrl0:slot1    sbus-upa   connected   configured  ok
```

除了 SCSI 控制器之外，还列出 SCSI 设备：

```
# cfgadm -al

Ap_Id              Type           Receptacle    Occupant      Condition
c0                 scsi-bus       connected     configured    unknown
c0::dsk/c0t14d0    disk           connected     configured    unknown
c0::dsk/c0t11d0    disk           connected     configured    unknown
c0::dsk/c0t8d0     disk           connected     configured    unknown
c0::dsk/c0t0d0     disk           connected     configured    unknown
c0::rmt/0          tape           connected     configured    unknown
sysctrl0:slot0     cpu/mem        connected     configured    ok
sysctrl0:slot1     sbus-upa       connected     configured    ok
```

如果控制器 `c0` 启用了 Solaris I/O 多路径功能并且 Solaris I/O 多路径功能支持所连接的磁盘和磁带设备，将输出：

```
# cfgadm -al

Ap_Id              Type           Receptacle    Occupant      Condition
c0                 scsi-bus       connected     configured    unknown
c0::11,0           disk-path      connected     configured    unknown
c0::14,0           disk-path      connected     configured    unknown
c0::8,0            disk-path      connected     configured    unknown
c0::0,0            disk-path      connected     configured    unknown
c0::a.0            tape-path      connected     configured    unknown
sysctrl0:slot0     cpu/mem        connected     configured    ok
sysctrl0:slot1     sbus-upa       connected     configured    ok
```

有关列出挂接点的更多信息，请参见 [cfgadm\(1M\)](#)。SCSI 控制器上的挂接点的插口状态和插卡状态的含义如下：

- `empty` (空)
不适用
- `disconnected` (断开)
总线停止 (总线上的 I/O 活动暂停)
- `connected`
总线处于活动状态
- `configured` (已配置)
总线上配置了一个或多个设备
- `unconfigured` (未配置)
总线上未配置任何设备

单个 SCSI 设备的对应状态为：

empty (空)

不适用

disconnected (断开)

设备附加到的总线处于停止状态

connected

设备附加到的总线处于活动状态

configured (已配置)

配置了设备或多路径 SCSI 设备的路径

unconfigured (未配置)

未配置设备或多路径 SCSI 设备的路径

选项

除了列出 (-l)，`cfgadm` 还定义了其他几种类型的操作。这些操作包括测试 (-t)、调用配置状态更改 (-c)、调用特定于硬件的函数 (-x) 以及获取配置管理帮助消息 (-h)。

-c *function*

以下通用命令是针对特定于 SCSI 硬件的库定义的：

对于 SCSI 控制器挂接点，支持以下配置状态更改操作：

connect

Unquiesce the SCSI bus.

disconnect

总线停止（暂停总线上的 I/O 活动）。

错误使用该命令可导致系统挂起。请参见“说明”部分。

configure

在 SCSI 总线上配置新设备。

unconfigure

取消对连接到总线的所有设备的配置。

以下通用命令是针对 SCSI 设备以及多路径 SCSI 设备的路径定义的：

configure

配置特定设备或多路径 SCSI 设备的特定路径。

unconfigure

取消对特定设备或多路径 SCSI 设备的特定路径的配置。

-f

如果该选项用于 `disconnect` 命令，则强制停止 SCSI 总线（如果硬件支持）。

错误使用该命令可导致系统挂起。请参见“说明”部分。

-h *ap_id*

SCSI 特定帮助可通过将帮助选项与任何 SCSI 挂接点一起使用来获取。

-o *hardware_option*

当前没有定义任何硬件专用选项。

-s *listing_option*

可使用 **select** 子选项列出 **scsi** 类的挂接点。有关其他信息，请参见 **cfgadm(1M)** 手册页。

-t *ap_id*

目前没有可用的测试命令。

-x *hardware_function*

下面的部分命令只适用于 SCSI 控制器，部分命令只适用于 SCSI 设备。

在下面，**controller_ap_id** 指 SCSI 控制器的 **ap_id**，例如，**c0**。**device_ap_id** 指 SCSI 设备的 **ap_id**，例如，**c0::dsk/c0dt3d0**。

定义了以下特定于硬件的函数：

insert_device *controller_ap_id*

向 SCSI 控制器 **controller_ap_id** 添加新设备。

该命令只用于交互式使用。

remove_device *device_ap_id*

删除 **device_ap_id** 设备。

该命令只用于交互式使用。

replace_device *device_ap_id*

删除 **device_ap_id** 设备并使用同类设备代替。

该命令只用于交互式使用。

reset_device *device_ap_id*

复位 **device_ap_id**。

reset_bus *controller_ap_id*

仅复位 **controller_ap_id** 的总线，而不复位附加到该总线的任何设备。

reset_all *controller_ap_id*

复位 **controller_ap_id** 的总线以及该总线上的所有设备。

locator [=on|off] *device_ap_id*

如果平台提供硬盘定位器 LED 指示灯，设置或获取该指示灯。如果未设置 [on|off] 子选项，则输出硬盘定位器的状态。

led[=LED,mode=on|off|blink] *device_ap_id*

如果未设置任何子参数，则该函数输出当前 LED 指示灯设置的列表。如果设置了子参数，则该函数设置插槽特定的 LED 指示灯的模式。

示例

示例1 配置磁盘

以下命令配置附加到 c0 控制器的磁盘：

```
# cfgadm -c configure c0::dsk/c0t3d0
```

示例2 取消配置磁盘

以下命令取消对附加到 c0 控制器的磁盘的配置。

```
# cfgadm -c unconfigure c0::dsk/c0t3d0
```

示例3 添加新设备

以下命令向 c0 控制器添加新设备：

```
# cfgadm -x insert_device c0
```

系统作出如下响应：

```
Adding device to SCSI HBA: /devices/sbus@1f,0/SUNW,fas@e,88000000
This operation will suspend activity on SCSI bus c0
Continue (yes/no)?
```

输入：

```
y
```

系统作出如下响应：

```
SCSI bus quiesced successfully.
It is now safe to proceed with hotplug operation.
Enter y if operation is complete or n to abort (yes/no)?
```

输入：

```
y
```

示例4 替换设备

以下命令替换附加到 c0 控制器的设备：

```
# cfgadm -x replace_device c0::dsk/c0t3d0
```

系统作出如下响应：

```
Replacing SCSI device: /devices/sbus@1f,0/SUNW,fas@e,88000000/sd@3,0
This operation will suspend activity on SCSI bus: c0
Continue (yes/no)?
```

输入：

```
y
```

系统作出如下响应：

示例4 替换设备 (续)

SCSI bus quiesced successfully.
 It is now safe to proceed with hotplug operation.
 Enter y if operation is complete or n to abort (yes/no)?

输入:

y

示例5 取消配置磁盘时遇到已挂载的文件系统

以下命令表明在取消配置磁盘时遇到已挂载的文件系统:

```
# cfgadm -c unconfigure c1::dsk/clt0d0
```

系统作出如下响应:

```
cfgadm: Component system is busy, try again: failed to offline:
/devices/pci@1f,4000/scsi@3,1/sd@1,0
      Resource                      Information
-----
/dev/dsk/clt0d0s0  mounted filesystem "/mnt"
```

示例6 显示磁盘定位器的值

以下命令显示磁盘定位器的值。该示例特定于 SPARC Enterprise Server 系列:

```
# cfgadm -x locator c0::dsk/c0t6d0
```

系统作出如下响应:

```
Disk                      Led
c0t6d0                    locator=on
```

示例7 设置磁盘定位器的值

以下命令设置磁盘定位器的值。该示例特定于 SPARC Enterprise Server 系列:

```
# cfgadm -x locator=off c0::dsk/c0t6d0
```

系统不输出任何信息作为响应。

示例8 配置多路径 SCSI 磁盘的路径

以下命令配置通过 c0 控制器连接的路径:

```
# cfgadm -c configure c0::2,0
```

示例9 取消对多路径 SCSI 磁盘的路径的配置

以下命令取消对通过 c0 控制器连接的路径的配置:

```
# cfgadm -c unconfigure c0::2,0
```

文件 `/usr/lib/cfgadm/scsi.so.1`
 适用于通用 SCSI 热插拔的特定于硬件的库

`/usr/platform/SPARC-Enterprise/lib/cfgadm/scsi.so.1`
 适用于通用 SCSI 热插拔的特定于平台的库

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/library

另请参见 [cfgadm\(1M\)](#)、[luxadm\(1M\)](#)、[config_admin\(3CFGADM\)](#)、[libcfgadm\(3LIB\)](#)、[attributes\(5\)](#)、[scsi_vhci\(1M\)](#)

附注 对于控制包含关键分区（如 `root (/)`、`/usr`、`swap` 或 `/var`）的磁盘的控制器，不支持 `disconnect`（停止）操作。不应在此类控制器上尝试执行 `disconnect` 操作。不正确的用法可能导致系统挂起且需要重新引导。

 如果控制器处于断开（停止）状态，则可能导致系统死锁。使用 `disconnect` 操作时需谨慎。在完成 DR 操作所需的最短时间内，控制器应保持断开状态。`disconnect` 命令只用于在系统运行时替换 SCSI 电缆。不应出于其他任何目的使用该命令。解决死锁（如果发生）的唯一方法是重新引导系统。

 并非所有 SCSI 控制器都支持热插拔操作。

警告 某些 SCSI 设备上的连接器不遵循 SCSI 热插拔规范。在此类设备上执行热插拔操作可能损坏 SCSI 总线上的硬件。有关其他信息，请参见硬件手册。

引用名

cfgadm_sdcard – 特定于 SD/MMC 硬件的 cfgadm 命令

用法概要

```
/usr/sbin/cfgadm [-f] [-y | -n] [-o hardware_options]
-c function ap_id[...]

/usr/sbin/cfgadm [-f] [-y | -n] [-o hardware_options]
-c hardware_function ap_id[...]

/usr/sbin/cfgadm [-v] [-a] [-s listing_options]
[-o hardware_options] [-l [ap_id | ap_type[...]]]

/usr/sbin/cfgadm [-v] [-o hardware_options] -h [ap_id[...]]
```

描述

特定于安全数字 (Secure Digital, SD) 和多媒体存储卡 (MultiMediaCard, MMC) 硬件的库 /usr/lib/cfgadm/sdcard.so.1 通过 [cfgadm\(1M\)](#) 命令提供 SD/MMC 热插拔功能。cfgadm 命令对挂接点执行操作，而挂接点是系统中可以动态重新配置硬件资源的位置。有关挂接点的信息，请参见 [cfgadm\(1M\)](#)。

每个 SD/MMC 插槽都通过设备树中的一个挂接点来表示。系统中连接并配置的卡设备显示为挂接点名称的扩展。“挂接点”和“SD/MMC 插槽”两个术语在本手册页中可互换使用。

挂接点通过 ap_id 命名。所有的 SD/MMC 挂接点均由以下形式的字符串组成：

```
sdcardX/[S][::dsk/cXtYd0]
```

其中：

X 是 SD/MMC 控制器号。

S 是控制器上的槽号（0 到 8）。

dev/cXtYd0 标识插入的内存卡。

Y 是目标编号。

通常，设备标识符源自 /dev 中的设备的对应逻辑链接。由于 SD/MMC 设备只支持一个 LUN (LUN 0)，因此，设备字符串的 "d" 组件总是设置为 0（零）。例如，SD/MMC 控制器 2 的插槽 4 的逻辑 ap_id 应为 sdcard2/4。如果 SD/MMC 媒体卡插入该挂接点并且该设备已配置，ap_id 可能为 sdcard2/4::dsk/c2t0d0。

标识设备的 cXtYd0 字符串与设备挂接点之间存在一对一的对应关系。

系统中挂接点的简单列表包含所有 SD/MMC 设备插槽以及附加的设备。例如：

```
#cfgadm -l
```

Ap_Id	Type	Receptacle	Occupant	Condition
sdcard0/0::dev/c2t0d0	sdcard	connected	configured	ok
sata0/1::dev/c0t1d0	disk	connected	configured	ok
sata0/2::dev/c0t2d0	cd-dvd	connected	configured	ok
sata0/3	sata-port	empty	unconfigured	ok

usb0/1	unknown	empty	unconfigured	ok
usb0/2	unknown	empty	unconfigured	ok

有关挂接点列表的更多信息，请参见 [cfgadm\(1M\)](#)。

SD/MMC 插槽中的挂接点的插口状态的含义如下：

empty（空）
SD/MMC 插槽通电且已启用。在该插槽未检测到任何设备。

disconnected（断开）
SD/MMC 插槽未启用，或检测到 SD/MMC 设备但未与该设备建立通信，或插槽出现故障。

connected（已连接）
在插槽中检测到 SD/MMC 设备并且已建立设备通信。

插卡（SD/MMC 插槽中插入的设备）状态的含义如下：

configured（已配置）
附加的 SD/MMC 设备已配置，且可供操作系统使用。

unconfigured（未配置）
未附加任何设备或 SD/MMC 插槽中插入的 SD/MMC 设备尚未配置。要进行配置，请运行 **cfgadm -c configure ap_id** 命令。

挂接点（SD/MMC 插槽）状况的含义如下：

ok（正常）
SD/MMC 插槽已通电并启用，随时可用。

failed（失败）
SD/MMC 插槽出现故障。该端口被系统禁用和/或断电。这种情况下，端口不可用且处于未知状态。该故障可能是由插入插槽的设备导致的。

unknown（未知）
SD/MMC 插槽已禁用且处于未知状况。

选项 除了列出 (-l)，**cfgadm** 命令还定义了其他几种类型的操作。这些操作包括调用配置状态更改 (-c)、调用特定于硬件的函数 (-x) 以及获取配置管理帮助消息 (-h)。

-c：对于 SD/MMC 插槽挂接点，支持以下配置状态更改操作：

connect（连接）
启用（激活）SD/MMC 插槽并建立与所附加设备之间的通信。该操作意味着给插槽通电（如有必要）。

disconnect（断开）
如果尚未取消插入的设备的配置，则取消其配置并禁用（取消激活）SD/MMC 插槽。后续 "connect" 命令可启用 SD/MMC 插槽操作但不会使设备处于 "configured"（已配置）状态。

以下状态更改操作适用于 SD/MMC 插槽中插入的 SD/MMC 卡：

configure（配置）

配置新设备以便由操作系统使用（如果尚未配置）。该命令意味着执行连接操作（如有必要）。

unconfigure（取消配置）

取消插入 SD/MMC 插槽的设备的配置（如果尚未取消配置）。

-f：不支持。

-h *ap_id*：SD/MMC 特定帮助可通过将帮助选项与任何 SD/MMC 挂接点一起使用来获取。

-l [-v]：按照 [cfgadm\(1M\)](#) 中所述使用 **-l** 选项。如果该选项与 **-v** 选项一起使用，"Information"（信息）字段将包含以下特定于 SD/MMC 的信息：

Mod（型号）：产品型号字符串

Rev（修订版）：产品修订版本号（主发行版.次发行版）

Date（日期）：制造日期（年、月）

SN：产品序列号（十六进制）

-o hardware_options：当前没有定义特定于硬件的选项。

-s listing_options：可使用 **select** 子选项列出 SD/MMC 类的挂接点。请参见 [cfgadm\(1M\)](#)。

-t *ap_id*：自检功能。SD/MMC 插槽不支持。

-x hardware_function：执行特定于硬件的函数。**sdcard_reset_slot *ap_id*** 指示复位 *ap_id* 指定的 SD/MMC 插槽。如果插槽中插入了 SD/MMC 设备，也会复位该设备。

-v：在详细模式下执行。

示例

示例 1—配置 SD/MMC 卡：

以下命令配置附加到 SD/MMC 控制器 0 上的插槽 0 的卡。仅当 SD/MMC 插槽中插入设备时，才应发出该命令。

```
# cfgadm -c configure sdcard0/0
```

示例 2—取消 SD/MMC 卡的配置：

以下命令取消 SD/MMC 控制器 0 上的插槽 3 中插入的卡的配置：

```
# cfgadm -c unconfigure sdcard0/3
```

示例 3—取消磁盘配置时遇到一个已挂载文件系统：

以下命令表明在取消配置磁盘时遇到已挂载的文件系统：

```
# cfgadm -c unconfigure sdcards1/5::dsk/c01t35d0
```

系统作出如下响应：

```
cfgadm: Component system is busy, try again: failed to offline:
/devices/pci@0,0/pci8086,244e@1e/pci1095,3124@1/sd@5,0
Resource                      Information
-----
/dev/dsk/c1t5d0s0    mounted filesystem "/mnt"
```

文件

/usr/lib/cfgadm/sdcards.so.1
适用于通用 SD/MMC 热插拔的特定于硬件的库。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/library

另请参见

[cfgadm\(1M\)](#)、[config_admin\(3CFGADM\)](#)、[libcfgadm\(3LIB\)](#)、[attributes\(5\)](#)、[sda\(7D\)](#)、[sdcards\(7D\)](#)、

附注

正常操作情况下，SD/MMC 媒体卡会在插入卡时自动配置。大多数管理员发现正常操作时无需使用该命令。

如果正在向 SD/MMC 卡写入内容，那么删除该设备前应先取消其配置，否则，将导致数据丢失。以只读模式挂载的设备可随时安全地删除。

配置了文件系统（[pcfs\(7FS\)](#) 除外）的设备始终应该在删除前先显式取消配置。

引用名 `cfgadm_shp` – 特定于 PCI Express 和标准 PCI 热插拔硬件的 `cfgadm` 命令

用法概要

```

/usr/sbin/cfgadm [-f] [-y | -n] [-v]
                [-o hardware_options] -c function ap_id [ap_id]

/usr/sbin/cfgadm [-f] [-y | -n] [-v]
                [-o hardware_options] -x hardware_function ap_id [ap_id]

/usr/sbin/cfgadm [-v] [-s listing_options]
                [-o hardware_options] -x hardware_function ap_id [ap_id]

/usr/sbin/cfgadm [-v] [-o hardware_options] -tap_id [ap_id]

/usr/sbin/cfgadm [-v] [-o hardware_function] -h [ap_id | ap_type]
```

描述

将 PCI Express 和标准 PCI 热插拔适配卡热插到允许热插拔系统中的相应可热插拔插槽中时，特定于 PCI Express 和标准 PCI 热插拔硬件的库 `/usr/lib/cfgadm/shp.so.1` 可通过 `cfgadm` 命令（请参见 [cfgadm\(1M\)](#)）提供相关支持。`cfgadm_pci` 库（请参见 [cfgadm_pci\(1M\)](#)）提供对其他 PCI 热插拔适配卡（除 PCI Express 和标准 PCI 热插拔卡之外的其他卡）的支持。除了本手册页中标注的以外，介于 PCI Express 热插拔和标准 PCI 热插拔之间的热插拔管理模型仍是相同的。

对于 PCI 热插拔，特定 PCI 总线上的每个热插拔槽均可由该 PCI 总线的一个挂接点来表示。

一个挂接点由一个插口和一个插卡组成。PCI 热插拔下的插口通常指物理可热插拔槽；插卡通常指插入插槽的 PCI 适配卡。

挂接点通过 `ap_id` 命名。`ap_id` 包括逻辑 `ap_id` 和物理 `ap_id`。物理 `ap_id` 基于物理路径名，例如：

```
/devices/pci@7c,0/pci10de,5d@:pcie2
```

而逻辑 `ap_id` 是更为简短、更加用户友好的名称，例如 `pcie2`。热插拔 PCI 的 `ap_type` 为 `pci`。

请注意，`ap_type` 与 `Type`（类型）字段中的信息不同。

PCI Express ap_id 命名 对于位于 PCI Express 分层结构（也就是说，父或祖先为 PCI Express 设备）中的挂接点（包括本身不属于 PCI Express 设备的挂接点），将使用下面显示的命名模式。

语法：

APID: *absolute-slot-path*
基本术语。

absolute-slot-path: *slot-path*[:*slot-path*[:*slotpath* ...]]
...其中 *fru-id* 表示包含 *slot-id* 的机箱 FRU（如果有）。

fru-id: *fru-type*[*serialid*#]
...其中 *fru-type* 为 PCI Express 扩展机箱的 "iob"，后接序列号 *serialid*#（如果提供）

```
slot-id : slot-name | device-type physical-slot# | \
nexus-driver-name nexus-driver-instance.\
device-type pci-device-number
```

...其中 *slot-name* 为平台或硬件自行分配的名称。*device-type* 为 *pcie*（对于 PCI Express 设备）或 *pci*（对于 PCI 设备）。*nexus-driver-name* 为设备组件的驱动程序名称，*physical-slot#* 为硬件槽号，而 *pci-device-number* 为按标准 PCI 命名规则命名的 PCI 设备编号。

首先将构造 *absolute-slot-path*，以便使用在用户看来更易于实际确认的术语来描述挂接点的拓扑位置。此 *absolute-slot-path* 包含由 :（冒号）分隔的 *slot-path* 组件。叶或最左侧的 *slot-path* 组件描述的是挂接点设备本身，而其右侧相邻的 *slot-path* 组件直至最右侧或最顶端的 *slot-path* 组件分别描述从父设备到根设备的各种设备。

每个 *slot-path* 均含 *slot-id*（可选择性地在前面添加 *fru-id*，该 ID 用于标识内含 *slot-id*（详见下）所描述设备的扩展机箱）。*fru-id* 包含 *fru-type*，后跟可选的 *serialid#*。PCI Express 扩展机箱类型的 *fru-type* 为 "iob"，而 *serialid#* 既可以从扩展机箱硬件中获取的 64 位十六进制数（表示原始序列号），也可以是大写的适用于 Sun 品牌扩展机箱的 ASCII 四字符序列。

每个 *slot-id* 均包含下述三个可能的格式之一：

slot-id 格式 (1)

slot-names

slot-id 格式 (2)

device-type physical-slot#

slot-id 格式 (3)

nexus-driver-name nexus-driver-instance device-type pci-device-number

在选择具体格式时，将按照格式编号从最低到最高的优先顺序进行选择，也可以自上而下地进行选择（如上所述）。如果无法成功构造某个格式，则会尝试下一个编号较高的格式。

slot-id 格式 (1) 中的 *slot-names* 是从设备树中对应节点的 *slot-names* 属性获取的，是由硬件或平台分配的名称。此格式尚未预定义或建立。

在 *slot-id* 格式 (2) 中，*device-type* 表示组件插槽的设备类型，其值为 *pcie*（对于 PCI Express）或 *pci*（对于 PCI），而从对应设备节点的 *physical-slot#* 属性获取的 *physical-slot#* 表示组件的硬件槽号。

当其他所有格式均无法成功构造时，将使用 *slot-id* 格式 (3)，并被认为是缺省格式。*nexus-driver-name* 是组件的驱动程序名称；*nexus-driver-instance* 是该驱动程序的实例；*device-type* 与格式 (2) 中描述的相同；*pci-device-number* 是所述的 PCI 设备编号，用于按照标准 PCI 命名规则命名的设备配置循环。

在对 *slot-path* 组件进行总结时，我们会发现该组件是对可选 FRU 组件（可能位于该组件之前）的扩展，具体说来，*slot-path* 将包含以下按顺序排列的格式之一：

- (1) [iob[serialid#].]
 slot-names
- (2) [iob[serialid#].]
 device_type physical_slot#
- (2) [iob[serialid#].]
 nexus-driver-name nexus-driver-instance.
- device_type pci-device-number*

最后，在 cfadm 中实际使用的 ap_id 名称的最终格式确定如下（按优先级顺序指定）：

ap_id 格式 (1)

如果 *absolute-slot-path* 的长度在 cfadm 的 ap_id 字段的固定长度限制之内，则只需使用 *absolute-slot-path*。

ap_id 格式 (2)

（*absolute-slot-path* 的长度超过 ap_id 的长度限制）如果最后一个 *slot_path* 组件包含在扩展机箱中，且该组件包含 *serialid#*，则使用这最后一个 *slot_path* 组件。此格式要求使用 *serialid#* 是为了确保 ap_id 在全局范围内都是唯一的。

ap_id 格式 (3)

（*absolute-slot-path* 的长度超过 ap_id 的长度限制）将使用最后一个 *slot_path* 组件的缺省格式，即 *slot-id* 格式 (3)。

不管使用哪一个最终的 ap_id 名称，*absolute-slot-path* 都将存储在能够使用 -s 或 -v 选项显示的 "Information"（信息）(info) 字段中。此信息可用于对任何使用 ap_id 格式 (2) 或 ap_id 格式 (3) 命名的 ap_id 进行物理定位。将 *absolute-slot-path* 存储在信息字段中时会进行稍许变换，即使用正斜杠 (/) 替换冒号 (:)，以便更准确地指示拓扑上下文。*absolute-slot-path* 可能包括叶或最右侧 *slot-path* 组件（直至板载主机槽）上面不属于热插拔组件的 *slot-path* 组件。

有关可热插拔组件的示例列表，请参见“示例”部分。

选项

所支持的选项如下：

-c *function*

以下函数可用于 PCI 可热插拔插槽：

configure (配置)

将 PCI 设备配置在 Solaris 使用的插槽中。

connect

将插槽与 PCI 总线相连。

disconnect

断开插槽与 PCI 总线之间的连接。

insert
不支持。

remove
不支持。

unconfigure
将 PCI 设备的资源从系统中逻辑移除。

-f
不支持。

-h *ap_id* | *ap_type*
显示特定于 PCI 热插拔的帮助消息。

-l *list*
列出 PCI 热插拔插槽的值。

-o *hardware_options*
当前没有定义任何硬件专用选项。

-s *listing_options*
与通用的 [cfgadm\(1M\)](#) 相同。

-t *ap_id*
此命令只适用于支持插槽测试功能的平台。

-v
在详细模式下执行。

如果 **-v** 选项与 **-l** 选项一起使用，**cfgadm** 命令输出有关挂接点的信息。对 PCI Express 分层结构中的挂接点，“Information”（信息）字段将包含挂接点的绝对插槽路径位置，包括插槽路径中每个组件的任何特定于硬件或平台的标记信息。插槽路径中的每个组件将通过 /（正斜杠）分隔。请参见上面的“PCI Express *ap_id* 命名”。有关不在 PCI Express 分层结构中的 PCI 热插拔挂接点的内容，请参见 [cfgadm_pci\(1M\)](#)。Type（类型）字段中的信息可使用 **-v** 选项列出，也可不使用该选项列出。插卡的 Type（类型）字段描述有关插槽的信息。有两个可能的值：

unknown（未知）

该插槽为空。如果该插槽中安装了卡，表明该卡未配置或卡上的设备未安装驱动程序。

subclass/board

该插槽中的卡是一个单功能或多功能设备。

subclass 是表示设备子类代码的字符串，例如 SCSI、ethernet、pci-isa 等。如果该卡为多功能设备，则将改为显示 **MULT**。

board 是表示设备的板类型的字符串。例如，hp 是用于 PCI 热插拔适配器的字符串。

-x *hardware_function*

执行特定于硬件的函数。这些特定于硬件的函数通常不会更改插口或插卡的状态。

支持以下 *hardware_function* :

`led=[led_sub_arg],mode=[mode_sub_arg]`

在没有子参数的情况下，将显示当前 LED 指示灯设置的列表。在有子参数的情况下，将为插槽设置特定 LED 指示灯的模式。

将 *led_sub_arg* 指定为 `fault`、`power`、`attn` 或 `active`。

将 *mode_sub_arg* 指定为 `on`、`off` 或 `blink`。

对于 PCI Express，只有 `power` 和 `attn` LED 指示灯是有效的，只有 `attn` LED 指示灯的状态可以更改。

更改 LED 指示灯的状态不会更改插口或插卡的状态。通常，LED 指示灯由热插拔控制器控制，无需用户干预。使用该命令用于测试目的。

注意 - 更改 LED 指示灯的状态可能导致插卡或插口的状态显示错误。

以下命令可显示 LED 指示灯的值：

```
example# cfgadm -x led pcie2
Ap_Id      Led
pcie2      power=on,fault=off,active=off,attn=off
```

以下命令将 `attn` LED 指示灯设置为 `blink`，以指示插槽的位置：

```
example# cfgadm -x led=attn,mode=blink pcie2
```

示例

示例1 显示每个插槽的值

以下命令可显示每个插槽的值：

```
example# cfgadm -l
```

Ap_Id	Type	Receptacle	Occupant	Condition
c0	scsi-bus	connected	configured	unknown
c1	scsi-bus	connected	unconfigured	unknown
c2	scsi-bus	connected	unconfigured	unknown
pcie7	etherne/hp	connected	configured	ok
pcie8	unknown	empty	unconfigured	unknown
pcie9	fibre/hp	connected	configured	ok

示例2 替换卡

以下命令列出所有支持 DR 的挂接点：

```
example# cfgadm
```

Type	Receptacle	Occupant	Condition
c0	scsi-bus	connected	configured
c1	scsi-bus	connected	unconfigured

示例2 替换卡 (续)

c2	scsi-bus	connected	unconfigured	unknown
pcie7	etherne/hp	connected	configured	ok
pcie8	unknown	empty	unconfigured	unknown
pcie9	fibre/hp	connected	configured	ok

以下命令可取消对 pcie7 所标识的卡的配置并以电子方式断开其连接：

```
example# cfgadm -c disconnect pcie7
```

该更改可通过输入以下命令验证：

```
example# cfgadm pcie7
Ap_Id      Type      Receptacle  Occupant    Condition
pcie7      unknown   disconnected unconfigured unknown
```

此时可换卡。以下命令可通过电子方式连接并配置替换卡：

```
example# cfgadm -c configure pcie7
```

该更改可通过输入以下命令验证：

```
example# cfgadm pcie7
Ap_Id      Type      Receptacle  Occupant    Condition
pcie7      etherne/hp connected   configured  ok
```

示例3 解释 PCI Express 拓扑中的 ApId

以下命令显示了 I/O 扩展机箱在连接到主机级可热插拔插槽后，其中的 PCI Express 和 PCI 挂接点的拓扑列表：

```
example# cfgadm -s cols=ap_id:info
Ap_Id      Information
iou#0-pci#0      Location: iou#0-pci#0
iou#0-pci#1      Location: iou#0-pci#1
iou#0-pci#1:iob.pci3      Location: iou#0-pci#1/iob.pci3
iou#0-pci#1:iob.pci4      Location: iou#0-pci#1/iob.pci4
iou#0-pci#2      Location: iou#0-pci#2
iou#0-pci#2:iob58071.pcie1      Location: iou#0-pci#2/iob58071.pcie1
iou#0-pci#2:iob58071.special      Location: iou#0-pci#2/iob58071.special
iou#0-pci#3      Location: iou#0-pci#3
iou#0-pci#3:iobBADF.pcie1      Location: iou#0-pci#3/iobBADF.pcie1
iou#0-pci#3:iobBADF.pcie2      Location: iou#0-pci#3/iobBADF.pcie2
iou#0-pci#3:iobBADF.pcie3      Location: iou#0-pci#3/iobBADF.pcie3
iou#0-pci#3:iobBADF.pci1      Location: iou#0-pci#3/iobBADF.pci1
iou#0-pci#3:iobBADF.pci2      Location: iou#0-pci#3/iobBADF.pci2
```

在本示例中，iou#0-pci#[0-3] 项表示系统中最顶端的可热插拔插槽。由于 iou#n-pci#n 格式与上述语法规范部分中描述的任何一种格式都不匹配，因此我们可以推断，此热插拔拓扑中基本组件的这种名称是通过 slot-names 属性由平台指定的。

示例 3 解释 PCI Express 拓扑中的 ApId (续)

上述输出中的插槽描述如下：

插槽 iou#0-pci#0
此插槽为空或其插卡未配置。

插槽 iou#0-pci#1
此插槽包含一个带两个可热插拔插槽 (pci3 和 pci4) 的扩展机箱。pci3 和 pci4 表示两个包含在该扩展机箱内的 PCI 插槽，其物理槽号分别为 3 和 4。此处的扩展机箱既没有也不导出 serial-id (序列号)。

插槽 iou#0-pci#2
此插槽包含一个十六进制 serial-id (序列号) 为 58071 的第三方扩展机箱。该扩展机箱内有两个可热插拔插槽 pcie1 和 special。pcie1 表示 PCI Express 插槽，其物理槽号为 1。插槽 special 带有一个源自平台、硬件或固件的标签。

插槽 iou#0-pci#3
此插槽包含一个 FRU 标识符为 BADF 的 Sun 扩展机箱。此扩展机箱包含三个 PCI Express 插槽 pcie1、pcie2 和 pcie3，物理槽号分别为 1、2、3，以及两个 PCI 插槽 pci1 和 pci2，物理槽号分别为 1 和 2。

以下命令显示了 I/O 扩展机箱（带连接的可热插拔和非热插拔主机槽）中的 PCI Express 和 PCI 挂接点的拓扑列表：

```
example# cfgadm -s cols=ap_id:info
Ap_Id                               Information
Slot1                               Location: Slot1
Slot2:iob4ffa56.pcie1               Location: Slot2/iob4ffa56.pcie1
Slot2:iob4ffa56.pcie2               Location: Slot2/iob4ffa56.pcie2
Slot5:iob3901.pci1                  Location: Slot2/iob3901.pci1
Slot5:iob3901.pci2                  Location: Slot2/iob3901.pci2
```

在本示例中，主机系统只有一个可热插拔插槽 Slot1。我们可以推断 Slot2 和 Slot5 不属于可热插拔插槽，因为这两个插槽在 cfgadm 中没有自行显示为挂接点。不过，Slot2 和 Slot5 都包含带可热插拔插槽的第三方扩展机箱。

以下命令显示了缺少特定设备属性的挂接点的拓扑列表：

```
example# cfgadm -s cols=ap_id:info
Ap_Id                               Information
px_pci7.pcie0                       Location: px_pci7.pcie0
px_pci11.pcie0                       Location: px_pci11.pcie0
px_pci11.pcie0:iob.pcie1              Location: px_pci11.pcie0/iob.pcie1
px_pci11.pcie0:iob.pcie2              Location: px_pci11.pcie0/iob.pcie2
px_pci11.pcie0:iob.pcie3              Location: px_pci11.pcie0/iob.pcie3
```

示例 3 解释 PCI Express 拓扑中的 ApId (续)

在本示例中，主机系统包含两个可热插拔插槽 `px_pci7.pcie0` 和 `px_pci11.pcie0`。在这种情况下，系统会将 `slot-id` 格式 (3) (缺省格式) 用于 *absolute-slot-path* 中的 *slot-path* 基本组件，因为框架无法获取足够的信息以生成其他优先级更高的、更具描述性的格式。

自右至左进行解释时，挂接点 `px_pci7.pcie0` 表示 PCI 设备编号为 0 (这并不表示物理槽号也是这一编号) 且已绑定到结点驱动程序 `px_pci` 实例 7 上的 PCI Express 插槽。类似地，挂接点 `px_pci11.pcie0` 表示 PCI 设备编号为 0 且已绑定到 `px_pci` 的驱动程序实例 11 上的 PCI Express 插槽。

在 `px_pci11.pcie0` 下是带三个可热插拔 PCI Express 插槽但没有 `serial-id` 的第三方扩展机箱。

以下命令显示了超过 `ApId` 字段长度限制的挂接点路径的拓扑列表：

```
example# cfgadm -s cols=ap_id:info
Ap_Id                               Information
pcie4                               Location: pcie4
pcie4:iobSUNW.pcie1                 Location: pcie4/iobSUNW.pcie1
pcie4:iobSUNW.pcie2                 Location: pcie4/iobSUNW.pcie2
iob8879c3f3.pci1                     Location: pcie4/iobSUNW.pcie2/iob8879c3f3.pci1
iob8879c3f3.pci2                     Location: pcie4/iobSUNW.pcie2/iob8879c3f3.pci2
iob8879c3f3.pci3                     Location: pcie4/iobSUNW.pcie2/iob8879c3f3.pci3
```

在本示例中，主机中只有一个可热插拔插槽 `pcie4`。在 `pcie4` 下连接的是一个 FRU 标识符为 `SUNW` 的 Sun 扩展机箱。嵌套在该扩展机箱 (`ApId pcie4:iobSUNW.pcie2`) PCI Express 插槽 `pcie2` 下的是另一个带三个可热插拔 PCI 插槽的扩展机箱。

对于以下项的 *absolute-slot-path* 格式：

```
pcie4/iobSUNW.pcie2/iob8879c3f3.pci1...3
```

由于其长度超过 `ApId` 字段的长度限制，且叶 *slot-path* 组件是全局唯一的，因此将使用 `ap_id` 格式 (2)，其中，*absolute-slot-path* 中的叶 *slot-path* 组件将用作最终 `ApId`。

以下命令显示了挂接点路径的拓扑列表，其中的挂接点路径超过 `ApId` 字段的长度限制，且缺少自行对叶 *slot-id* 进行唯一标识所需的足够信息 (例如，缺少 `serial-id` (序列号))：

```
example# cfgadm -s cols=ap_id:info
Ap_Id                               Information
pcie4                               Location: pcie4
pcie4:iob4567812345678.pcie3        Location: pcie4/iob4567812345678.pcie3
px_pci20.pcie0
```

示例 3 解释 PCI Express 拓扑中的 ApId (续)

```
Location: pcie4/iob4567812345678.pcie3/iob.pcie1
px_pci21.pcie0
Location: pcie4/iob4567812345678.pcie3/iob.pcie2
```

在本示例中，主机中只有一个可热插拔插槽 `pcie4`。在 `pcie4` 下连接的是十六进制 `serial-id`（序列号）为 `4567812345678` 的第三方扩展机箱。嵌套在该扩展机箱 (`ApId` `pcie4:iob4567812345678.pcie3`) PCI Express 插槽 `pcie3` 下的是另一个带两个可热插拔 PCI Express 插槽但没有 `serial-id`（序列号）的第三方扩展机箱。

对于以下项的 `absolute-slot-path` 格式：

```
pcie4/iob4567812345678.pcie3/iob.pcie1...2
```

由于其长度超过 `ApId` 字段的长度限制，且叶 `slot-path` 组件不是全局唯一的，因此将使用 `ap_id` 格式 (3)。在 `ap_id` 格式 (2) 中，将使用 `absolute-slot-path` 中叶 `slot-path` 组件的 `slot-id` 格式 (3)（缺省格式）作为最终 `ApId`。

叶组件 `.../iob.pcie1` 的缺省格式或 `slot-id` 格式 (3) 表示设备编号为 `0` 且已绑定到 `px_pci` 的驱动程序实例 `20` 的 PCI Express 插槽。类似地，叶组件 `.../iob.pcie2` 的缺省格式表示设备编号为 `0` 且已绑定到 `px_pci` 的驱动程序实例 `21` 的 PCI Express 插槽。

文件

`/usr/lib/cfgadm/shp.so.1`
适用于 PCI Express 和标准 PCI 热插拔的特定于硬件的库。

属性

有关以下属性的说明，请参见 `attributes(5)`：

属性类型	属性值
可用性	system/library
接口稳定性	Uncommitted（未确定）

另请参见

[cfgadm\(1M\)](#)、[cfgadm_pci\(1M\)](#)、[hotplugd\(1M\)](#)、[config_admin\(3CFGADM\)](#)、[libcfcfgadm\(3LIB\)](#)、

附注

`cfgadm_shp` 库依赖于 `hotplug` 服务，而该服务由 `smf(5)` 进行管理，其 FMRI（故障管理资源标识符）如下：

```
svc:/system/hotplug:default
```

要想正常运行 `cfgadm_shp` 库，必须启用该服务。有关详细信息，请参见 [hotplugd\(1M\)](#)。

引用名	cfgadm_sysctrl – EXX00 系统板管理				
用法概要	<pre>/usr/sbin/cfgadm -c function [-f] [-o disable-at-boot enable-at-boot] [-n -y] sysctrl0:slot# ... /usr/sbin/cfgadm -x quiesce-test sysctrl0:slot# /usr/sbin/cfgadm -x insert-test remove-test sysctrl0:slot# ... /usr/sbin/cfgadm -x set-condition-test=# sysctrl0:slot# ... /usr/sbin/cfgadm [-l] -o disable-at-boot enable-at-boot sysctrl0:slot# ...</pre>				
描述	sysctrl 硬件特定库 <code>/usr/platform/sun4u/lib/cfgadm/sysctrl.so.1</code> 提供动态重新配置功能，用于在 E6X00、E5X00、E4X00 以及 E3X00 系统中配置和断开系统板。可将 I/O 板和 CPU 板插入针对 Solaris 配置的运行中系统上的插槽，而无需重新引导。也可以从正在运行的系统断开和移除这两种类型的板，而无需重新引导。 系统插槽在设备树中显示为挂接点，系统机箱中的每个实际插槽对应一个挂接点。如果插槽中没有插入板，则插口状态为 <code>empty</code>（空）。如果板已断电并可随时移除，则插口状态为 <code>disconnected</code>（断开）。如果板已通电并连接到系统总线，则插口状态为 <code>connected</code>（已连接）。 插口状态为 <code>empty</code>（空）或 <code>disconnected</code>（断开）时，插卡状态为 <code>unconfigured</code>（未配置）。插口状态为 <code>connected</code>（已连接）时，插卡状态为 <code>unconfigured</code>（未配置）或 <code>configured</code>（已配置）。 处于 <code>configured</code>（已配置）状态时，板上的设备可供 Solaris 使用。处于 <code>unconfigured</code>（未配置）状态时，板上的设备不可用。 插入板会使插口状态从 <code>empty</code>（空）更改为 <code>disconnected</code>（断开）。移除板会使插口状态从 <code>disconnected</code>（断开）更改为 <code>empty</code>（空）。移除处于 <code>connected</code>（已连接）状态的板会导致操作系统崩溃，并可能导致系统的永久性损坏。				
选项	有关选项的更完整的描述，请参见 cfgadm(1M)。 所支持的选项如下： <table><tr><td><code>-c function</code></td><td>执行状态更改函数。将函数指定为 <code>connect</code>、<code>disconnect</code>、<code>configure</code> 或 <code>unconfigure</code>。</td></tr><tr><td><code>configure</code></td><td>将插卡状态更改为 <code>configured</code>（已配置）。</td></tr></table> 如果插口状态为 <code>disconnected</code>（断开），<code>configure</code> 函数会首先尝试连接该插口。<code>configure</code> 函数可遍历 <code>connect</code> 函数创建的 OBP 设备树，并创建 Solaris 设备树节	<code>-c function</code>	执行状态更改函数。将函数指定为 <code>connect</code> 、 <code>disconnect</code> 、 <code>configure</code> 或 <code>unconfigure</code> 。	<code>configure</code>	将插卡状态更改为 <code>configured</code> （已配置）。
<code>-c function</code>	执行状态更改函数。将函数指定为 <code>connect</code> 、 <code>disconnect</code> 、 <code>configure</code> 或 <code>unconfigure</code> 。				
<code>configure</code>	将插卡状态更改为 <code>configured</code> （已配置）。				

点，从而根据需要连接设备。对于 CPU/内存板，`configure` 会将 CPU 添加到处于 `powered-off`（已断电）状态的 CPU 列表中。这些操作对于 `psrinfo(1M)` 和 `psradm(1M)` 命令是可见的。为 CPU/内存板发布了两个内存挂接点。通过 `mount(1M)` 和 `ifconfig(1M)` 可使用新板上的 I/O 设备。要使用 CPU，请使用 `psradm -n` 使新处理器联机。使用 `cfgadm_ac(1M)` 可测试和配置内存区。

connect 将插口状态更改为 `connected`（已连接）。

当总线信号已连接并且板已测试时，要更改插口状态必须冻结系统总线。可通过执行 `quiesce` 操作冻结总线，该操作会停止所有进程活动并暂停所有驱动程序。由于 `quiesce` 操作以及随后的恢复操作比较耗时，而且并非所有驱动程序都支持这些操作，因此提供了 `-x quiesce-test`。当系统总线冻结时，固件会对正在连接的板进行测试。对于 I/O 板，该操作可在很短的时间内完成，而对于 CPU/内存板，由于需要进行 CPU 外部高速缓存测试，因此比较耗时。该操作过程中不进行内存测试。系统在 `quiesce` 操作开始前会提示用户进行确认。使用 `-y` 或 `-n` 选项应答该提示。如果板标记为 `disabled-at-boot`，则 `connect` 操作被拒绝，除非指定了强制标志 `-f` 或引导时启用标志 `-o enable-at-boot`。请参见 `-l`。

disconnect 将插口状态更改为 `disconnected`（断开）。

如果插卡状态为 `configured`（已配置），`disconnect` 函数会首先尝试

	取消对该插卡的配置。<code>disconnect</code> 操作不要求执行 <code>quiesce</code> 操作，因此耗时较短。板会断电，可随时移除。
<code>unconfigure</code>	将插卡状态更改为 <code>unconfigured</code>（未配置）。 在此过程中，会使板上的设备对 Solaris 不可见。I/O 板上的 I/O 设备将从 Solaris 设备树中删除。对于任何正在使用中的设备，会停止取消配置过程并报告设备正在使用。必须在使设备处于不忙状态后重试取消配置操作。对于 CPU/内存板，发出板取消配置操作前，必须已将内存更改为未配置状态。板上的 CPU 会脱机、断电并从 Solaris CPU 列表中删除。不能使与进程绑定的 CPU 脱机。有关使 CPU 脱机的更多信息，请参见 <code>psradm(1M)</code>、<code>psrinfo(1M)</code>、<code>pbind(1M)</code> 和 <code>p_online(2)</code>。
<code>-f</code>	强制将块连接到 <code>disabled-board-list</code> 非易失性变量中标记为 <code>disabled-at-boot</code> 的板。请参见《Platform Notes: Sun Enterprise 6x00/5x00/4x00/3x00 Systems》。
<code>-l</code>	列出选项。按 <code>cfgadm(1M)</code> 中的描述支持该选项。 <i>type</i> 字段可以是 <code>cpu/mem</code>、<code>mem</code>、<code>dual-sbus</code>、<code>sbus-upa</code>、<code>dual-pci</code>、<code>soc+sbus</code>、<code>s</code> 或 <code>unknown</code> 之一。 硬件特定的 <code>info</code> 字段的设置如下：<code>[disabled at boot]</code> <code>[non-detachable]</code> <code>[100 MHz capable]</code> 对于 <code>sbus-upa</code> 和 <code>soc+upa</code> 类型的板，首先显示以下附加信息：<code>[single buffered ffb double buffered ffb no ffb installed]</code> 对于磁盘类型的板，首先显示以下附加信息：<code>{target: # no disk}</code> <code>{target: # no disk}</code>

-o disable-at-boot enable-at-boot	修改 disabled-board-list 非易失性变量的状态。将 -o 选项与 -c <i>function</i> 或 -l 选项结合使用。 将 -o enable-at-boot 与 -c connect 结合使用可在连接 disabled-at-boot 板时覆盖块。										
-x insert-test remove-test	执行测试。 指定 remove-test 将指定插槽的驱动程序状态从 disconnected（断开）更改为 empty（空），而无需在自动测试过程中物理移除板。 指定 insert-test 将通过 remove-test 命令显示为 empty（空）的插槽的驱动程序状态更改为 disconnected（断开），就好像之前插入过板一样。										
-x quiesce-test sysctrl0:slot1	执行测试。 允许执行板的连接操作所需的 quiesce 操作。执行该测试可确认在当前软件和硬件配置下，可以停止系统。不能停止的设备或进程的名称将显示在错误消息中。该命令适用于任何有效的板挂接点，但由于所有系统都有插槽 1，因此推荐采用给定的形式。										
-x set-condition-test=#	执行测试。 允许设置系统板挂接点的状态，以便测试状态更改命令的策略逻辑。新的设置为指示如下状态值之一的数字： <table><tr><td>0</td><td>unknown</td></tr><tr><td>1</td><td>ok</td></tr><tr><td>2</td><td>failing</td></tr><tr><td>3</td><td>failed</td></tr><tr><td>4</td><td>unusable</td></tr></table>	0	unknown	1	ok	2	failing	3	failed	4	unusable
0	unknown										
1	ok										
2	failing										
3	failed										
4	unusable										

操作数

支持以下操作数：	
sysctrl0:slot#	EXX00 系统板的挂接点由 sysctrl 驱动程序的实例 0 (sysctrl0) 发布。挂接点名称的编号为 slot0 到 slot15。将 # 指定为 0 到 15 之间的数字，表示槽号。此格式符合 cfgadm(1M) 中提供的逻辑 ap_id 规范。对应的物理 ap_id 已在“文件”FILES 部分中列出。

文件 /usr/platform/sun4u/lib/cfgadm/sysctrl.so.1
 硬件特定库

 /devices/central@1f,0/fhc@0,f8800000/clock-board@0,900000:slot*
 挂接点

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/library/platform

另请参见 [cfgadm\(1M\)](#)、[cfgadm_ac\(1M\)](#)、[ifconfig\(1M\)](#)、[mount\(1M\)](#)、[pbind\(1M\)](#)、[psradm\(1M\)](#)、[psrinfo\(1M\)](#)

 《Sun Enterprise 6x00, 5x00, 4x00 and 3x00 Systems Dynamic Reconfiguration User's Guide》

 《Platform Notes: Sun Enterprise 6x00/5x00/4x00/3x00 Systems》

附注 有关 EXX00 系统 CPU/内存板动态重新配置的更多详细信息，请参见 《Sun Enterprise 6x00, 5x00, 4x00 and 3x00 Systems Dynamic Reconfiguration User's Guide》。

引用名

cfgadm_usb – USB 硬件特定的 cfgadm 命令

用法概要

```
/usr/sbin/cfgadm [-f] [-y | -n] [-v] -c function ap_id...

/usr/sbin/cfgadm -f [-y | -n] [-v] [-o hardware_options]
-x hardware_function ap_id...

/usr/sbin/cfgadm -v [-a] [-s listing_option]
-l [ap_id | ap_type...]

/usr/sbin/cfgadm -v -h [ap_id]...
```

描述

通用串行总线 (Universal Serial Bus, USB) 硬件特定库 `/usr/lib/cfgadm/usb.so.1` 提供了通过 `cfgadm(1M)` 命令管理 USB 设备的功能。`cfgadm` 在挂接点上运行。有关挂接点的详细信息，请参见 `cfgadm(1M)`。

对于 USB 管理，支持的挂接点只有连接到 USB 总线的集线器的端口。

挂接点通过挂接点 ID (`ap_id`) 进行命名。USB 总线是分层的，因此 `ap_id` 也是分层的。USB 集线器端口的编号从 1 到 n 。所有 USB `ap_id` 均由以下格式的字符串组成：

```
usbN/A[.B[.C[...]]]
```

其中

N 是系统上的第 N 个 USB 主机控制器，
 A 是根（顶层）集线器的端口 A 。
 B 是插入它上面的集线器端口 A 中的集线器的端口 B 。
 C 是插入它上面的集线器端口 B 中的集线器的端口 C ，依此类推。

例如，USB 控制器 0（唯一的控制器）的根集线器上的第一个端口具有以下逻辑 `ap_id`：

```
usb0/1
```

同样，插入第一个 USB 控制器的根集线器上第一个端口中的第一个外部集线器上的第二个端口具有以下逻辑 `ap_id`：

```
usb0/1.2
```

例如，如果 `ap_id` 为 `usb0/1.4.3.4`，则表示插入集线器（插入集线器（插入系统上的第一个 USB 主机控制器的根集线器的端口 1）端口 4）端口 3 的集线器的端口 4。

```
example# cfgadm -l
```

Ap_Id	Type	Receptacle	Occupant	Condition
usb0/1	USB-hub	connected	configured	ok
usb0/2	unknown	empty	unconfigured	ok
usb0/1.1	USB-storage	connected	configured	ok
usb0/1.2	unknown	empty	unconfigured	ok
usb0/1.3	unknown	empty	unconfigured	ok
usb0/1.4	USB-device	connected	configured	ok

USB2.0 芯片具有一个 EHCI 主机 USB2.0 主机控制器和多个伴随 USB 1.x 主机控制器（OHCI 或 UHCI 主机控制器）。

插入 USB2.0 设备后，它将在可能没有一对一映射到系统上外部物理端口号的 EHCI 逻辑端口上显示。插入 USB1.x 设备后，EHCI 主机控制器将该设备重新路由到伴随主机控制器，并且该设备将在伴随控制器的逻辑端口号上显示。

逻辑端口号到物理端口号的映射可以变得很复杂。例如：

```
% cfgadm
```

Ap_Id	Type	Receptacle	Occupant	Condition
c0	scsi-bus	connected	configured	unknown
usb0/1	usb-mouse	connected	configured	ok
usb0/2	usb-kbd	connected	configured	ok
usb0/3	unknown	empty	unconfigured	ok
usb0/4	usb-hub	connected	configured	ok
usb0/4.1	unknown	empty	unconfigured	ok
usb0/4.2	unknown	empty	unconfigured	ok
usb0/4.3	unknown	empty	unconfigured	ok
usb0/4.4	usb-storage	connected	configured	ok
usb1/1	unknown	empty	unconfigured	ok
usb1/2	unknown	empty	unconfigured	ok
usb1/3	unknown	empty	unconfigured	ok
usb2/1	unknown	empty	unconfigured	ok
usb2/2	usb-device	connected	configured	ok
usb3/1	unknown	empty	unconfigured	ok
usb3/2	unknown	empty	unconfigured	ok
usb3/3	unknown	empty	unconfigured	ok
usb3/4	unknown	empty	unconfigured	ok
usb3/5	unknown	empty	unconfigured	ok

在本例中，usb0 是板载 USB 1.x 主机控制器。usb1 和 usb2 是伴随 OHCI USB1.x 主机控制器，usb3 是 EHCI USB2.0 主机控制器。

下表中显示了此 USB2.0 芯片有些混乱的路由：

logical port number	physical port number
-----	-----
usb1/1	internal port 1
usb1/2	external port 1
usb1/3	external port 3
usb2/1	internal port 2
usb2/2	external port 2
usb3/1	internal port 1
usb3/2	internal port 2
usb3/3	external port 1
usb3/4	external port 2
usb3/5	external port 3

遗憾的是，具体的路由通常只能通过实验确定。

USB 端口上挂接点的插口状态含义如下：

connected (已连接)

USB 端口已通电且启用。USB 设备插入该端口中。该设备在逻辑上连接到 USB 总线。

disconnected (断开)

USB 端口已通电且启用。USB 设备插入到该端口中。该设备已在逻辑上与 USB 总线断开连接（使用 `cfgadm -c disconnect` 命令）。

empty (空)

USB 端口已通电，但没有插入任何设备。

USB 端口上的设备在 USB 端口挂接点的插卡状态含义如下：

configured (已配置)

USB 端口上的 USB 设备已配置且可供 Solaris 使用。

unconfigured (未配置)

使用 `cfgadm -c unconfigure` 命令使 USB 端口上的 USB 设备显式脱机，或没有成功配置该设备而无法用于 Solaris（例如，没有驱动程序或设备存在问题）。

挂接点的状态含义如下：

ok (正常)

正常状态—可供使用。

FAILED (失败)

未使用。

failed (失败)

未使用。

unusable (不可用)

在设备已被应用程序打开的情况下，用户物理移除了该设备（可能有未完成的 I/O）。用户需要重新插入同一物理设备并正确关闭应用程序，然后再次移除该设备。此操作完成之前，端口不能配置其他插入的设备。

如果无法将原始设备重新插入端口，请参见《[Oracle Solaris 管理：常见任务](#)》了解清除此挂接点状态的说明。

unknown (未知)

未使用。

可以随时对 USB 设备进行热插拔，系统会检测该事件并采取相应的操作。

在从 USB 移除插口上的设备之前，不必将插口转换为 **disconnected**（断开）状态。但是，建议不要热移除当前正在使用的设备（如当前由卷管理器或某些其他应用程序打开的可移除磁盘）。

选项 `cfgadm` 可定义多种类型的操作。这些操作包括调用配置状态更改 (-c)、调用硬件特定函数 (-x) 以及获取配置管理帮助消息 (-h)。

如果上述任一操作失败，设备和挂接点可能不会处于期望的状态。使用 `cfgadm -l` 命令可显示设备的当前状态。

有关其他所有选项的含义，可参见 `cfgadm(1M)` 中的定义。

所支持的选项如下：

-c function

以下通用命令是针对 USB 硬件特定库定义的。支持以下配置状态更改操作：

configure

如果已有 USB 设备插入端口，则此命令将尝试配置该设备并完成所有设置，使其可供 Solaris 使用。如果需要，此命令可执行暗示的 `connect`（与 `disconnect` 相反）。如果已经配置了该端口上的设备，则此命令不完成任何操作，而是返回错误消息。成功执行此命令后，该设备可供 Solaris 使用。

disconnect

对 `ap_id` 执行 `unconfigure`（如果不处于 `unconfigured`（未配置）），然后将插口转换为 `disconnected`（断开）状态，即使设备仍插在端口中。发出 `cfgadm -c configure` 命令或以物理方式热插入设备可使设备的插口状态恢复为 `connected`（已连接），并使其插卡状态恢复为 `configured`（已配置），前提是找到驱动程序且枚举和配置设备时未出现问题。

unconfigure

使插入端口中的设备不可供 Solaris 使用（使其脱机）。如果成功，`cfgadm` 会将该 `ap_id` 的插卡状态报告为 `unconfigured`（未配置）。向 `ap_id` 发出 `configure` 命令（如果成功）可使其插卡状态恢复为 `configured`（已配置，即联机），因为这会以物理方式将设备热插入端口。

-f

不支持。

-h ap_id

通过对任何 USB 挂接点使用帮助选项可获取 USB 特定的帮助。

-l[v]

按照 `cfgadm(1M)` 中的描述使用 `-l` 选项。如果该选项与 `-v` 选项一起使用，`Information` 字段将包含以下 USB 特定信息：

- `Mfg`：生产商字符串 (`iManufacturer`)
- `Product`：产品字符串 (`iProduct`)
- `NConfigs`：设备支持的配置总数 (`bNumConfigurations`)。
- `Config`：十进制形式的当前配置设置（配置索引，非配置值）。
- 当前配置的配置字符串描述符 (`iConfiguration`)

有关这些字段的描述，请参见通用串行总线规范。

-o hardware_options

只有硬件特定的命令 `-x usb_config` 支持硬件选项。有关可用选项的说明，请参见以下关于该命令的描述。

-s listing_options

可使用 `select` 子选项列出 USB 类的挂接点。请参见 [cfgadm\(1M\)](#)。

-x hardware_function

定义了以下硬件特定的函数：

usb_config -o config=*n*

此命令要求使用 `-o` 选项指定必需的 `config` 值。

将 `ap_id` 上的多配置 USB 设备的 USB 配置设置为配置索引 `n`。此后该设备都设置为此配置，而且在设备的重新引导、热移除和取消配置/配置中将一直保持此设置。

`n` 的有效值范围是 0 到 (`Nconfigs - 1`)。可以通过先后执行 `disconnect` 和 `configure` 来复位设备。`configure` 操作会将设备配置为新的配置设置。

如果上述任一步骤失败，则配置文件和设备将恢复为之前的状态，并显示错误消息。

usb_reset

执行设备的软件复位（重新枚举）。这等效于移除设备然后再次插入。如果集线器支持单个端口的关开机循环，则集线器上的端口将进行关开机循环。

如果连接的设备是集线器，则此函数将复位该集线器以及其作为根的树中的任何设备。

如果上述任一步骤失败，设备将恢复为之前的状态，并显示错误消息。

状态表：挂接点状态与命令：

Valid states:

empty/unconfigured	→ no device connected
disconnected/unconfigured	→ logically disconnected, unavailable, devinfo node removed, device physically connected
connected/unconfigured	→ logically connected, unavailable, devinfo node present
connected/configured	→ connected, available

下表说明操作或命令导致的状态转换：

current state	operation	new state
-----	-----	-----
empty/ unconfigured:		
	device plugged in:	connected/configured or connected/unconfigured (if enumeration failed)
	device removed:	n/a
	cfgadm -c unconfigure:	empty/unconfigured
	cfgadm -c configure:	empty/unconfigured
	cfgadm -c disconnect:	empty/unconfigured (no-op and error)
disconnected/ unconfigured:		
	device plugged in:	n/a
	device removed:	empty/unconfigured
	cfgadm -c unconfigure:	disconnected/unconfigured
	cfgadm -c configure:	connected/configured, or connected/unconfigured (if reenumeration failed)
	cfgadm -c disconnect:	disconnected/unconfigured
connected/unconfigured:		
	device plugged in:	n/a
	device removed:	empty/unconfigured
	cfgadm -c unconfigure:	connected/unconfigured
	cfgadm -c configure:	connected/configured, or connected/unconfigured (if reenumeration failed)
	cfgadm -c disconnect:	disconnected/unconfigured
connected/configured:		
	device plugged in:	n/a
	device removed:	empty/unconfigured or connected/configured, but with ap condition 'unusable' if device was open when removed
	cfgadm -c unconfigure:	connected/unconfigured
	cfgadm -c configure:	connected/configured
	cfgadm -c disconnect:	disconnected/unconfigured

示例

示例1 列出所有 USB 设备的状态
以下命令列出系统上所有 USB 设备的状态：

```
# cfgadm
Ap_Id      Type      Receptacle  Occupant    Condition
usb0/1     USB-hub   connected   configured  ok
```


示例1 列出所有 USB 设备的状态 (续)

usb0/2	unknown	empty	unconfigured	ok
usb0/1.1	USB-storage	connected	configured	ok
usb0/1.2	unknown	empty	unconfigured	ok
usb0/1.3	unknown	empty	unconfigured	ok
usb0/1.4	USB-device	connected	configured	ok

请注意，cfgadm 将 ap_id usb0/1.4 上的 USB-device 设备视为单个单元，因为它当前不能控制单个接口。

示例2 列出没有插入任何设备的端口的状态

以下命令列出没有插入任何设备的端口的状态：

```
example# cfgadm -l usb0/1.3
```

Ap_Id	Type	Receptacle	Occupant	Condition
usb0/1.3	unknown	empty	unconfigured	ok

示例3 列出插入设备的同一端口的状态

以下命令列出以物理方式插入成功配置的设备后，同一端口的状态：

```
example# cfgadm -l usb0/1.3
```

Ap_Id	Type	Receptacle	Occupant	Condition
usb0/1.3	USB-hub	connected	configured	ok

示例4 取消对现有 USB 设备的配置

以下命令取消对连接到 usb0/1.3 的 USB 设备的配置，然后显示 ap_id 的状态：

```
example# cfgadm -c unconfigure usb0/1.3
Unconfigure the device: /devices/pci@0,0/pci8086,7112@7,2/hub@2:2.3
This operation suspends activity on the USB device
Continue (yes/no)?
```

Enter:

y

```
example# cfgadm -l usb0/1.3
```

Ap_Id	Type	Receptacle	Occupant	Condition
usb0/1.3	unknown	connected	unconfigured	ok

示例5 取消配置并在逻辑上断开现有 USB 设备

以下命令取消连接到 usb0/1.3 的 USB 设备的配置，并在逻辑上断开其连接：

```
example# cfgadm -c disconnect usb0/1.3
Disconnect the device: /devices/pci@0,0/pci8086,7112@7,2/hub@2:2.3
This operation suspends activity on the USB device
```

示例5 取消配置并在逻辑上断开现有 USB 设备 (续)

Continue (yes/no)?

Enter:

y

example# **cfgadm -l usb0/1.3**

Ap_Id	Type	Receptacle	Occupant	Condition
usb0/1.3	unknown	disconnected	unconfigured	ok

disconnect 暗示 **cfgadm** 首先执行 **unconfigure**。现在插口状态显示为 **disconnected**，即使设备在物理上仍处于连接状态。在这种情况下，以物理方式热插入或**对 ap_id** 使用 **cfgadm -c configure** 可以使其恢复联机。

示例6 配置之前未配置的 USB 设备

以下命令配置之前连接到 **usb0/1.3** 的 USB 设备：

example # **cfgadm -yc configure usb0/1.3**

example# **cfgadm -l usb0/1.3**

Ap_Id	Type	Receptacle	Occupant	Condition
usb0/1.3	unknown	connected	configured	ok

示例7 复位 USB 设备

以下命令复位 USB 设备：

example# **cfgadm -x usb_reset usb0/1.3**

Reset the device: /devices/pci@0,0/pci8086,7112@7,2/hub@2:2.3

This operation suspends activity on the USB device

Continue (yes/no)?

Enter:

y

示例8 显示有关 USB 设备的详细信息

以下命令显示有关 USB 设备的详细信息。该设备在 **Information** 字段中显示以下 USB 特定的信息：

- 生厂商字符串：Iomega
- 产品字符串：USB Zip 250
- 支持的配置数：1
- 当前活动的配置：0
- 配置 0 的配置字符串描述符：缺省

example# **cfgadm -lv usb0/1.5**

Ap_Id	Receptacle	Occupant	Condition	Information

示例 8 显示有关 USB 设备的详细信息 (续)

```
When      Type      Busy      Phys_Id
usb0/1.5  connected  configured  ok      Mfg:"Io
mega" Product:"USB Zip 250" NConfigs:1 Config:0 : Default

example# cfgadm -l -s "cols=ap_id:info" usb0/1.5
Ap_Id      Information
usb0/1.5   Mfg:"Iomega" Product:"USB Zip 250"
NConfigs:1 Config:0 : Default
```

示例 9 显示有关所有 USB 设备的详细信息

以下命令显示有关系统上所有 USB 设备的详细信息：

```
example# cfgadm -l -s "select=class(usb),cols=ap_id:info"
Ap_Id      Information
usb0/1      Mfg:<undefined> Product:<undefined>
NConfigs:1 Config:0 <no cfg str descr>
usb0/2
usb0/1.1    Mfg:<undefined> Product:<undefined>
NConfigs:1 Config:0 <no cfg str descr>
usb0/1.2
usb0/1.3
usb0/1.4    Mfg:"Wizard" Product:"Modem/ISDN"
NConfigs:3 Config:1 : V.90 Analog Modem
usb0/1.5    Mfg:"Iomega" Product:"USB Zip 250"
NConfigs:1 Config:0 : Default
usb0/1.6    Mfg:"SOLID YEAR" Product:"SOLID YEAR
USB"NConfigs:1 Config:0 <no cfg str descr>
usb0/1.7
```

仅包含 ap_id 的行表示空端口。可以过滤出这些端口。以下示例仅列出连接有设备的 USB ap_id 以及有关这些设备的信息。

```
example# cfgadm -l -s "select=class(usb),cols=ap_id:info" | grep Mfg
usb0/1      Mfg:<undefined> Product:<undefined>
NConfigs:1 Config:0 <no cfg str descr>
usb0/1.1    Mfg:<undefined> Product:<undefined>
NConfigs:1 Config:0 <no cfg str descr>
usb0/1.4    Mfg:"Wizard" Product:"Modem/ISDN"
NConfigs:3 Config:1 : V.90 Analog Modem
usb0/1.5    Mfg:"Iomega" Product:"USB Zip 250"
NConfigs:1 Config:0 : Default
usb0/1.6    Mfg:"SOLID YEAR" Product:"SOLID YEAR USB"
Config:0 <no cfg str descr>
```

示例 10 列出有关多配置 USB 设备的信息

以下示例列出有关多配置 USB 设备的信息。

请注意 NConfigs 字段：该设备可用的配置有 0、1 和 2（0 到 (NConfigs-1)）。

```
example# cfgadm -l -s "cols=ap_id:info" usb0/1.4
Ap_Id                               Information
usb0/1.4                           Mfg:"Wizard"  Product:"Modem/ISDN"
NConfigs:3  Config:1 V.90 Analog Modem"
```

示例 11 设置多配置 USB 设备的当前配置

以下示例设置多配置 USB 设备的当前配置：

```
example# cfgadm -o config=2 -x usb_config usb0/1.4
Setting the device: /devices/pci@1f,2000/usb@1/device@3
to USB configuration 2
This operation suspends activity on the USB device
Continue (yes/no)?
```

```
Enter:

y

USB configuration changed successfully.
```

相同类型的多个设备位于同一总线时，应该检查设备路径以确保引用的是正确的设备实例。在 Information 字段提供了该信息。

文件

/usr/lib/cfgadm/usb.so.1

用于管理通用 USB 设备的硬件特定库

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/library

另请参见

[cfgadm\(1M\)](#)、[config_admin\(3CFGADM\)](#)、[attributes\(5\)](#)、[scsa2usb\(7D\)](#)、[usba\(7D\)](#)

通用串行总线 1.1 规范 (www.usb.org)

《Oracle Solaris 管理：常见任务》

附注

[cfgadm\(1M\)](#) 无法取消配置、断开、复位或更改当前由任何应用程序打开的任何 USB 设备的配置。如果分层结构中有某个设备由应用程序打开，则也无法在集线器上执行这些操作。有关取消对当前正在使用的 USB 海量存储设备的配置的信息，请参见 [scsa2usb\(7D\)](#)。

只有超级用户才能对挂接点执行任何函数。但是，如果只是列出挂接点，那么非超级用户也可完成。

引用名 chat – automated conversational exchange tool

用法概要 chat [*options*] *script*

描述 The chat program implements a conversational text-based exchange between the computer and any serial device, including (but not limited to) a modem, an ISDN TA, and the remote peer itself, establishing a connection between the Point-To-Point Protocol daemon (pppd) and the remote pppd process.

选项 The chat command supports the following options:

- f <*chat file*> Read the chat script from the chat file. This option is mutually exclusive with the chat script parameters. You must have read access to use the file. Multiple lines are permitted in the file. Use the space or horizontal tab characters to separate the strings.
- t <*timeout*> Set the timeout for the expected string to be received. If the string is not received within the time limit, the reply string is not sent. If specified, a 'subexpect' (alternate reply) string can be sent. Otherwise, if no alternate reply strings remain, the chat script fails.. A failed script will cause the chat program to terminate with a non-zero error code.
- r <*report file*> Set the file for output of the report strings. If you use the keyword REPORT, the resulting strings are written to this file. If the -r option is not used and you use the REPORT keyword, the stderr file is used for the report strings.
- e Start with the echo option turned on. You turn echo on or off at specific points in the chat script using the ECHO keyword. When echoing is enabled, all output from the modem is echoed to stderr.
- E Enables environment variable substitution within chat scripts using the standard \$xxx syntax.
- v Request that the chat script execute in a verbose mode. The chat program logs the execution state of the chat script as well as all text received from the modem and output strings sent to the modem. The default is to log through [syslog\(3C\)](#) with facility local2; the logging method is alterable using the -S and -s options.
- V Request that the chat script be executed in a stderr verbose mode. The chat program logs all text received from the modem and output strings sent to the modem to stderr. stderr is usually the local console at the station running the chat or pppd program.
- s Use stderr. Log messages from -v and error messages are sent to stderr.

-S	Do not use syslog. By default, error messages are set to syslog. This option prevents log messages from -v and error messages from being sent to syslog.
-T <phone number>	Pass in an arbitrary string (usually a telephone number) that will be substituted for the \T substitution metacharacter in a send string.
-U <phone number 2>	Pass in a second string (usually a telephone number) that will be substituted for the \U substitution metacharacter in a send string. This is useful when dialing an ISDN terminal adapter that requires two numbers.
script	If the script is not specified in a file with the -f option, the script is included as parameters to the chat program.

扩展描述

Chat Script

The chat script defines communications. A script consists of one or more "expect-send" pairs of strings separated by spaces, with an optional "subexpect-subsend" string pair, separated by a dash (as in the following example:)

```
ogin:-BREAK-ogin: ppp ssword: hello2u2
```

The example indicates that the chat program should expect the string "ogin:". If it fails to receive a login prompt within the time interval allotted, it sends a break sequence to the remote and then expects the string "ogin:". If the first "ogin:" is received, the break sequence is not generated.

Upon receiving the login prompt, the chat program sends the string "ppp" and then expects the prompt "ssword:". When the password prompt is received, it sends the password hello2u2.

A carriage return is normally sent following the reply string. It is not expected in the "expect" string unless it is specifically requested by using the \r character sequence.

The expect sequence should contain only what is needed to identify the received data. Because it's stored on a disk file, it should not contain variable information. Generally it is not acceptable to look for time strings, network identification strings, or other variable pieces of data as an expect string.

To correct for characters that are corrupted during the initial sequence, look for the string "ogin:" rather than "login:". The leading "l" character may be received in error, creating problems in finding the string. For this reason, scripts look for "ogin:" rather than "login:" and "ssword:" rather than "password:".

An example of a simple script follows:

```
ogin: ppp ssword: hello2u2
```

The example can be interpreted as: expect ogin:, send ppp, expect ...ssword:, send hello2u2.

When login to a remote peer is necessary, simple scripts are rare. At minimum, you should include sub-expect sequences in case the original string is not received. For example, consider the following script:

```
ogin:--ogin: ppp ssword: hello2u2
```

This script is more effective than the simple one used earlier. The string looks for the same login prompt; however, if one is not received, a single return sequence is sent and then the script looks for login: again. If line noise obscures the first login prompt, send the empty line to generate a login prompt again.

Comments

Comments can be embedded in the chat script. Comment lines are ignored by the chat program. A comment starts with the hash (“#”) character in column one. If a # character is expected as the first character of the expect sequence, quote the expect string. If you want to wait for a prompt that starts with a # character, write something like this:

```
# Now wait for the prompt and send logout string
'# ' logout
```

Sending Data From A File

If the string to send begins with an at sign (“@”), the remainder of the string is interpreted as the name of the file that contains the string. If the last character of the data read is a newline, it is removed. The file can be a named pipe (or fifo) instead of a regular file. This enables chat to communicate with another program, for example, a program to prompt the user and receive a password typed in.

Abort

Many modems report the status of a call as a string. These status strings are often “CONNECTED” or “NO CARRIER” or “BUSY.” If the modem fails to connect to the remote, you can terminate the script. Abort strings may be specified in the script using the ABORT sequence. For example:

```
ABORT BUSY ABORT 'NO CARRIER' '' ATZ OK ATDT5551212 CONNECT
```

This sequence expects nothing and sends the string ATZ. The expected response is the string OK. When OK is received, the string ATDT5551212 dials the telephone. The expected string is CONNECT. If CONNECT is received, the remainder of the script is executed. When the modem finds a busy telephone, it sends the string BUSY, causing the string to match the abort character sequence. The script fails because it found a match to the abort string. If the NO CARRIER string is received, it aborts for the same reason.

Clr_Abort

The CLR_ABORT sequence clears previously set ABORT strings. ABORT strings are kept in an array of a pre-determined size; CLR_ABORT reclaims the space for cleared entries, enabling new strings to use that space.

Say

The SAY string enables the script to send strings to a user at a terminal via standard error. If chat is being run by pppd and pppd is running as a daemon (detached from its controlling terminal), standard error is normally redirected to the `/etc/ppp/connect-errors` file.

SAY strings must be enclosed in single or double quotes. If carriage return and line feed are required for the output, you must explicitly add them to your string.

The SAY string can provide progress messages to users even with “ECHO OFF.” For example, add a line similar to the following to the script:

```
ABORT BUSY
ECHO OFF
SAY "Dialing your ISP...\n"
'' ATDT5551212
TIMEOUT 120
SAY "Waiting up to 2 minutes for connection ..."
CONNECT ''
SAY "Connected, now logging in ...\n"
ogin: account
ssword: pass
$ \c
SAY "Logged in OK ... \n"
```

This sequence hides script detail while presenting the SAY string to the user. In this case, you will see:

```
Dialing your ISP...
Waiting up to 2 minutes for connection...Connected, now logging in...
Logged in OK ...
```

Report

REPORT is similar to the ABORT string. With REPORT, however, strings and all characters to the next control character (such as a carriage return), are written to the report file.

REPORT strings can be used to isolate a modem's transmission rate from its CONNECT string and return the value to the chat user. Analysis of the REPORT string logic occurs in conjunction with other string processing, such as looking for the expect string. It's possible to use the same string for a REPORT and ABORT sequence, but probably not useful.

Report strings may be specified in the script using the REPORT sequence. For example:

```
REPORT CONNECT
ABORT BUSY
ATDT5551212 CONNECT
ogin: account
```

The above sequence expects nothing, then sends the string ATDT5551212 to dial the telephone. The expected string is CONNECT. If CONNECT is received, the remainder of the

script is executed. In addition, the program writes the string CONNECT to the report file (specified by -r) in addition to any characters that follow.

Clr_Report CLR_REPORT clears previously set REPORT strings. REPORT strings are kept in an array of a pre-determined size; CLR_REPORT reclaims the space for cleared entries so that new strings can use that space.

Echo ECHO determines if modem output is echoed to stderr. This option may be set with the -e option, but can also be controlled by the ECHO keyword. The "expect-send" pair ECHO ON enables echoing, and ECHO OFF disables it. With ECHO, you can select which parts of the conversation should be visible. In the following script:

```
ABORT 'BUSY'
ABORT 'NO CARRIER'
"" AT&F
OK\r\n ATD1234567
\r\n \c
ECHO ON
CONNECT \c
ogin: account
```

All output resulting from modem configuration and dialing is not visible, but output is echoed beginning with the CONNECT (or BUSY) message.

Hangup The HANGUP option determines if a modem hangup is considered as an error. HANGUP is useful for dialing systems that hang up and call your system back. HANGUP can be ON or OFF. When HANGUP is set to OFF and the modem hangs up (for example, following the first stage of logging in to a callback system), chat continues running the script (for example, waiting for the incoming call and second stage login prompt). When the incoming call is connected, use the HANGUP ON string to reinstall normal hang up signal behavior. An example of a simple script follows:

```
ABORT 'BUSY'
"" AT&F
OK\r\n ATD1234567
\r\n \c
CONNECT \c
'Callback login:' call_back_ID
HANGUP OFF
ABORT "Bad Login"
'Callback Password:' Call_back_password
TIMEOUT 120
CONNECT \c
HANGUP ON
ABORT "NO CARRIER"
ogin:--BREAK--ogin: real_account
```

Timeout	The initial timeout value is 45 seconds. Use the <code>-t</code> parameter to change the initial timeout value. To change the timeout value for the next expect string, the following example can be used: <pre>''AT&F OK ATDT5551212 CONNECT \c TIMEOUT 10 ogin:--ogin: username TIMEOUT 5 assword: hello2u2</pre> The example changes the timeout to ten seconds when it expects the login: prompt. The timeout is changed to five seconds when it looks for the password prompt. Once changed, the timeout value remains in effect until it is changed again.										
EOT	The EOT special reply string instructs the chat program to send an EOT character to the remote. This is equivalent to using <code>^D\c</code> as the reply string. The EOT string normally indicates the end-of-file character sequence. A return character is not sent following the EOT. The EOT sequence can be embedded into the send string using the sequence <code>^D</code>.										
BREAK	The BREAK special reply string sends a break condition. The break is a special transmitter signal. Many UNIX systems handle break by cycling through available bit rates, and sending break is often needed when the remote system does not support autobaud. BREAK is equivalent to using <code>\K\c</code> as the reply string. You embed the break sequence into the send string using the <code>\K</code> sequence.										
Escape Sequences	Expect and reply strings can contain escape sequences. Reply strings accept all escape sequences, while expect strings accept most sequences. A list of escape sequences is presented below. Sequences that are not accepted by expect strings are indicated. <table> <tr> <td><code>''</code></td><td>Expects or sends a null string. If you send a null string, chat sends the return character. If you expect a null string, chat proceeds to the reply string without waiting. This sequence can be a pair of apostrophes or quote mark characters.</td></tr> <tr> <td><code>\b</code></td><td>Represents a backspace character.</td></tr> <tr> <td><code>\c</code></td><td>Suppresses the newline at the end of the reply string. This is the only method to send a string without a trailing return character. This sequence must be at the end of the send string. For example, the sequence <code>hello\c</code> will simply send the characters h, e, l, l, o. (Not valid in expect.)</td></tr> <tr> <td><code>\d</code></td><td>Delay for one second. The program uses <code>sleep(1)</code> which delays to a maximum of one second. (Not valid in expect.)</td></tr> <tr> <td><code>\K</code></td><td>Insert a BREAK. (Not valid in expect.)</td></tr> </table>	<code>''</code>	Expects or sends a null string. If you send a null string, chat sends the return character. If you expect a null string, chat proceeds to the reply string without waiting. This sequence can be a pair of apostrophes or quote mark characters.	<code>\b</code>	Represents a backspace character.	<code>\c</code>	Suppresses the newline at the end of the reply string. This is the only method to send a string without a trailing return character. This sequence must be at the end of the send string. For example, the sequence <code>hello\c</code> will simply send the characters h, e, l, l, o. (Not valid in expect.)	<code>\d</code>	Delay for one second. The program uses <code>sleep(1)</code> which delays to a maximum of one second. (Not valid in expect.)	<code>\K</code>	Insert a BREAK. (Not valid in expect.)
<code>''</code>	Expects or sends a null string. If you send a null string, chat sends the return character. If you expect a null string, chat proceeds to the reply string without waiting. This sequence can be a pair of apostrophes or quote mark characters.										
<code>\b</code>	Represents a backspace character.										
<code>\c</code>	Suppresses the newline at the end of the reply string. This is the only method to send a string without a trailing return character. This sequence must be at the end of the send string. For example, the sequence <code>hello\c</code> will simply send the characters h, e, l, l, o. (Not valid in expect.)										
<code>\d</code>	Delay for one second. The program uses <code>sleep(1)</code> which delays to a maximum of one second. (Not valid in expect.)										
<code>\K</code>	Insert a BREAK. (Not valid in expect.)										

<code>\n</code>	Send a newline or linefeed character.
<code>\N</code>	Send a null character. The same sequence may be represented by <code>\0</code> . (Not valid in expect.)
<code>\p</code>	Pause for 1/10th of a second. (Not valid in expect.)
<code>\q</code>	Suppress writing the string to syslog. The string <code>?????</code> is written to the log in its place. (Not valid in expect.)
<code>\r</code>	Send or expect a carriage return.
<code>\s</code>	Represents a space character in the string. Can be used when it is not desirable to quote the strings which contains spaces. The sequence <code>'HI TIM'</code> and <code>HI\sTIM</code> are the same.
<code>\t</code>	Send or expect a tab character.
<code>\T</code>	Send the phone number string as specified with the <code>-T</code> option. (Not valid in expect.)
<code>\U</code>	Send the phone number 2 string as specified with the <code>-U</code> option. (Not valid in expect.)
<code>\\</code>	Send or expect a backslash character.
<code>\ddd</code>	Collapse the octal digits (<code>ddd</code>) into a single ASCII character and send that character. (<code>\000</code> is not valid in an expect string.)
<code>^C</code>	Substitute the sequence with the control character represented by <code>C</code> . For example, the character <code>DC1 (17)</code> is shown as <code>^Q</code> . (Some characters are not valid in expect.)

环境变量

Environment variables are available within chat scripts if the `-E` option is specified on the command line. The metacharacter `$` introduces the name of the environment variable to substitute. If the substitution fails because the requested environment variable is not set, nothing is replaced for the variable.

退出状态

The chat program terminates with the following completion codes:

- `0` Normal program termination. Indicates that the script was executed without error to normal conclusion.
- `1` One or more of the parameters are invalid or an expect string was too large for the internal buffers. Indicates that the program was not properly executed.
- `2` An error occurred during the execution of the program. This may be due to a read or write operation failing or chat receiving a signal such as `SIGINT`.
- `3` A timeout event occurred when there was an expect string without having a `"-subsend"` string. This indicates that you may not have programmed the script correctly for the condition or that an unexpected event occurred and the expected string could not be found.

- 4 The first string marked as an ABORT condition occurred.
- 5 The second string marked as an ABORT condition occurred.
- 6 The third string marked as an ABORT condition occurred.
- 7 The fourth string marked as an ABORT condition occurred.
- ... The other termination codes are also strings marked as an ABORT condition.

To determine which event terminated the script, use the termination code. It is possible to decide if the string "BUSY" was received from the modem versus "NO DIALTONE." While the first event may be retried, the second probably will not succeed during a retry.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/ppp
Interface Stability	Committed

另请参见

[sleep\(1\)](#), [uucp\(1C\)](#), [pppd\(1M\)](#), [uucico\(1M\)](#), [syslog\(3C\)](#), [attributes\(5\)](#)

Additional information on chat scripts are available with UUCP documentation. The chat script format was taken from scripts used by the `uucico` program.

- 引用名

check-hostname – check if sendmail can determine the system's fully-qualified host name
- 用法概要

`/usr/sbin/check-hostname`
- 描述

The `check-hostname` script is a migration aid for [sendmail\(1M\)](#). This script tries to determine the local host's fully-qualified host name (FQHN) in a manner similar to [sendmail\(1M\)](#). If `check-hostname` is able to determine the FQHN of the local host, it reports success. Otherwise, `check-hostname` reports how to reconfigure the system so that the FQHN can be properly determined.
- 文件

`/etc/hosts`Host name database

`/etc/nsswitch.conf`Name service switch configuration file

`/etc/resolv.conf`Configuration file for name server routines
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/smtp/sendmail
Interface Stability	Committed

另请参见

[domainname\(1M\)](#), [sendmail\(1M\)](#), [hosts\(4\)](#), [attributes\(5\)](#)

引用名	check-permissions – check permissions on mail rerouting files
用法概要	<code>/usr/sbin/check-permissions [login]</code>
描述	The <code>check-permissions</code> script is intended as a migration aid for sendmail(1M). It checks the <code>/etc/mail/sendmail.cf</code> file for all configured alias files, and checks the alias files for <code>:include:</code> files. It also checks for certain <code>.forward</code> files. For each file that <code>check-permissions</code> checks, it verifies that none of the parent directories are group- or world-writable. If any directories are overly permissive, it is reported. Otherwise it reports that no unsafe directories were found. As to which <code>.forward</code> files are checked, it depends on the arguments included on the command line. If no argument is given, the current user's home directory is checked for the presence of a <code>.forward</code> file. If any arguments are given, they are assumed to be valid logins, and the home directory of each one is checked. If the special argument <code>ALL</code> is given, the <code>passwd</code> entry in the <code>/etc/nsswitch.conf</code> file is checked, and all password entries that can be obtained through the switch file are checked. In large domains, this can be time-consuming.
操作数	The following operands are supported: <i>login</i> Where <i>login</i> is a valid user name, checks the home directory for <i>login</i>. <code>ALL</code> Checks the home directory of <i>all</i> users.
文件	<code>/etc/mail/sendmail.cf</code> Defines environment for <code>sendmail</code> <code>/etc/mail/aliases</code> Ascii mail aliases file
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/smtp/sendmail
Interface Stability	Committed

另请参见 [getent\(1M\)](#), [sendmail\(1M\)](#), [aliases\(4\)](#), [attributes\(5\)](#)

引用名	chk_encodings – check the label encodings file syntax
用法概要	<code>/usr/sbin/chk_encodings [-a] [-c <i>maxclass</i>] [<i>pathname</i>]</code>
描述	chk_encodings checks the syntax of the label-encodings file that is specified by <i>pathname</i>. With the -a option, chk_encodings also prints a semantic analysis of the label-encodings file that is specified by <i>pathname</i>. If <i>pathname</i> is not specified, chk_encodings checks and analyzes the <code>/etc/security/tsol/label_encodings</code> file. If label-encodings file analysis was requested, whatever analysis can be provided is written to the standard output file even if errors were found.
选项	-a Provide a semantic analysis of the label encodings file. -c <i>maxclass</i> Accept a maximum classification value of <i>maxclass</i> (default 255) in the label encodings file CLASSIFICATIONS section.
退出状态	When successful, chk_encodings returns an exit status of 0 (true) and writes to the standard output file a confirmation that no errors were found in <i>pathname</i> . Otherwise, chk_encodings returns an exit status of nonzero (false) and writes an error diagnostic to the standard output file.
文件	<code>/etc/security/tsol/label_encodings</code> The label encodings file contains the classification names, words, constraints, and values for the defined labels of this system.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/trusted
Interface Stability	See below.
Standard	DDS-2600-6216-93, 《Compartmented Mode Workstation Labeling: Encodings Format》, September 1993

The command output is Not-an-Interface. The command invocation is Committed for systems that implement the DIA MAC policy.

另请参见 [label_encodings\(4\)](#), [attributes\(5\)](#), [labels\(5\)](#)

《Trusted Extensions Label Administration》中的“[How to Analyze and Verify the label_encodings File](#)”

附注

The functionality described on this manual page is available only if the system is configured with Trusted Extensions.

This file is part of the Defense Intelligence Agency (DIA) Mandatory Access Control (MAC) policy. This file might not be applicable to other MAC policies that might be developed for future releases of Solaris Trusted Extensions software.

引用名

chroot – change root directory for a command

用法概要

```
/usr/sbin/chroot newroot command
```

描述

The `chroot` utility causes *command* to be executed relative to *newroot*. The meaning of any initial slashes (`/`) in the path names is changed to *newroot* for *command* and any of its child processes. Upon execution, the initial working directory is *newroot*.

Notice that redirecting the output of *command* to a file,

```
chroot newroot command >x
```

will create the file `x` relative to the original root of *command*, not the new one.

The new root path name is always relative to the current root. Even if a `chroot` is currently in effect, the *newroot* argument is relative to the current root of the running process.

This command can be run only by the super-user.

返回值

The exit status of `chroot` is the return value of *command*.

示例

示例 1 Using the `chroot` Utility

The `chroot` utility provides an easy way to extract tar files (see [tar\(1\)](#)) written with absolute filenames to a different location. It is necessary to copy the shared libraries used by tar (see [ldd\(1\)](#)) to the *newroot* filesystem.

```
example# mkdir /tmp/lib; cd /lib
example# cp ld.so.1 libc.so.1 libcmd.so.1 libdl.so.1 \
libsec.so.1 /tmp/lib
example# cp /usr/bin/tar /tmp
example# dd if=/dev/rmt/0 | chroot /tmp tar xvf -
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[cd\(1\)](#), [tar\(1\)](#), [chroot\(2\)](#), [ttyname\(3C\)](#), [attributes\(5\)](#)

附注

Exercise extreme caution when referencing device files in the new root file system.

References by routines such as [ttyname\(3C\)](#) to `stdin`, `stdout`, and `stderr` will find that the device associated with the file descriptor is unknown after `chroot` is run.

引用名	cimworkshop – start the Sun WBEM CIM WorkShop application										
用法概要	<code>/usr/sadm/bin/cimworkshop</code>										
描述	The <code>cimworkshop</code> command starts Sun WBEM CIM WorkShop, a graphical user interface that enables you to create, modify, and view the classes and instances that describe the managed resources on your system. Managed resources are described using a standard information model called Common Information Model (CIM). A CIM class is a computer representation, or model, of a type of managed resource, such as a printer, disk drive, or CPU. A CIM instance is a particular managed resource that belongs to a particular class. Instances contain actual data. Objects can be shared by any WBEM-enabled system, device, or application. CIM objects are grouped into meaningful collections called schema. One or more schemas can be stored in directory-like structures called namespaces. The CIM WorkShop application displays a Login dialog box. Context help is displayed on the left side of the CIM WorkShop dialog boxes. When you click on a field, the help content changes to describe the selected field. By default, CIM WorkShop uses the RMI protocol to connect to the CIM Object Manager on the local host, in the default namespace, <code>root\cimv2</code>. You can select HTTP if you want to communicate to a CIM Object Manager using the standard XML/HTTP protocol from the Desktop Management Task Force. When a connection is established, all classes contained in the default namespace are displayed in the left side of the CIM WorkShop window. The name of the current namespace is listed in the tool bar. All programming operations are performed within a namespace. Four namespaces are created in a root namespace during installation: <table><tr><td><code>cimv2</code></td><td>Contains the default CIM classes that represent managed resources on your system.</td></tr><tr><td><code>security</code></td><td>Contains the security classes used by the CIM Object Manager to represent access rights for users and namespaces.</td></tr><tr><td><code>system</code></td><td>Contains properties for configuring the CIM Object Manager.</td></tr><tr><td><code>snmp</code></td><td>Contains pre-defined SNMP-related classes and all SNMP MOF files that are compiled.</td></tr></table> The <code>cimworkshop</code> application allows you to perform the following tasks: <table><tr><td>Create, view, and change namespaces.</td><td>Use the CIM WorkShop application to view all namespaces. A namespace is a directory-like structure that can store CIM classes and instances.</td></tr></table>	<code>cimv2</code>	Contains the default CIM classes that represent managed resources on your system.	<code>security</code>	Contains the security classes used by the CIM Object Manager to represent access rights for users and namespaces.	<code>system</code>	Contains properties for configuring the CIM Object Manager.	<code>snmp</code>	Contains pre-defined SNMP-related classes and all SNMP MOF files that are compiled.	Create, view, and change namespaces.	Use the CIM WorkShop application to view all namespaces. A namespace is a directory-like structure that can store CIM classes and instances.
<code>cimv2</code>	Contains the default CIM classes that represent managed resources on your system.										
<code>security</code>	Contains the security classes used by the CIM Object Manager to represent access rights for users and namespaces.										
<code>system</code>	Contains properties for configuring the CIM Object Manager.										
<code>snmp</code>	Contains pre-defined SNMP-related classes and all SNMP MOF files that are compiled.										
Create, view, and change namespaces.	Use the CIM WorkShop application to view all namespaces. A namespace is a directory-like structure that can store CIM classes and instances.										

- Create, delete, and view CIM classes.

You cannot modify the unique attributes of the classes that make up the CIM and Solaris Schema. You can create a new instance or subclass of the class and modify the desired attributes in that instance or subclass.
- Create, modify, delete, and view CIM instances.

You can add instances to a class and modify its inherited properties or create new properties. You can also change the property values of a CIM instance.
- Invoke methods.

You can set input values for a parameter of a method and invoke the method.

When CIM WorkShop connects to the CIM Object Manager in a particular namespace, all subsequent operations occur within that namespace. When you connect to a namespace, you can access the classes and instances in that namespace (if they exist) and in any namespaces contained in that namespace.

When you use CIM WorkShop to view CIM data, the WBEM system validates your login information on the current host. By default, a validated WBEM user is granted read access to the CIM Schema. The CIM Schema describes managed objects on your system in a standard format that all WBEM-enabled systems and applications can interpret.

- Read Only

Allows read-only access to CIM Schema objects. Users with this privilege can retrieve instances and classes, but cannot create, delete, or modify CIM objects.
- Read/Write

Allows full read, write, and delete access to all CIM classes and instances.
- Write

Allows write and delete, but not read access to all CIM classes and instances.
- None

Allows no access to CIM classes and instances.

用法

The `cimworkshop` command is not a tool for a distributed environment. Rather, this command is used for local administration on the machine on which the CIM Object Manager is running.

退出状态

The `cimworkshop` utility terminates with exit status 0.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWwbdev

另请参见

[mofcomp\(1M\)](#), [wbemlogviewer\(1M\)](#), [init.wbem\(1M\)](#), [attributes\(5\)](#)

引用名	clear_locks – clear locks held on behalf of an NFS client
用法概要	<code>/usr/sbin/clear_locks [-s] <i>hostname</i></code>
描述	The <code>clear_locks</code> command removes all file, record, and share locks created by the <i>hostname</i> and held on the current host, regardless of which process created or owns the locks. This command can be run only by the super-user. This command should only be used to repair the rare case of a client crashing and failing to clear held locks. Clearing locks held by an active client may cause applications to fail in an unexpected manner.
选项	<code>-s</code> Remove all locks created by the current machine and held by the server <i>hostname</i> .
操作数	The following operands are supported: <i>hostname</i> name of host server
退出状态	0 Successful operation. 1 If not root. 2 Usage error. 3 If unable to contact server (RPC).
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [fcntl\(2\)](#), [attributes\(5\)](#)

引用名	clinfo – display cluster information
用法概要	clinfo [-nh]
描述	The <code>clinfo</code> command displays cluster configuration information about the node from which the command is executed. Without arguments, <code>clinfo</code> returns an exit status of 0 if the node is configured and booted as part of a cluster. Otherwise, <code>clinfo</code> returns an exit status of 1.
选项	The following options are supported: -h Displays the highest node number allowed to be configured. This is different from the maximum number of nodes supported in a given cluster. The current highest configured node number can change immediately after the command returns since new nodes can be dynamically added to a running cluster. For example, <code>clinfo -h</code> might return 64, meaning that the highest number you can use to identify a node is 64. See the Sun Cluster 3.0 System Administration Guide for a description of utilities you can use to determine the number of nodes in a cluster. -n Prints the number of the node from which <code>clinfo</code> is executed.
退出状态	The following exit values are returned: 0 Successful completion. 1 An error occurred. This is usually because the node is not configured or booted as part of a cluster.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [attributes\(5\)](#)

引用名

clri, dcopy – clear inode

用法概要

```
clri [-F FSType] [-V] special i-number
dcopy [-F FSType] [-V] special i-number
```

描述

clri writes zeros on the inodes with the decimal *i-number* on the file system stored on *special*. After clri, any blocks in the affected file show up as missing in an [fsck\(1M\)](#) of *special*.

Read and write permission is required on the specified file system device. The inode becomes allocatable.

The primary purpose of this routine is to remove a file that for some reason appears in no directory. If it is used to zap an inode that does appear in a directory, care should be taken to track down the entry and remove it. Otherwise, when the inode is reallocated to some new file, the old entry will still point to that file. At that point, removing the old entry will destroy the new file. The new entry will again point to an unallocated inode, so the whole cycle is likely to be repeated again and again.

dcopy is a symbolic link to clri.

选项

-F *FSType*

Specify the *FSType* on which to operate. The *FSType* should either be specified here or be determinable from `/etc/vfstab` by matching *special* with an entry in the table, or by consulting `/etc/default/fs`.

-V

Echo the complete command line, but do not execute the command. The command line is generated by using the options and arguments provided by the user and adding to them information derived from `/etc/vfstab`. This option should be used to verify and validate the command line.

用法

See [largefile\(5\)](#) for the description of the behavior of clri and dcopy when encountering files greater than or equal to 2 Gbyte (2³¹ bytes).

文件

`/etc/default/fs`Default local file system type

`/etc/vfstab`List of default parameters for each file system

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[fsck\(1M\)](#), [vfstab\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

附注

This command might not be supported for all *FSTypes*.

引用名

consadm – select or display devices used as auxiliary console devices

用法概要

```
/usr/sbin/consadm
/usr/sbin/consadm [-a device. . .] [-p]
/usr/sbin/consadm [-d device. . .] [-p]
/usr/sbin/consadm [-p]
```

描述

consadm selects the hardware *device* or devices to be used as auxiliary console devices, or displays the current device. Only superusers are allowed to make or display auxiliary console device selections.

Auxiliary console devices receive copies of console messages, and can be used as the console during single user mode. In particular, they receive kernel messages and messages directed to /dev/sysmsg. On Solaris x86 based systems they can also be used for interaction with the bootstrap.

By default, selecting a display device to be used as an auxiliary console device selects that device for the duration the system remains up. If the administrator needs the selection to persist across reboots the -p option can be specified.

consadm runs a daemon in the background, monitoring auxiliary console devices. Any devices that are disconnected (hang up, lose carrier) are removed from the auxiliary console device list, though not from the persistent list. While auxiliary console devices may have been removed from the device list receiving copies of console messages, those messages will always continue to be displayed by the default console device.

The daemon will not run if it finds there are not any auxiliary devices configured to monitor. Likewise, after the last auxiliary console is removed, the daemon will shut itself down. Therefore the daemon persists for only as long as auxiliary console devices remain active.

See [eeprom\(1M\)](#) for instructions on assigning an auxiliary console device as the system console.

选项

The following options are supported:

- a *device* Adds *device* to the list of auxiliary console devices. Specify *device* as the path name to the device or devices to be added to the auxiliary console device list.
 - d *device* Removes *device* from the list of auxiliary console devices. Specify *device* as the path name to the device or devices to be removed from the auxiliary console device list.
 - p Prints the list of auxiliary consoles that will be auxiliary across reboots.
- When invoked with the -a or -d options, tells the application to make the change persist across reboot.

示例

示例 1 Adding to the list of devices that will receive console messages

The following command adds /dev/term/a to the list of devices that will receive console messages.

```
example# consadm -a /dev/term/a
```

示例 2 Removing from the list of devices that will receive console messages

The following command removes /dev/term/a from the list of devices that will receive console messages. This includes removal from the persistent list.

```
example# consadm -d -p /dev/term/a
```

示例 3 Printing the list of devices selected as auxiliary console devices

The following command prints the name or names of the device or devices currently selected as auxiliary console devices.

```
example# consadm
```

环境变量

See [environ\(5\)](#) for descriptions of the following environment variables that affect the execution of consadm: LC_CTYPE, LC_MESSAGES, and NLSPATH.

退出状态

The following exit values are returned:

0 Successful completion.

>0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[kmdb\(1\)](#), [svcs\(1\)](#), [eeprom\(1M\)](#), [svcadm\(1M\)](#), [syslogd\(1M\)](#), [environ\(5\)](#), [attributes\(5\)](#), [smf\(5\)](#), [sysmsg\(7D\)](#), [console\(7D\)](#)

附注

Auxiliary console devices are not usable for kmdb or firmware I/O, do not receive panic messages, and do not receive output directed to /dev/console.

The consadm service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/system/consadm
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名 console-reset – force console login to display

用法概要 `svc:/system/console-reset`

描述 The console- reset service runs if X is not running on console and the console is ready for interactive use. Under these conditions, console- reset turns off the boot-time graphics, forcing the console login prompt to display.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Uncommitted

另请参见 [svcs\(1\)](#), [svcadm\(1M\)](#), [vbiosd\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The console- reset service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/system/console-reset:default`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名 conv_lp – convert LP configuration

用法概要 conv_lp [-d *dir*] [-f *file*]

描述 conv_lp reads LP printer configuration information from a directory and converts it to an output file for use with print client software.

选项 The following options are supported:

- d*dir* The root (/) directory from which LP configuration information is read. The default is root (/).
- f*file* The output file to which conv_lp writes the converted LP configuration information.

退出状态 The following exit values are returned:

- 0 Successful completion.
- non-zero An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	print/lp/print-client-commands

另请参见 [attributes\(5\)](#)

引用名	conv_lpd – convert LPD configuration
用法概要	conv_lpd [-c printcap] [-n] <i>file</i>
描述	conv_lpd converts LPD printer configuration information from <i>file</i> to a printcap file. <i>file</i> specifies the name of the input file, and is in printcap format.
选项	The following options are supported: -c printcap Specifies the type of output file produced by the conversion. -c printcap converts to a printcap file. -n Preserves the namelist during the conversion.
操作数	The following operands are supported: <i>file</i> The file to be converted.
退出状态	The following exit values are returned: 0 Successful completion. non-zero An error occurred.
文件	/etc/printcap SunOS 4.x printer capability database.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	print/lp/print-client-commands

另请参见 [lpset\(1M\)](#), [attributes\(5\)](#)

引用名	coreadm – 核心文件管理
用法概要	<pre>coreadm [-g pattern] [-G content] [-i pattern] [-I content] [-d option]... [-e option]... coreadm [-p pattern] [-P content] [pid]...</pre>
描述	coreadm 指定异常终止的进程生成的核心文件的名称和位置。请参见 core(4)。 仅属于“维护和修复”RBAC 配置文件的用户和角色可以执行第一种格式的 SYNOPSIS（用法概要）。此用法概要可以配置系统范围的核心文件选项，包括全局核心文件名称模式和 init(1M) 进程的核心文件名称模式。所有设置都会永久保存，并在引导时应用。 非特权用户可以执行第二种格式的 SYNOPSIS（用法概要）。此语法格式可指定操作系统用于生成每进程核心文件的文件名称模式和核心文件内容。 核心文件名称模式是包含以前导字符 % 指定的嵌入变量的常规文件系统路径名。变量基于操作系统生成核心文件时的有效值进行扩展。可能的嵌入变量如下所示： %d 可执行文件目录名，最多包含 MAXPATHLEN 个字符 %f 可执行文件名，最多包含 MAXCOMLEN 个字符 %g 有效组 ID %m 机器名 (uname -m) %n 系统节点名 (uname -n) %p 进程 ID %t time(2) 的十进制值 %u 有效用户 ID %z 执行的进程所在的区域名称 (zonename) %% 字面值 % 例如，核心文件名称模式 <code>/var/cores/core.%f.%p</code> 对于进程 ID 为 1234 的 <code>foo</code> 命令，将生成核心文件名称 <code>/var/cores/core.foo.1234</code>。

通过使用一系列标记来标识进程的二进制图像的各个部分，可以指定核心文件内容描述。

anon

匿名专用映射，包括非主线程栈的线程栈

ctf

装入的对象文件的 CTF 类型信息部分

data

可写的专用文件映射

dism

DISM 映射

heap

进程堆

ism

ISM 映射

rodata

只读专用文件映射

shanon

匿名共享映射

shfile

文件支持的共享映射

shm

System V 共享内存

stack

进程栈

syntab

装入的对象文件的符号表部分

text

可读且可执行的专用文件映射

此外，可以使用标记 **all** 来表示核心文件应该包含进程的二进制图像的所有这些部分。可以使用标记 **none** 来指示将不包含任何映射。**default** 标记指示包含系统缺省内容 (**stack+heap+shm+ism+dism+text+data+rodata+anon+shanon+ctf+syntab**)。无论映射内容是什么，核心文件中始终存在 **/proc** 文件系统数据结构。

可以使用 **+** 和 **-** 来串联标记。例如，核心文件内容 **default-ism** 将生成包含缺省映射集但不含任何锁定共享内存映射的核心文件。

不带参数的 **coreadm** 命令将报告当前系统配置，例如：

```
$ coreadm
  global core file pattern: /var/cores/core.%f.%p
  global core file content: all
  init core file pattern: core
  init core file content: default
    global core dumps: enabled
    per-process core dumps: enabled
  global setid core dumps: enabled
per-process setid core dumps: disabled
  global core dump logging: disabled
```

仅带进程 ID 列表的 `coreadm` 命令将报告每个进程的每进程核心文件名称模式，例如：

```
$ coreadm 278 5678
  278:   core.%f.%p default
  5678:  /home/george/cores/%f.%p.%t all-ism
```

只有进程所有者或具有 `proc_owner` 特权的用户才能以这种方式查询进程。

当进程进行核心转储时，最多会生成三个核心文件：一个在每进程的位置，一个在系统范围的全局位置，还有一个在运行的进程所在区域的全局位置（如果该进程在本地区域而非全局区域中运行）。将根据对应位置的有效选项生成每个核心文件。

生成全局核心文件后，该文件将具有 `600` 模式，并由超级用户拥有。非特权用户不能检查此类文件。

将通过进程凭证在模式 `600` 下创建普通的每进程核心文件。进程所有者可以检查此类文件。

一直是或者自从上一次 `exec(2)` 之后成为 `setuid` 或 `setgid` 的进程存在与核心转储相关的安全问题。同样，最初具有超级用户特权但因 `setuid(2)` 失去这些特权的进程也存在与核心转储相关的安全问题。上述任何一种类型的进程都可以在该进程的当前非特权所有者无法访问的地址空间中包含敏感信息。如果启用 `setid` 核心文件，则生成的文件将处于模式 `600` 并由超级用户拥有。

选项

所支持的选项如下：

`-d option...`

禁用指定的核心文件选项。有关可能的选项的描述，请参见 `-e option`。

可以在命令行上指定多个 `-e` 和 `-d` 选项。只有属于“维护和修复”RBAC 配置文件的用户和角色才能使用此选项。

`-e option...`

启用指定的核心文件选项。将 `option` 指定为以下类型之一：

`global`

允许使用全局核心模式的核心转储。

`global-setid`

允许使用全局核心模式的 `set-id` 核心转储。

log

尝试生成全局核心文件时，生成 `syslog(3C)` 消息。

process

允许使用每进程核心模式的核心转储。

proc-setid

允许使用每进程核心模式的 `set-id` 核心转储。

可以在命令行上指定多个 `-e` 和 `-d` 选项。只有属于“维护和修复”RBAC 配置文件的用户和角色才能使用此选项。

-g *pattern*

将全局核心文件名称模式设置为 *pattern*。模式必须以 `/` 开头，可以包含描述中介绍的任何特殊 `%` 变量。

只有属于“维护和修复”RBAC 配置文件的用户和角色才能使用此选项。

-G *content*

将全局核心文件内容设置为 *content*。必须使用描述中介绍的标记来指定内容。

只有属于“维护和修复”RBAC 配置文件的用户和角色才能使用此选项。

-i *pattern*

将缺省的每进程核心文件名称设置为 *pattern*。此选项更改每进程模式仍设置为缺省值的任何进程的每进程模式。已设置了每进程模式的进程或从一个已设置每进程模式的进程派生的进程（使用 `-p` 选项）不受影响。重新引导后仍会保持此缺省值。

只有属于“维护和修复”RBAC 配置文件的用户和角色才能使用此选项。

-I *content*

将缺省的每进程核心文件内容设置为 *content*。此选项更改每进程内容仍设置为缺省值的任何进程的每进程内容。已设置了每进程内容的进程或从一个已设置每进程内容的进程派生的进程（使用 `-p` 选项）不受影响。重新引导后仍会保持此缺省值。

只有属于“维护和修复”RBAC 配置文件的用户和角色才能使用此选项。

-p *pattern*

针对每个指定的进程 ID，将每进程核心文件名称模式设置为 *pattern*。模式可以包含描述中介绍的任何特殊 `%` 变量，无需以 `/` 开头。如果模式不是以 `/` 开头，则会相对于进程生成核心文件时使用的目录对其进行评估。

非特权用户只能将 `-p` 选项应用于该用户所拥有的进程。具有 `proc_owner` 特权的用户可以将该选项应用于任何进程。受影响进程的未来子进程将继承每进程核心文件名称模式。请参见 `fork(2)`。

如果没有指定进程 ID，则 `-p` 选项将在父进程（通常是运行 `coreadm` 的 `shell`）上将每进程核心文件名称模式设置为 *pattern*。

-P *content*

针对每个指定的进程 ID，将每进程核心文件内容设置为 *content*。必须使用**描述**中介绍的标记来指定内容。

非特权用户只能将 **-p** 选项应用于该用户所拥有的进程。具有 **proc_owner** 特权的用户可以将该选项应用于任何进程。受影响进程的未来子进程将继承每进程核心文件名称模式。请参见 **fork(2)**。

如果没有指定进程 ID，则 **-P** 选项将在父进程（通常是运行 **coreadm** 的 **shell**）上将每进程文件内容设置为 *content*。

操作数

支持下列操作数：

pid
进程 ID

示例

示例 1 设置核心文件名称模式

从用户的 **\$HOME/.profile** 或 **\$HOME/.login** 执行时，以下命令将为登录会话期间运行的所有进程设置核心文件名称模式：

```
example$ coreadm -p core.%f.%p
```

请注意，由于省略了进程 ID，因此将在当前运行的 **shell** 中设置每进程核心文件名称模式，且该模式将被所有子进程继承。

示例 2 将用户文件转储至子目录

以下命令将用户的所有核心转储转储至起始目录的 **corefiles** 子目录中，通过系统节点名称进行区分。该命令对于使用多个不同机器但具有一个共享起始目录的用户非常有用。

```
example$ coreadm -p $HOME/corefiles/%n.%f.%p 1234
```

示例 3 选择全局核心文件系统信息库

以下命令将系统设置为仅当从 **/usr/bin** 或 **/usr/sbin** 运行可执行文件时，才在全局系统信息库中生成核心文件。

```
example# mkdir -p /var/cores/usr/bin
example# mkdir -p /var/cores/usr/sbin
example# coreadm -G all -g /var/cores/%d/%f.%p.%n
```

文件

/var/cores
为全局核心文件存储提供的目录。

退出状态

将返回下列退出值：

0
成功完成。

- 1 获取或修改系统核心文件配置时发生致命错误。
- 2 指定的命令行选项无效。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见 [gcore\(1\)](#)、[pfexec\(1\)](#)、[svcs\(1\)](#)、[init\(1M\)](#)、[svcadm\(1M\)](#)、[exec\(2\)](#)、[fork\(2\)](#)、[setuid\(2\)](#)、[time\(2\)](#)

附注 在本地（非全局）区域中，全局设置应用于该区域中运行的进程。此外，全局区域设置还应用于其他任何区域中运行的进程。

术语**全局设置**指的是应用于系统或区域整体的设置，但并不一定表示这些设置对全局区域有效。

coreadm 服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

```
svc:/system/coreadm:default
```

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

已经获得 `solaris.smf.manage.coreadm` 和 `solaris.smf.value.coreadm` 授权的用户、角色或配置文件也可以使用 `-g`、`-G`、`-i`、`-I`、`-e` 和 `-d` 选项。

引用名

cpustat – monitor system behavior using CPU performance counters

用法概要

```
cpustat -c eventspec [-c eventspec]... [-p period] [-T u | d ]
      [-sntD] [interval [count]]

cpustat -h
```

描述

The cpustat utility allows CPU performance counters to be used to monitor the overall behavior of the CPUs in the system.

If *interval* is specified, cpustat samples activity every *interval* seconds, repeating forever. If a *count* is specified, the statistics are repeated *count* times. If neither are specified, an interval of five seconds is used, and there is no limit to the number of samples that are taken.

选项

The following options are supported:

-c *eventspec*

Specifies a set of events for the CPU performance counters to monitor. The syntax of these event specifications is:

```
[picn=]eventn[,attr[n][=val]][, [picn=]eventn
      [,attr[n][=val]],...]
```

You can use the -h option to obtain a list of available events and attributes. This causes generation of the usage message. You can omit an explicit counter assignment, in which case cpustat attempts to choose a capable counter automatically.

Attribute values can be expressed in hexadecimal, octal, or decimal notation, in a format suitable for [strtoll\(3C\)](#). An attribute present in the event specification without an explicit value receives a default value of 1. An attribute without a corresponding counter number is applied to all counters in the specification.

The semantics of these event specifications can be determined by reading the CPU manufacturer's documentation for the events.

Multiple -c options can be specified, in which case the command cycles between the different event settings on each sample.

-D

Enables debug mode.

-h

Prints an extensive help message on how to use the utility and how to program the processor-dependent counters.

-n

Omits all header output (useful if cpustat is the beginning of a pipeline).

-p *period*

Causes cpustat to cycle through the list of *eventspecs* every *period* seconds. The tool sleeps after each cycle until *period* seconds have elapsed since the first *eventspec* was measured.

When this option is present, the optional *count* parameter specifies the number of total cycles to make (instead of the number of total samples to take). If *period* is less than the number of *eventspecs* times *interval*, the tool acts as if *period* is 0.

-s

Creates an idle soaker thread to spin while system-only *eventspecs* are bound. One idle soaker thread is bound to each CPU in the current processor set. System-only *eventspecs* contain both the *nouser* and the *sys* tokens and measure events that occur while the CPU is operating in privileged mode. This option prevents the kernel's idle loop from running and triggering system-mode events.

-T u | d

Display a time stamp.

Specify *u* for a printed representation of the internal representation of time. See [time\(2\)](#).
Specify *d* for standard date format. See [date\(1\)](#).

-t

Prints an additional column of processor cycle counts, if available on the current architecture.

用法

A closely related utility, [cputrack\(1\)](#), can be used to monitor the behavior of individual applications with little or no interference from other activities on the system.

The `cpustat` utility must be run by the super-user, as there is an intrinsic conflict between the use of the CPU performance counters system-wide by `cpustat` and the use of the CPU performance counters to monitor an individual process (for example, by `cputrack`.)

Once any instance of this utility has started, no further per-process or per-LWP use of the counters is allowed until the last instance of the utility terminates.

The times printed by the command correspond to the wallclock time when the hardware counters were actually sampled, instead of when the program told the kernel to sample them. The time is derived from the same timebase as [gethrtime\(3C\)](#).

The processor cycle counts enabled by the `-t` option always apply to both user and system modes, regardless of the settings applied to the performance counter registers.

On some hardware platforms running in system mode using the “*sys*” token, the counters are implemented using 32-bit registers. While the kernel attempts to catch all overflows to synthesize 64-bit counters, because of hardware implementation restrictions, overflows can be lost unless the sampling interval is kept short enough. The events most prone to wrap are those that count processor clock cycles. If such an event is of interest, sampling should occur frequently so that less than 4 billion clock cycles can occur between samples.

The output of `cpustat` is designed to be readily parseable by [nawk\(1\)](#) and [perl\(1\)](#), thereby allowing performance tools to be composed by embedding `cpustat` in scripts. Alternatively, tools can be constructed directly using the same APIs that `cpustat` is built upon using the facilities of [libcpc\(3LIB\)](#). See [cpc\(3CPC\)](#).

The cpustat utility only monitors the CPUs that are accessible to it in the current processor set. Thus, several instances of the utility can be running on the CPUs in different processor sets. See [prset\(1M\)](#) for more information about processor sets.

Because cpustat uses LWPs bound to CPUs, the utility might have to be terminated before the configuration of the relevant processor can be changed.

示例

SPARC

示例 1 Measuring External Cache References and Misses

The following example measures misses and references in the external cache. These occur while the processor is operating in user mode on an UltraSPARC machine.

```
example% cpustat -c EC_ref,EC_misses 1 3
```

time	cpu	event	pic0	pic1
1.008	0	tick	69284	1647
1.008	1	tick	43284	1175
2.008	0	tick	179576	1834
2.008	1	tick	202022	12046
3.008	0	tick	93262	384
3.008	1	tick	63649	1118
3.008	2	total	651077	18204

x86

示例 2 Measuring Branch Prediction Success on Pentium 4

The following example measures branch mispredictions and total branch instructions in user and system mode on a Pentium 4 machine.

```
example% cpustat -c \  
pic12=branch_retired,emask12=0x4,pic14=branch_retired,\  
emask14=0xf,sys 1 3
```

time	cpu	event	pic12	pic14
1.010	1	tick	458	684
1.010	0	tick	305	511
2.010	0	tick	181	269
2.010	1	tick	469	684
3.010	0	tick	182	269
3.010	1	tick	468	684
3.010	2	total	2063	3101

示例 3 Counting Memory Accesses on Opteron

The following example determines the number of memory accesses made through each memory controller on an Opteron, broken down by internal memory latency:

```
cpustat -c \  
pic0=NB_mem_ctrlr_page_access,umask0=0x01, \  
pic1=NB_mem_ctrlr_page_access,umask1=0x02, \  
sys 1 3
```

示例 3 Counting Memory Accesses on Opteron (续)

```
pic2=NB_mem_ctrlr_page_access,umask2=0x04,sys \
1

time cpu event      pic0      pic1      pic2
1.003 0 tick      41976     53519     7720
1.003 1 tick      5589      19402      731
2.003 1 tick      6011      17005      658
2.003 0 tick      43944     45473     7338
3.003 1 tick      7105      20177      762
3.003 0 tick      47045     48025     7119
4.003 0 tick      43224     46296     6694
4.003 1 tick      5366      19114      652
```

警告

By running the cpustat command, the super-user forcibly invalidates all existing performance counter context. This can in turn cause all invocations of the cputrack command, and other users of performance counter context, to exit prematurely with unspecified errors.

If cpustat is invoked on a system that has CPU performance counters which are not supported by Solaris, the following message appears:

```
cpustat: cannot access performance counters - Operation not applicable
```

This error message implies that cpc_open() has failed and is documented in [cpc_open\(3CPC\)](#). Review this documentation for more information about the problem and possible solutions.

If a short interval is requested, cpustat might not be able to keep up with the desired sample rate. In this case, some samples might be dropped.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	diagnostic/cpu-counters
Interface Stability	Committed

另请参见

[cputrack\(1\)](#), [nawk\(1\)](#), [perl\(1\)](#), [iostat\(1M\)](#), [prstat\(1M\)](#), [psrset\(1M\)](#), [vmstat\(1M\)](#), [cpc\(3CPC\)](#), [cpc_open\(3CPC\)](#), [cpc_bind_cpu\(3CPC\)](#), [gethrtime\(3C\)](#), [strtoll\(3C\)](#), [libcpc\(3LIB\)](#), [attributes\(5\)](#)

附注

When cpustat is run on a Pentium 4 with HyperThreading enabled, a CPC set is bound to only one logical CPU of each physical CPU. See [cpc_bind_cpu\(3CPC\)](#).

引用名

croinfo, diskinfo – query and display information about 1) chassis, receptacle, and occupants or 2) disk occupants of bay receptacles

用法概要

```
croinfo [-h] [-v] [-I cro_db] [-o fields] [-O fields]
        [-P product-id] [-C chassis-id] [-A alias-id]
        [-R receptacle-name] [-T receptacle-type] [-t occupant-type]
        [-D devchassis-path] [-d occupant-devices] [-p occupant-paths]
        [-c occupant-compdev] [-i occupant-devid] [-m occupant-mfg]
        [-e occupant-model] [-n occupant-part] [-s occupant-serial]
        [-f occupant-firm] [-1 occupant-misc-1] [-2 occupant-misc-2]
        [-3 occupant-misc-3]

diskinfo [-h] [-v] [-I cro_db] [-o fields] [-O fields]
        [-P product-id] [-C chassis-id] [-A alias-id]
        [-R receptacle-name] [-T receptacle-type] [-t occupant-type]
        [-D devchassis-path] [-d occupant-devices] [-p occupant-paths]
        [-c occupant-compdev] [-i occupant-devid] [-m occupant-mfg]
        [-e occupant-model] [-n occupant-part] [-s occupant-serial]
        [-f occupant-firm] [-1 occupant-misc-1] [-2 occupant-misc-2]
        [-3 occupant-misc-3]

croinfo -?

diskinfo -?
```

描述

The diskinfo and croinfo utility share the same binary executable. At runtime, the utility checks to see how it was invoked, and adjusts defaults.

The croinfo utility allows users to query and display specific aspects of a system's configuration. Queries are performed against a record-oriented dataset that captures the relationship between physical location and various aspects of the device currently at that physical location. This relationship is expressed in terms of Chassis, Receptacle, and Occupant (thus the cro prefix).

Records in a CRO dataset are composed of multiple, named fields with each record having a potentially unique field value. An angle-bracket reference, such as *product-id*, is referring to a specific *field-name*. For a given record, a field value is either undefined (empty) or defined with a set of indexed string values. Some defined field values have just one string value, while others can have multiple string values.

Each *field-name* defined is associated with a separate *field-char* character. By convention, uppercase *field-char* characters are associated for chassis and receptacle information fields, and lowercase *field-char* characters are associated with occupant information fields. For each *field-char* character, a separate *-field-char field-name-RE* flag regular expression filter option is provided. This allows the user to customize queries to display information about specific aspects of the configuration. Records that match all regular expressions are selected, in dataset order, for display. For multiple string values, only one index value needs to match for the field to match.

By default, only a minimal number of default output fields are displayed. You can override the default with either `-o fields` for human-readable output or with `-O fields` for parseable output. In both cases, output fields can be specified using either the short-hand *field-char*[...] notation or in the more descriptive *field-name*[...] notation.

For human-readable use, by using the `-o fields` option, the user can override the default fields, and output any fields, in the specified order, in a column-aligned whitespace separated format. In general, output will be one line of output per matching record with undefined (empty) *field-name* values displayed as a hyphen (–). If, however, a displayed record has a multiple string *value* field, then multiple lines of output are produced with any secondary non-multiple string values fields showing a colon (:).

For scripting, by using the `-O fields` option, the user can override the default output fields and output any fields, in the specified order, in a parseable colon-separated format with whitespace removed and column headers suppressed. Output will be one line of output per matching record with undefined (empty) field-name values displayed with no value. If a displayed record *field-name* has a multiple string *value* field, then all the values are concatenated, separated by a comma. Any occurrence of a colon or a comma in a value is escaped with a leading backslash (\). To make scripts more legible, use of the `-O field-name`[...] notation is encouraged.

The `-o` and `-O` options are mutually exclusive.

If the `-h` option is used, or scripting output format is requested by using `-o`, the column headers for output fields are suppressed.

A Chassis is identified by a specific *product-id* and *chassis-id*. The *product-id* relates to a specific chassis-level product, like a system chassis Sun-Fire-X4200-M2 or a storage chassis SUN-Storage-J4410. For a given *product-id* value, the *chassis-id* defines a unique serial number.

A specific *product-id.chassis-id* combination can have an “managed” location-oriented *alias-id* defined by the administrator, using [fmadm\(1M\)](#) that provides installation-specific location information about where a chassis is physically located. This might include such information as building, room, rack, and U-number range within a rack.

In addition to the managed location-oriented *alias-id* defined using [fmadm\(1M\)](#), system chassis always have one well-known alias called SYS that can be used to identify receptacles that are internal to the system chassis.

Within a chassis, each receptacle has a unique *receptacle-name* that should match the physical silk-screen label designation for the receptacle. Each receptacle also has a *receptacle-type*, which helps define acceptable *occupant-types*.

When a receptacle is occupied, use the `-f` flag definitions for available *occupant-information*. Of particular interest is the `-c` *occupant-compdev* occupant information: it describes the common component of the public / dev name associated with the occupant device. For disks, this is the whole-disk `c#t#d#` name.

The CRO dataset order is associated with the *devchassis-path* of the record, which corresponds to the `/dev/chassis` name space maintained by [devchassisd\(1M\)](#). That ordering places records associated with the well-known SYS internal alias first, and records with BOOT receptacle-name first in SYS. This is done to ensure that, when applicable, information about the typical boot device is provided first.

For `croinfo`, the default output is in `-o Dtc` format, and all CRO records are shown.

For `diskinfo`, the default output is in `-o Dc` format, and a `-T bay receptacle-type` filter is applied. The meanings of the *occupant-misc-#* fields also take on a disk-specific interpretation: *misc-1* is capacity, and *misc-2* is target-port information. These defaults allow the `diskinfo` command to query the relationship between chassis, bay receptacles, and their *disk* occupants, while ignoring other CRO information.

选项

For each record *field-name* defined, a separate `-field-char field-name-RE` flag regular expression filter can be specified. For a given *field-name*, if no specific `-field-char field-name-RE` filter is defined, all CRO records match.

This allows the user to customize queries to display information about specific aspects of the configuration. CRO records that match all of the specified *field-name* regular expressions (as in [regex\(3C\)](#)) will be selected for display, with specific fields output controlled by means of `-o`, `-O`, or the default.

-P *product-id*

The *product-id* specifies the product identifier of an enumerated chassis. The *product-id* might be exposed in the `/dev/chassis` name space. For storage products that do not have an established [fmadm\(1M\)](#) managed *alias-id*, the *product-id* is visible in the [devchassis\(7FS\)](#) `/dev/chassis` name space.

Example system *product-id* value: Sun-Fire-X4200-M2

Example storage *product-id* value: SUN-Storage-J4410

-C *chassis-id*

The *chassis-id* specifies the serial number of a product chassis. The *chassis-id* might be exposed in the `/dev/chassis` name space. For storage products that do not have an established [fmadm\(1M\)](#) managed *alias-id*, the *product-id* is visible in the [devchassis\(7FS\)](#) `/dev/chassis` name space.

Example *chassis-id* value: 0818QAJ002

-A *alias-id*

An *alias-id* value can be the well-known alias value of SYS, for system internal devices. In addition, an *alias-id* value can be a managed alias, defined by the administrator using [fmadm\(1M\)](#). The intended use of a managed alias is to define the physical location of the *product-id.chassis-id*. The *alias-id* is exposed in the `/dev/chassis` name space.

Example well-known *alias-id* value: SYS

Example managed *alias-id* value: RACK29.U01-04

-R *receptacle-name*

For a specific *product-id*, the unique *receptacle-name* defines location of a specific receptacle in a chassis. The *receptacle-name* should be identical to a silk-screen label on the physical chassis, and should also match product documentation. The *receptacle-name* can have multiple path components, such as SYS/HD0. The receptacle-name is exposed in the /dev/chassis name space.

Example *receptacle-name* value: SYS/HD0

-T *receptacle-type*

Example *receptacle-type* value: bay

-t *occupant-type*

A receptacle without an occupant has an undefined (empty) *occupant-type* value, shown as a hyphen (-).

Example *occupant-type* value: disk

-D *devchassis-path*

Example *devchassis-path* value: /dev/chassis/SYS/HD0/disk

-d *occupant-devices*

A receptacle without an occupant has an undefined (empty) *occupant-devices* value, shown as a hyphen (-).

Example *occupant-devices* value:

/devices/scsi_vhci/disk@g5000c500101ba0a3

-p *occupant-paths*

A receptacle without an occupant has an undefined (empty) *occupant-paths* value, shown as a hyphen (-).

Example *occupant-paths* value:

devices/pci@0,0/pci10de,5d@pci11f8,8001@0/iport@f/disk@w5000c500101ba0a1,0

-c *occupant-compdev*

A receptacle without an occupant has an undefined (empty) *occupant-compdev* value, shown as a hyphen (-).

Example *occupant-compdev* value: c0t5000C500101BA0A3d0

-i *occupant-devid*

A receptacle without an occupant has an undefined (empty) *occupant-devid* value, shown as a hyphen (-).

Example *occupant-devid* value: id1,sd@n5000c500101ba0a3

-m *occupant-mfg*

A receptacle without an occupant has an undefined (empty) *occupant-mfg* value, shown as a hyphen (-).

Example *occupant-mfg* value: SEAGATE

-e *occupant-model*

A receptacle without an occupant has an undefined (empty) *occupant-model*, shown as a hyphen (-).

Example *occupant-model* value: ST32000SSSUN2.0T

-n *occupant-part*

A receptacle without an occupant has an undefined (empty) *occupant-part* value, shown as a hyphen (-).

Example *occupant-part* value: SEAGATE-ST32000SSSUN2.0T

-s *occupant-serial*

A receptacle without an occupant has an undefined (empty) *occupant-serial* value, shown as a hyphen (-).

Example *occupant-serial* value: 000949L09C8L_____9WM09C8L

-f *occupant-firm*

A receptacle without an occupant has an undefined (empty) *occupant-firm* value, shown as a hyphen (-).

Example *occupant-firm* value: 0313

-1 *occupant-misc-1*

A receptacle without an occupant has an undefined (empty) *occupant-misc-1* value, shown as a hyphen (-).

-2 *occupant-misc-2*

A receptacle without an occupant has an undefined (empty) *occupant-misc-2* value, shown as a hyphen (-).

-3 *occupant-misc-3*

A receptacle without an occupant has an undefined (empty) *occupant-misc-3* value, shown as a hyphen (-).

-?

Display usage information.

Output field-name
Control Options

-o *fields*

Output specified fields, in order, in human-readable format.

For *croinfo*, default output is in -o *Dtc* format. For *diskinfo*, default output is in -o *Dc* format.

-O *fields*

Output specified fields, in order, in parseable format.

-h

Do not output *field-name* column headers

-v

Display verbose header that includes various information about when the CRO dataset was created. This option is of particular use with -I and is used to specify a non-standard source for the CRO dataset.

**Dataset Selection
Option**

-I *cro_db*

Data file from which to obtain CRO dataset information.

示例

In some of the following examples, example output wraps in an 80-character-wide display.

示例 1 Determining Where a Disk is Located

The following command determines where a disk is located:

```
# croinfo -c c0t5000C500101BA0A3d0
D:devchassis-path                t:occupant-type
-----
/dev/chassis/RACK29.U01-04/DISK_00/disk  disk

c:occupant-compdev
-----
c0t5000C500101BA0A3d0
```

示例 2 Reporting Internal Disks

The following command reports the *receptacle-name* and the *occupant-compdev* of internal disks, that is, disks that are associated with the well-known SYS alias.

```
# diskinfo -A SYS -o Rc
R:receptacle-name  c:occupant-compdev
-----
SYS/HD0            c8t0d0
SYS/HD1            c8t1d0
SYS/HD2            -
SYS/HD3            -
```

Note that the SYS/HD2 and SYS/HD3 receptacles are empty.

The same command, in scripting output mode, would produce:

```
# diskinfo -A SYS -O receptacle-name,occupant-compdev
SYS/HD0:c8t0d0
SYS/HD1:c8t1d0
SYS/HD2:
SYS/HD3:
```

示例 3 Reporting Disks in a Specific Enclosure

The following command reports all the disks within a chassis with a specific *product-id* and *chassis-id* value.

```
# diskinfo -P SUN-Storage-J4410 -C SUN-Storage-J4410 -o Rc
R:receptacle-name      c:occupant-compdev
-----
DISK_00                c0t5000C500101BA0A3d0
DISK_01                c0t5000C500101B95BBd0
DISK_02                -
```

示例 4 Reporting Physical Path Information

The following command reports physical path information for a specific disk.

```
# croinfo -c c0t5000C500101BA0A3d0 -o cp
c:occupant-compdev
-----
c0t5000C500101BA0A3d0

p:occupant-paths
-----
/devices/pci@0,0/pci10de,5d@pci11f8,8001@0/iport@f/disk@w5000c500101ba0a1,0
/devices/pci@7b,0/pci10de,5d@pci11f8,8001@0/iport@f/disk@w5000c500101ba0a2,0
```

Note that *occupant-paths* has multiple string values.

示例 5 Making Inventory of Disks

The following example reports how many of a specific type of disk are available using *occupant-part*.

```
# for i in `croinfo -h -o n | sort -u`
> do
> echo $i "\\t\\c";croinfo -h -n $i | wc -l
> done
SEAGATE-ST330055SSUN300G          3
SEAGATE-ST330056SSUN300G          19
SEAGATE-ST345056SSUN450G          5
```

示例 6 Locating a Specific Type of Disk

The following command reports where disks of a specific type are located, what their *ctd* name is (by means of *occupant-compdev*), and what firmware level they are at.

```
# croinfo -n SEAGATE-ST330055SSUN300G -o Dcf
D:Devchassis              c:component
-----
/dev/chassis/RACK29.U29-32/SCSI_Device__11/disk c0t5000C50007DD49F7d0
/dev/chassis/RACK29.U33-36/SCSI_Device__18/disk c0t5000C50008F7FB4Fd0
/dev/chassis/RACK29.U33-36/SCSI_Device__19/disk c0t5000C50007DD412Fd0
```

示例 6 Locating a Specific Type of Disk (续)

```
f:firm
-----
0892
0892
0892
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	See below

The interface stability of `croinfo` and `diskinfo` is Committed. The interface stability of command output is Not-an-Interface.

另请参见

[devchassisd\(1M\)](#), [fmadm\(1M\)](#), [fmd\(1M\)](#), [attributes\(5\)](#), [devchassis\(7FS\)](#)

The SCSI Storage Interfaces committee website, <http://www.t10.org>

SCSI Primary Commands-4, SPC4; SCSI Enclosure Services-2, SES2, Serial Attached SCSI-2, SAS2

附注

`croinfo` representation depends on the ability of [fmd\(1M\)](#) to enumerate system topology and accurately represent associated chassis, receptacles, and occupants. These dependencies might extend through [fmd\(1M\)](#) and require that connected hardware, and its associated firmware, comply with specific standards. For disk bays, this requires that storage chassis behave in a T10 standards-compliant (SPC4 and SES2) fashion. Storage chassis that do not respond appropriately might not report chassis, bays, or disk nodes correctly. Specifically, `diskinfo` requires that chassis support SES diagnostic page 0xa (Additional Element Status) and set the Element Index Present (EIP) bit to 1. Enclosures that do not meet this criterion will not be fully enumerated, and thus will not be properly represented.

引用名	cron – clock daemon
用法概要	/usr/sbin/cron
描述	cron starts a process that executes commands at specified dates and times. You can specify regularly scheduled commands to cron according to instructions found in crontab files in the directory /var/spool/cron/crontabs. Users can submit their own crontab file using the crontab(1) command. Commands which are to be executed only once can be submitted using the at(1) command. cron only examines crontab or at command files during its own process initialization phase and when the crontab or at command is run. This reduces the overhead of checking for new or changed files at regularly scheduled intervals. As cron never exits, it should be executed only once. This is done routinely by way of the svc:/system/cron:default service. The file /etc/cron.d/FIFO file is used as a lock file to prevent the execution of more than one instance of cron. cron captures the output of the job's stdout and stderr streams, and, if it is not empty, mails the output to the user. If the job does not produce output, no mail is sent to the user. An exception is if the job is an at(1) job and the -m option was specified when the job was submitted. cron and at jobs are not executed if your account is locked. Jobs and processes execute. The shadow(4) file defines which accounts are not locked and will have their jobs and processes executed.
Setting cron Jobs Across Timezones	The timezone of the cron daemon sets the system-wide timezone for cron entries. This, in turn, is by set by default system-wide using /etc/default/init. The timezone for cron entries can be overridden in a user's crontab file; see crontab(1). If some form of <i>daylight savings</i> or <i>summer/winter time</i> is in effect, then jobs scheduled during the switchover period could be executed once, twice, or not at all.
Setting cron Defaults	To keep a log of all actions taken by cron, you must specify CRONLOG=YES in the /etc/default/cron file. If you specify CRONLOG=NO, no logging is done. Keeping the log is a user configurable option since cron usually creates huge log files. You can specify the PATH for user cron jobs by using PATH= in /etc/default/cron. You can set the PATH for root cron jobs using SUPATH= in /etc/default/cron. Carefully consider the security implications of setting PATH and SUPATH. Example /etc/default/cron file: <pre>CRONLOG=YES PATH=/usr/bin:/usr/ucb:</pre> This example enables logging and sets the default PATH used by non-root jobs to /usr/bin:/usr/ucb:. Root jobs continue to use /usr/sbin:/usr/bin.

The cron log file is periodically rotated by [logadm\(1M\)](#).

文件

/etc/cron.d

Main cron directory

/etc/cron.d/FIFO

Lock file

/etc/default/cron

cron default settings file

/var/cron/log

cron history information

/var/spool/cron

Spool area

/etc/cron.d/queuedefs

Queue description file for at, batch, and cron

/etc/logadm.conf

Configuration file for logadm

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[svcs\(1\)](#), [at\(1\)](#), [crontab\(1\)](#), [sh\(1\)](#), [logadm\(1M\)](#), [svcadm\(1M\)](#), [queuedefs\(4\)](#), [shadow\(4\)](#), [attributes\(5\)](#), [rbac\(5\)](#), [smf\(5\)](#), [smf_security\(5\)](#)

附注

The cron service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/system/cron:default`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command. Most administrative actions may be delegated to users with the `solaris.smf.manage.cron` authorization (see [rbac\(5\)](#) and [smf_security\(5\)](#)).

诊断

A history of all actions taken by cron is stored in `/var/cron/log` and possibly in `/var/cron/olog`.

引用名

cryptoadm – 加密框架管理

用法概要

```

cryptoadm list [-mpv] [provider=provider-name]
    [mechanism=mechanism-list]

cryptoadm disable
    provider=provider-name mechanism=mechanism-list | random | all

cryptoadm enable
    provider=provider-name mechanism=mechanism-list | random | all

cryptoadm install provider=provider-name

cryptoadm install provider=provider-name
    [mechanism=mechanism-list]

cryptoadm uninstall provider=provider-name

cryptoadm unload provider=provider-name

cryptoadm disable fips-140

cryptoadm enable fips-140

cryptoadm list fips-140

cryptoadm refresh

cryptoadm start

cryptoadm stop

cryptoadm --help

```

描述

cryptoadm 实用程序可以显示系统的加密提供器信息，配置每个提供器的机制策略，以及安装或卸载加密提供器。加密框架支持三种类型的提供器：用户级提供器（PKCS11 共享库）、内核软件提供器（可装入的内核软件模块）以及内核硬件提供器（加密硬件设备）。

对于内核软件提供器，**cryptoadm** 实用程序提供了 **unload** 子命令。该子命令可指示内核卸载内核软件提供器。

对于加密框架的 **metaslot**，**cryptoadm** 实用程序提供了子命令，用于启用和禁用 **metaslot** 的功能、列出 **metaslot** 的配置、指定备用持久性对象存储以及配置 **metaslot** 的机制策略。

cryptoadm 实用程序提供了子命令，用于启用和禁用加密框架中的 FIPS-140 模式。此外，它还提供了 **list** 子命令用来显示 FIPS-140 模式的当前状态。

管理员会发现，使用 **syslog** 实用程序（参见 [syslogd\(1M\)](#) 和 [logadm\(1M\)](#)）来维护加密子系统非常有用。日志记录在以下情况下尤其有用：

- 如果内核级守护进程停用，则所有应用程序都会发生故障。您可以通过系统日志了解这一情况，并使用 [svcadm\(1M\)](#) 重新启动 **svc:/system/cryptosvc** 服务。

- 如果在框架中插入了故障提供器，您可以通过系统日志了解这一情况，然后将其从框架中删除。

`cryptoadm` 命令只能由特权用户运行，但下列子命令或选项除外。

- 子命令 `list`，任何选项
- 子命令 `--help`

选项

`cryptoadm` 实用程序具有如下所示的各种子命令和选项的组合。

`cryptoadm list`
显示已安装提供器的列表。

`cryptoadm list metaslot`
显示 `metaslot` 的系统范围配置。

`cryptoadm list -m [ provider=provider-name | metaslot ]`
显示可用于已安装的提供器或 `metaslot` 的机制的列表。如果指定了提供器，则显示指定提供器的名称以及可用于该提供器的机制列表。如果指定了 `metaslot` 关键字，则显示可用于 `metaslot` 的机制列表。

`cryptoadm list -p [ provider=provider-name | metaslot ]`
显示已安装提供器的机制策略（即哪些机制可用，哪些机制不可用）。此外，还会显示提供器功能策略或 `metaslot`。如果指定了提供器，则仅显示对其强制实施了机制策略的提供器的名称。如果指定了 `metaslot` 关键字，则显示对 `metaslot` 强制实施的机制策略。

`cryptoadm list -v provider=provider-name | metaslot`
如果指定了提供器，则显示指定提供器的详细信息。如果指定了 `metaslot` 关键字，则显示有关 `metaslot` 的详细信息。

`-v`
对于上述各种 `list` 子命令（`list -p` 除外），使用 `-v`（即详细）选项可获取有关提供器、机制和插槽的详细信息。

`cryptoadm disable provider=provider-name`
`[ mechanism=mechanism-list | provider-feature ... | all ]`
禁用为提供器指定的机制或提供器功能。有关 *mechanism*、*provider-feature* 和 `all` 关键字的描述，请参见“操作数”部分。

`cryptoadm [ mechanism=mechanism-list ] [ auto-key-migrate ]`
禁用加密框架中的 `metaslot` 功能，或者禁用 `metaslot` 的某些功能。如果没有指定任何操作数，此命令将禁用加密框架中的 `metaslot` 功能。如果指定了机制列表，则禁用为 `metaslot` 指定的机制。如果禁用了 `metaslot` 的所有机制，则 `metaslot` 也将被禁用。有关机制的描述，请参见“操作数”部分。如果指定了 `auto-key-migrate` 关键字，即使必须执行加密操作，该关键字也将禁止将敏感令牌对象迁移到其他插槽。有关 `auto-key-migrate` 的描述，请参见“操作数”部分。

`cryptoadm enable provider=provider-name`
`[ mechanism=mechanism-list | provider-feature ... | all ]`

启用为提供器指定的机制或提供器功能。有关 *mechanism*、*provider-feature* 和 *all* 关键字的描述，请参见“操作数”部分。

```
cryptoadm enable metaslot [ mechanism=mechanism-list ] |
[ [ token=token-label] [ slot=slot-description] |
default-keystore ] | [ auto-key-migrate ]
```

如果没有指定任何操作数，此命令将启用加密框架中的 *metaslot* 功能。如果指定了机制列表，则仅启用为 *metaslot* 指定的机制。如果指定了 *token-label*，则将使用指定的令牌作为持久性对象存储。如果指定了 *slot-description*，则将使用指定的插槽作为持久性对象存储。如果同时指定了 *token-label* 和 *slot-description*，则将使用带有匹配的令牌标签和插槽描述的提供器作为持久性对象存储。如果指定了 *default-keystore* 关键字，*metaslot* 将使用缺省的持久性对象存储。如果指定了 *auto-key-migrate* 关键字，敏感令牌对象将根据需要自动迁移到其他插槽，以完成某些特定加密操作。有关机制、令牌、插槽、*default-keystore* 和 *auto-key-migrate* 的描述，请参见“操作数”部分。

```
cryptoadm install provider=provider-name
```

将用户级提供器安装到系统中。*provider* 操作数必须是对应共享库的绝对路径名。如果一个库存在 32 位和 64 位两个版本，则该命令只能使用包含 *\$ISA* 的路径名参数运行一次。请注意，*\$ISA* 不是对环境变量的引用。另请注意，*\$ISA* 必须加引号（使用单引号，例如，'*\$ISA*'），或者，必须对 *\$* 进行转义，以免被 *shell* 错误地扩展。用户级框架可将 *\$ISA* 扩展为空字符串或体系结构特定的目录，例如 *sparcv9*。

安装用户级提供器时，推荐的方法是为该提供器构建一个软件包。有关更多信息，请参见 Solaris 开发者安全性指南。

```
cryptoadm install provider=provider-name
mechanism=mechanism-list
```

将内核软件提供器安装到系统中。提供器只应包含基名。*mechanism-list* 操作数指定了该提供器要支持的机制的完整列表。

安装内核级提供器时，推荐的方法是为提供器构建一个软件包。有关更多信息，请参见 Solaris 开发者安全性指南。

```
cryptoadm uninstall provider=provider-name
```

从系统中卸载指定的 *provider* 以及关联的机制策略。此子命令仅适用于用户级提供器或内核软件提供器。

```
cryptoadm unload provider=provider-name
```

卸载由 *provider* 指定的内核软件模块。

```
cryptoadm disable fips-140
```

禁用加密框架中以及硬件提供器的 FIPS-140 模式。

```
cryptoadm enable fips-140
```

启用加密框架中以及硬件提供器的 FIPS-140 模式。此子命令不会禁用用户级 *pkcs11* *softtoken* 库以及内核软件提供器中的未经 FIPS 批准的算法。框架使用者只能使用经 FIPS 批准的算法。

执行此子命令以后，系统会立即发出一条消息，告知管理员添加任何不在允许范围内的插件可能会违反 FIPS 遵从性，并会检查这些插件的安全策略。

系统将需要重新引导来执行“通电自检”，包括加密算法测试和软件完整性测试。

cryptoadm list fips-140

显示加密框架中以及硬件提供器的 FIPS-140 模式的当前设置。FIPS-140 模式的状态为 **enabled** 或 **disabled**。缺省 FIPS-140 模式为 **disabled**。

cryptoadm refresh

cryptoadm start

cryptoadm stop

供 **smf(5)** 使用的专用接口，不得直接使用这些接口。

cryptoadm -help

显示命令用法。

操作数

provider=provider-name

用户级提供器（PKCS#11 共享库）、内核软件提供器（可装入的内核软件模块）或内核硬件提供器（加密硬件设备）。

provider 操作数的有效值是格式为 **cryptoadm list** 的命令输出项中的一项。用户级提供器的 *provider* 操作数是对应共享库的绝对路径名。内核软件提供器的 *provider* 操作数仅包含基名。内核硬件提供器的 *provider* 操作数采用 *name/number* 格式。

mechanism=mechanism-list

包含一个或多个 PKCS #11 机制的由逗号分隔的列表。PKCS #11 规范中定义了实现加密操作的进程。您可以使用 **all** 来替换 *mechanism-list*，以指定提供器的所有机制。请参见下面关于 **all** 关键字的讨论。

provider-feature

给定提供器的加密框架功能。目前只接受 **random** 功能。对于用户级提供器，禁用随机功能将使提供器无法使用 PKCS #11 例程 **C_GenerateRandom** 和 **C_SeedRandom**。对于内核提供器，禁用随机功能将导致 **/dev/random** 无法从该提供器收集随机数。

all

可以将关键字 **all** 与 **disable** 和 **enable** 子命令结合使用，以便对所有提供器功能进行操作。

token=token-label

加密框架中某个提供器的令牌的标签。

令牌操作数的有效值是 **cryptoadm list -v** 命令输出中“Token Label”下显示的项。

slot=slot-description

加密框架中某个提供器的插槽的描述。

插槽操作数的有效值是 **cryptoadm list -v** 命令输出中“Description”下显示的项。

default-keystore

关键字 **default-keystore** 仅对 **metaslot** 有效。通过指定此关键字，可将 **metaslot** 的持久性对象存储重新设置为使用缺省存储。

auto-key-migrate

关键字 **auto-key-migrate** 仅对 **metaslot** 有效。通过指定此关键字，可以配置是否允许 **metaslot** 将敏感令牌对象从令牌对象插槽移至其他插槽来执行加密操作。

可通过下述两种方式将关键字 **all** 与 **disable** 和 **enable** 子命令结合使用：

- 您可以使用 **all** 来替换 **mechanism=mechanism-list** 及其他任何提供器功能，如下所示：

```
# cryptoadm enable provider=dca/0 all
```

此命令可启用提供器的机制以及其他任何提供器功能，如 **random**。

- 您也可以将 **all** 用作 **mechanism** 的参数，如下所示：

```
# cryptoadm enable provider=des mechanism=all
```

此命令可启用提供器的所有机制，但不会启用其他任何提供器功能，如 **random**。

示例

示例1 显示系统中已安装提供器的列表

以下命令显示所有已安装提供器的列表：

```
example% cryptoadm list
user-level providers:
/usr/lib/security/$ISA/pkcs11_kernel.so
/usr/lib/security/$ISA/pkcs11_softtoken.so
/opt/lib/libcryptoki.so.1
/opt/system/core-osonn/lib/$ISA/libpkcs11.so.1
```

kernel software providers:

```
des
aes
bfish
sha1
md5
```

kernel hardware providers:

```
dca/0
```

示例2 显示 **md5** 提供器的机制列表

以下命令是 **list** 子命令的变体形式：

```
example% cryptoadm list -m provider=md5
md5: CKM_MD5,CKM_MD5_HMAC,CKM_MD5_HMAC_GENERAL
```

示例3 禁用内核软件提供器的特定机制

以下命令禁用内核软件提供器 des 的 CKM_DES3_ECB 机制和 CKM_DES3_CBC 机制：

```
example# cryptoadm disable provider=des
```

示例4 显示提供器的机制策略

以下命令显示 des 提供器的机制策略：

```
example% cryptoadm list -p provider=des
des: All mechanisms are enabled, except CKM_DES3_ECB, CKM_DES3_CBC
```

示例5 启用提供器的特定机制

以下命令启用内核软件提供器 des 的 CKM_DES3_ECB 机制：

```
example# cryptoadm enable provider=des mechanism=CKM_DES3_ECB
```

示例6 安装用户级提供器

以下命令安装用户级提供器：

```
example# cryptoadm install provider=/opt/lib/libcryptoki.so.1
```

示例7 安装包含 32 位和 64 位版本的用户级提供器

以下命令安装包含 32 位和 64 位版本的用户级提供器：

```
example# cryptoadm install \
provider=/opt/system/core-osonn/lib/'$ISA'/libpkcs11.so.1
```

示例8 卸载提供器

以下命令卸载 md5 提供器：

```
example# cryptoadm uninstall provider=md5
```

示例9 禁用 metaslot

以下命令禁用加密框架中的 metaslot 功能。

```
example# cryptoadm disable metaslot
```

示例10 指定 metaslot 使用指定的令牌作为持久性对象存储

以下命令指定 metaslot 使用 Venus 令牌作为持久性对象存储。

```
example# cryptoadm enable metaslot token="SUNW,venus"
```

退出状态

将返回下列退出值：

0

成功完成。

>0
出现错误。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	请参见下文。

start、stop 和 refresh 选项是专用接口。其他所有选项及实用程序名称均为 Committed（已确定）。

另请参见 [logadm\(1M\)](#)、[svcadm\(1M\)](#)、[syslogd\(1M\)](#)、[libpkcs11\(3LIB\)](#)、[exec_attr\(4\)](#)、[prof_attr\(4\)](#)、[attribu](#)
《Oracle Solaris 管理：安全服务》

Solaris 开发者安全性指南

附注 即使已明确制定硬件提供器的策略（换句话说，已禁用其中的某些机制），且该硬件提供器已分离，仍然会列出该硬件提供器的策略。

cryptoadm 会假定为每个用户级提供器至少提供一个 32 位共享对象。如果同时提供 32 位和 64 位共享对象，这两个版本必须提供相同的功能。这两个版本也使用同一个机制策略。

引用名	cvcd – virtual console daemon						
用法概要	<pre>/platform/platform_name/cvcd [-a auth] [-e encr] [-u esp_auth]</pre>						
描述	The virtual console daemon, cvcd, is a server process that supports the network console provided on some platforms. The cvcd daemon accepts network console connections from a remote host (only one host at any given time). Console input is read from this connection and forwarded to cvc(7D) by way of cvcredir(7D). Similarly, console output is read from cvcredir(7D) and forwarded across the network console connection. If cvcd dies, console traffic is automatically rerouted through an internal hardware interface. The cvcd daemon normally starts at system boot time. Each domain supports only one cvcd process at a time. 注意 – On Sun Enterprise 10000 domains, cvcd uses a configuration file (/etc/ssphostname) to determine the name of the host from which network console connections are allowed. If the remote console host is renamed, you must edit the configuration file to reflect that change. The cvcd daemon supports per-socket IP Security Architecture (IPsec) through the options described below. See ipsec(7P).						
选项	The cvcd daemon supports the options listed below. <table> <tr> <td>-a auth</td><td>Controls the IPsec Authentication Header (AH) algorithm. <i>auth</i> can be one of none, md5, or sha1.</td></tr> <tr> <td>-e encr</td><td>Controls the IPsec Encapsulating Security Payload (ESP) encryption algorithm. <i>encr</i> can be one of none, des, or 3des.</td></tr> <tr> <td>-u esp_auth</td><td>Controls the IPsec Encapsulating Security Payload (ESP) authentication algorithm. <i>esp_auth</i> can be one of none, md5, or sha1.</td></tr> </table>	-a auth	Controls the IPsec Authentication Header (AH) algorithm. <i>auth</i> can be one of none, md5, or sha1.	-e encr	Controls the IPsec Encapsulating Security Payload (ESP) encryption algorithm. <i>encr</i> can be one of none, des, or 3des.	-u esp_auth	Controls the IPsec Encapsulating Security Payload (ESP) authentication algorithm. <i>esp_auth</i> can be one of none, md5, or sha1.
-a auth	Controls the IPsec Authentication Header (AH) algorithm. <i>auth</i> can be one of none, md5, or sha1.						
-e encr	Controls the IPsec Encapsulating Security Payload (ESP) encryption algorithm. <i>encr</i> can be one of none, des, or 3des.						
-u esp_auth	Controls the IPsec Encapsulating Security Payload (ESP) authentication algorithm. <i>esp_auth</i> can be one of none, md5, or sha1.						
操作数	The following operands are supported: <table> <tr> <td>platform_name</td><td>The official Sun platform name used in packaging and code. For example, for Sun Fire 15K servers, the <i>platform_name</i> would be SUNW, Sun-Fire-15000.</td></tr> </table>	platform_name	The official Sun platform name used in packaging and code. For example, for Sun Fire 15K servers, the <i>platform_name</i> would be SUNW, Sun-Fire-15000.				
platform_name	The official Sun platform name used in packaging and code. For example, for Sun Fire 15K servers, the <i>platform_name</i> would be SUNW, Sun-Fire-15000.						
示例	示例 1 Setting an IPsec Option The command below sets the value of the IPsec Authentication Header algorithm to md5. As a result of this command, cvcd will use the HMAC-MD5 authentication algorithm. <pre># svccfg -s svc:/system/cvc setprop cvc/ah_auth = "md5" # svccfg -s svc:/system/cvc setprop cvc/esp_encr = "none" # svccfg -s svc:/system/cvc setprop cvc/esp_auth = "none" # svcadm refresh svc:/system/cvc</pre>						

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	Sun Enterprise 10000 servers, Sun Fire High-End Systems
Availability	system/network-console

另请参见 [svcs\(1\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [services\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [cvc\(7D\)](#), [cvcredir\(7D\)](#), [ipsec\(7P\)](#)

Sun Enterprise 10000 SSP Reference Manual
System Management Services (SMS) Reference Manual

附注 The cvcd service is managed by the service management facility, [smf\(5\)](#), under the fault management resource identifier (FMRI):

`svc:/system/cvc`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#) or [svccfg\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	datadm – maintain DAT static registry file
用法概要	<pre>/usr/bin/datadm [-v] [-u] [-a service_provider.conf] [-r service_provider.conf]</pre>
描述	The datadm utility maintains the DAT static registry file, dat.conf(4). This administrative configuration program allows uDAPL service providers to add and remove themselves to the <code>dat.conf</code> file. You can add or remove interface adapters that a service provider supports from a system after its installation. You can use datadm to update the <code>dat.conf</code> file to reflect the current state of the system. A new set of interface adapters for all the service providers currently installed is regenerated.
选项	The following options are supported: -a service_provider.conf Enumerate each device entry in the service_provider.conf(4) file into a list of interface adapters, that is, interfaces to external network that are available to uDAPL consumers. -r service_provider.conf Remove the list of interface adapters that corresponds to the device entry in the service_provider.conf(4) file. -u Update the <code>dat.conf</code> to reflect the current state of the system with an up to date set of interface adapters for the service providers that are currently listed in the DAT static registry. -v Display the DAT static registry file, <code>dat.conf</code>.
示例	示例 1 Enumerating a Device Entry The following example enumerates a device entry in the service_provider.conf(4) file into interface adapters in the dat.conf(4) file. Assume that SUNW has a service provider library that supports the device <code>tavor</code>. It has a service_provider.conf(4) file installed in the directory <code>/usr/share/dat/SUNWudaplt.conf</code> with a single entry as follows: <pre>driver_name=tavor u1.2 nonthreadsafe default\ udapl_tavor.so.1 SUNW.1.0 ""</pre> <code>tavor</code> is an Infiniband Host Channel Adapter with two ports. Both IB ports exist in a single IB partition, <code>0x8001</code>. If an IB partition is created and plumbed to each port (with the names <code>p8001.ibd0</code> and <code>p8001.ibd1</code>), there will be two IB partition instances. See dladm(1M) for more information on creating IB partition data links.

示例 1 Enumerating a Device Entry (续)

```
# dladm show-part

LINK          PKEY    OVER    STATE    FLAGS
p8001.ibd0    8001    ibd0    unknown  ----
p8001.ibd1    8001    ibd1    unknown  ----
```

Running the command:

```
# datadm -a /usr/share/dat/SUNWudaplt.conf

...appends two new entries (if they do not already exist) in the /etc/dat/dat.conf file:

p8001.ibd0 u1.2 nonthreadsafe default udapl_tavor.so.1 SUNW.1.0 ""
"driver_name=tavor"
p8001.ibd1 u1.2 nonthreadsafe default udapl_tavor.so.1 SUNW.1.0 ""
"driver_name=tavor"
```

示例 2 Updating the dat.conf to Reflect the Current State of the System

A new IB partition, 0x8002, is added to the above example covering port 1 of the Host Channel Adapter. If a new IB partition is created on the port 1/partition 0x8002 with the partition link name specified as p8002.ibd0, there will be a third IB partition instance: p8002.ibd0.

```
# dladm show-part

LINK          PKEY    OVER    STATE    FLAGS
p8001.ibd0    8001    ibd0    unknown  ----
p8001.ibd1    8001    ibd1    unknown  ----
p8002.ibd0    8000    ibd0    unknown  ----
```

Running datadm -u command, updates the /etc/dat/dat.conf file with a new entry added reflecting the current state of the system.

datadm -v shows that there are now three entries in the /etc/dat/dat.conf file:

```
p8001.ibd0 u1.2 nonthreadsafe default udapl_tavor.so.1 SUNW.1.0 ""
"driver_name=tavor"
p8001.ibd1 u1.2 nonthreadsafe default udapl_tavor.so.1 SUNW.1.0 ""
"driver_name=tavor"
p8002.ibd0 u1.2 nonthreadsafe default udapl_tavor.so.1 SUNW.1.0 ""
"driver_name=tavor"
```

文件 /etc/dat/dat.conf
DAT static registry file

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/io/infiniband/udapl
Interface Stability	Committed

另请参见

[dladm\(1M\)](#), [pkgadd\(1M\)](#), [pkgrm\(1M\)](#), [libdat\(3LIB\)](#), [dat.conf\(4\)](#),
[service_provider.conf\(4\)](#), [attributes\(5\)](#)

引用名 dcs – domain configuration server

用法概要 `/usr/lib/dcs [-s sessions] [ [-a auth] [-e encr] [-u esp_auth] ] [-l]`

描述 The Domain Configuration Server (DCS) is a daemon process that runs on Sun servers that support remote Dynamic Reconfiguration (DR) clients. It is started by the Service Management Facility (see [smf\(5\)](#)) when the first DR request is received from a client connecting to the network service `sun-dr`. After the DCS accepts a DR request, it uses the [libcfgadm\(3LIB\)](#) interface to execute the DR operation. After the operation is performed, the results are returned to the client.

The DCS listens on the network service labeled `sun-dr`. Its underlying protocol is TCP. It is invoked as a server program by the SMF using the TCP transport. The fault management resource identifier (FMRI) for DCS is:

```
svc:/platform/sun4u/dcs:default
```

If you disable this service, DR operations initiated from a remote host fail. There is no negative impact on the server.

Security for the DCS connection is provided differently based upon the architecture of the system. The SMF specifies the correct options when invoking the DCS daemon, based upon the current architecture. For all architectures, security is provided on a per-connection basis.

The DCS daemon has no security options that are applicable when used on a Sun Enterprise 10000 system. So there are no options applicable to that architecture.

The security options for Sun Fire high-end systems are based on IPsec options defined as SMF properties. These options include the `-a auth`, `-e encr`, and `-u esp_auth` options, and can be set using the [svccfg\(1M\)](#) command. These options must match the IPsec policies defined for DCS on the system controller. Refer to the `kmd(1M)` man page in the System Management Services (SMS) Reference Manual. The `kmd(1M)` man page is not part of the SunOS man page collection.

Security on SPARC Enterprise Servers is not configurable. The DCS daemon uses a platform-specific library to configure its security options when running on such systems. The `-l` option is provided by the SMF when invoking the DCS daemon on SPARC Enterprise Servers. No other security options to the DCS daemon should be used on SPARC Enterprise Servers.

选项 The following options are supported:

- `-a auth` Controls the IPsec Authentication Header (AH) algorithm. *auth* can be one of none, md5, or sha1.
- `-e encr` Controls the IPsec Encapsulating Security Payload (ESP) encryption algorithm. *encr* can be one of none, des, or 3des.

- l Enables the use of platform-specific security options on SPARC Enterprise Servers.
- s *sessions* Sets the number of active sessions that the DCS allows at any one time. When the limit is reached, the DCS stops accepting connections until active sessions complete the execution of their DR operation. If this option is not specified, a default value of 128 is used.
- u *esp_auth* Controls the IPsec Encapsulating Security Payload (ESP) authentication algorithm. *esp_auth* can be one of none, md5, or sha1.

示例

示例 1 Setting an IPsec Option

The following command sets the Authentication Header algorithm for the DCS daemon to use the HMAC-MD5 authentication algorithm. These settings are only applicable for using the DCS daemon on a Sun Fire high-end system.

```
# svccfg -s svc:/platform/sun4u/dcs setprop dcs/ah_auth = "md5"
# svccfg -s svc:/platform/sun4u/dcs setprop dcs/esp_encr = "none"
# svccfg -s svc:/platform/sun4u/dcs setprop dcs/esp_auth = "none"
# svcadm refresh svc:/platform/sun4u/dcs
```

错误

The DCS uses [syslog\(3C\)](#) to report status and error messages. All of the messages are logged with the LOG_DAEMON facility. Error messages are logged with the LOG_ERR and LOG_NOTICE priorities, and informational messages are logged with the LOG_INFO priority. The default entries in the `/etc/syslog.conf` file log all of the DCS error messages to the `/var/adm/messages` log.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/domain-configuration/sparc-enterprise, SUNWdcsr
Interface Stability	Committed

另请参见

[svcs\(1\)](#), [cfgadm_sbd\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [syslog\(3C\)](#), [config_admin\(3CFGADM\)](#), [libcfgadm\(3LIB\)](#), [syslog.conf\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [dr\(7d\)](#)

附注

The dcs service is managed by the service management facility, [smf\(5\)](#), under the fault management resource identifier (FMRI):

```
svc:/platform/sun4u/dcs:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名 dd – convert and copy a file

用法概要 /usr/bin/dd [*operand=value*]...

描述 The dd utility copies the specified input file to the specified output with possible conversions. The standard input and output are used by default. The input and output block sizes may be specified to take advantage of raw physical I/O. Sizes are specified in bytes; a number may end with k, b, or w to specify multiplication by 1024, 512, or 2, respectively. Numbers may also be separated by x to indicate multiplication.

The dd utility reads the input one block at a time, using the specified input block size. dd then processes the block of data actually returned, which could be smaller than the requested block size. dd applies any conversions that have been specified and writes the resulting data to the output in blocks of the specified output block size.

cbs is used only if *ascii*, *asciib*, *unblock*, *ebcdic*, *ebcdicb*, *ibm*, *ibmb*, or *block* conversion is specified. In the first two cases, cbs characters are copied into the conversion buffer, any specified character mapping is done, trailing blanks are trimmed, and a *NEWLINE* is added before sending the line to output. In the last three cases, characters up to *NEWLINE* are read into the conversion buffer and blanks are added to make up an output record of size cbs. ASCII files are presumed to contain *NEWLINE* characters. If cbs is unspecified or 0, the *ascii*, *asciib*, *ebcdic*, *ebcdicb*, *ibm*, and *ibmb* options convert the character set without changing the input file's block structure. The *unblock* and *block* options become a simple file copy.

After completion, dd reports the number of whole and partial input and output blocks.

操作数 The following operands are supported:

<i>if=file</i>	Specifies the input path. Standard input is the default.
<i>of=file</i>	Specifies the output path. Standard output is the default. If the <i>seek=expr</i> conversion is not also specified, the output file will be truncated before the copy begins, unless <i>conv=notrunc</i> is specified. If <i>seek=expr</i> is specified, but <i>conv=notrunc</i> is not, the effect of the copy will be to preserve the blocks in the output file over which dd seeks, but no other portion of the output file will be preserved. (If the size of the seek plus the size of the input file is less than the previous size of the output file, the output file is shortened by the copy.)
<i>ibs=n</i>	Specifies the input block size in <i>n</i> bytes (default is 512).
<i>obs=n</i>	Specifies the output block size in <i>n</i> bytes (default is 512).
<i>bs=n</i>	Sets both input and output block sizes to <i>n</i> bytes, superseding <i>ibs=</i> and <i>obs=</i> . If no conversion other than <i>sync</i> , <i>noerror</i> , and <i>notrunc</i> is specified, each input block is copied to the output as a single block without aggregating short blocks.

<code>cbs=<i>n</i></code>	Specifies the conversion block size for <code>block</code> and <code>unblock</code> in bytes by <i>n</i> (default is 0). If <code>cbs=</code> is omitted or given a value of 0, using <code>block</code> or <code>unblock</code> produces unspecified results. This option is used only if ASCII or EBCDIC conversion is specified. For the <code>ascii</code> and <code>asciib</code> operands, the input is handled as described for the <code>unblock</code> operand except that characters are converted to ASCII before the trailing SPACE characters are deleted. For the <code>ebcdic</code>, <code>ebcdicb</code>, <code>ibm</code>, and <code>ibmb</code> operands, the input is handled as described for the <code>block</code> operand except that the characters are converted to EBCDIC or IBM EBCDIC after the trailing SPACE characters are added.								
<code>files=<i>n</i></code>	Copies and concatenates <i>n</i> input files before terminating (makes sense only where input is a magnetic tape or similar device).								
<code>skip=<i>n</i></code>	Skips <i>n</i> input blocks (using the specified input block size) before starting to copy. On seekable files, the implementation reads the blocks or seeks past them. On non-seekable files, the blocks are read and the data is discarded.								
<code>iseek=<i>n</i></code>	Seeks <i>n</i> blocks from beginning of input file before copying (appropriate for disk files, where <code>skip</code> can be incredibly slow).								
<code>oseek=<i>n</i></code>	Seeks <i>n</i> blocks from beginning of output file before copying.								
<code>seek=<i>n</i></code>	Skips <i>n</i> blocks (using the specified output block size) from beginning of output file before copying. On non-seekable files, existing blocks are read and space from the current end-of-file to the specified offset, if any, is filled with null bytes. On seekable files, the implementation seeks to the specified offset or reads the blocks as described for non-seekable files.								
<code>count=<i>n</i></code>	Copies only <i>n</i> input blocks.								
<code>conv=<i>value</i>[, <i>value</i>. . .]</code>	Where <i>values</i> are comma-separated symbols from the following list: <table><tr><td><code>ascii</code></td><td>Converts EBCDIC to ASCII.</td></tr><tr><td><code>asciib</code></td><td>Converts EBCDIC to ASCII using BSD-compatible character translations.</td></tr><tr><td><code>ebcdic</code></td><td>Converts ASCII to EBCDIC. If converting fixed-length ASCII records without NEWLINES, sets up a pipeline with <code>dd conv=unblock</code> beforehand.</td></tr><tr><td><code>ebcdicb</code></td><td>Converts ASCII to EBCDIC using BSD-compatible character translations. If converting fixed-length</td></tr></table>	<code>ascii</code>	Converts EBCDIC to ASCII.	<code>asciib</code>	Converts EBCDIC to ASCII using BSD-compatible character translations.	<code>ebcdic</code>	Converts ASCII to EBCDIC. If converting fixed-length ASCII records without NEWLINES, sets up a pipeline with <code>dd conv=unblock</code> beforehand.	<code>ebcdicb</code>	Converts ASCII to EBCDIC using BSD-compatible character translations. If converting fixed-length
<code>ascii</code>	Converts EBCDIC to ASCII.								
<code>asciib</code>	Converts EBCDIC to ASCII using BSD-compatible character translations.								
<code>ebcdic</code>	Converts ASCII to EBCDIC. If converting fixed-length ASCII records without NEWLINES, sets up a pipeline with <code>dd conv=unblock</code> beforehand.								
<code>ebcdicb</code>	Converts ASCII to EBCDIC using BSD-compatible character translations. If converting fixed-length								

- ASCII records without `NEWLINEs`, sets up a pipeline with `dd conv=unblock` beforehand.
- `ibm` Slightly different map of ASCII to EBCDIC. If converting fixed-length ASCII records without `NEWLINEs`, sets up a pipeline with `dd conv=unblock` beforehand.
- `ibmb` Slightly different map of ASCII to EBCDIC using BSD-compatible character translations. If converting fixed-length ASCII records without `NEWLINEs`, sets up a pipeline with `dd conv=unblock` beforehand.

The `ascii` (or `asciib`), `ebcdic` (or `ebcdicb`), and `ibm` (or `ibmb`) values are mutually exclusive.

- `block` Treats the input as a sequence of `NEWLINE`-terminated or EOF-terminated variable-length records independent of the input block boundaries. Each record is converted to a record with a fixed length specified by the conversion block size. Any `NEWLINE` character is removed from the input line. `SPACE` characters are appended to lines that are shorter than their conversion block size to fill the block. Lines that are longer than the conversion block size are truncated to the largest number of characters that will fit into that size. The number of truncated lines is reported.
- `unblock` Converts fixed-length records to variable length. Reads a number of bytes equal to the conversion block size (or the number of bytes remaining in the input, if less than the conversion block size), delete all trailing `SPACE` characters, and append a `NEWLINE` character.

The `block` and `unblock` values are mutually exclusive.

- `lcase` Maps upper-case characters specified by the `LC_CTYPE` keyword to `lower` to the corresponding lower-case character. Characters for which no mapping is specified are not modified by this conversion.
- `ucase` Maps lower-case characters specified by the `LC_CTYPE` keyword to `upper` to the corresponding upper-case character. Characters for which no mapping is specified are not modified by this conversion.

The `lcase` and `ucase` symbols are mutually exclusive.

<code>swab</code>	Swaps every pair of input bytes. If the current input record is an odd number of bytes, the last byte in the input record is ignored.
<code>noerror</code>	Does not stop processing on an input error. When an input error occurs, a diagnostic message is written on standard error, followed by the current input and output block counts in the same format as used at completion. If the <code>sync</code> conversion is specified, the missing input is replaced with null bytes and processed normally. Otherwise, the input block will be omitted from the output.
<code>notrunc</code>	Does not truncate the output file. Preserves blocks in the output file not explicitly written by this invocation of <code>dd</code> . (See also the preceding <code>of=file</code> operand.)
<code>sync</code>	Pads every input block to the size of the <code>ibs=</code> buffer, appending null bytes. (If either <code>block</code> or <code>unblock</code> is also specified, appends SPACE characters, rather than null bytes.)

If operands other than `conv=` are specified more than once, the last specified operand=*value* is used.

For the `bs=`, `cbs=`, `ibs=`, and `obs=` operands, the application must supply an expression specifying a size in bytes. The expression, `expr`, can be:

- 1. a positive decimal number
- 2. a positive decimal number followed by `k`, specifying multiplication by 1024
- 3. a positive decimal number followed by `b`, specifying multiplication by 512
- 4. two or more positive decimal numbers (with or without `k` or `b`) separated by `x`, specifying the product of the indicated values.

All of the operands will be processed before any input is read.

用法

See [largefile\(5\)](#) for the description of the behavior of `dd` when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).

示例

示例 1 Copying from one tape drive to another

The following example copies from tape drive `0` to tape drive `1`, using a common historical device naming convention.

示例 1 Copying from one tape drive to another (续)

```
example% dd if=/dev/rmt/0h of=/dev/rmt/1h
```

示例 2 Stripping the first 10 bytes from standard input

The following example strips the first 10 bytes from standard input:

```
example% dd ibs=10 skip=1
```

示例 3 Reading a tape into an ASCII file

This example reads an EBCDIC tape blocked ten 80-byte EBCDIC card images per block into the ASCII file x:

```
example% dd if=/dev/tape of=x ibs=800 cbs=80 conv=ascii,lcase
```

示例 4 Using conv=sync to write to tape

The following example uses conv=sync when writing to a tape:

```
example% tar cvf - . | compress | dd obs=1024k of=/dev/rmt/0 conv=sync
```

环境变量

See [environ\(5\)](#) for descriptions of the following environment variables that affect the execution of dd: LANG, LC_ALL, LC_CTYPE, LC_MESSAGES, and NLSPATH.

退出状态

The following exit values are returned:

- 0 The input file was copied successfully.
- >0 An error occurred.

If an input error is detected and the noerror conversion has not been specified, any partial output block will be written to the output file, a diagnostic message will be written, and the copy operation will be discontinued. If some other error is detected, a diagnostic message will be written and the copy operation will be discontinued.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed
Standard	See standards(5) .

另请参见

[cp\(1\)](#), [sed\(1\)](#), [tr\(1\)](#), [attributes\(5\)](#), [environ\(5\)](#), [largefile\(5\)](#), [standards\(5\)](#)

诊断

f+p records in(out) numbers of full and partial blocks read(written)

附注

Do not use dd to copy files between file systems having different block sizes.

Using a blocked device to copy a file will result in extra nulls being added to the file to pad the final block to the block boundary.

When dd reads from a pipe, using the `ibs=X` and `obs=Y` operands, the output will always be blocked in chunks of size Y. When `bs=Z` is used, the output blocks will be whatever was available to be read from the pipe at the time.

When using dd to copy files to a tape device, the file size must be a multiple of the device sector size (for example, 512 Kbyte). To copy files of arbitrary size to a tape device, use [tar\(1\)](#) or [cpio\(1\)](#).

For SIGINT, dd writes status information to standard error before exiting. It takes the standard action for all other signals.

引用名	ddu – 基于 GUI 的设备驱动实用程序
用法概要	<code>/usr/bin/ddu [--silent]</code>
描述	基于 GUI 的设备驱动实用程序 <code>ddu</code> 提供有关运行 Oracle Solaris 操作系统的系统上的设备信息。使用设备驱动实用程序，您可以连接到映像包管理系统 (Image Packaging System, IPS)，并为没有连接驱动程序的设备搜索设备驱动程序。这样，通过 <code>ddu</code>，用户就有机会安装缺少的驱动程序。 <code>ddu</code> 是具有易于使用的界面的 GUI。它有一个基于文本的对应项 ddu-text(1M)，相关信息可参见其手册页。
选项	支持以下选项： --静音 <code>ddu</code> 以无提示方式运行。如果它发现缺少驱动程序的设备，会发布一条通知。单击该通知将显示 DDU GUI。
示例	示例 1 启动 ddu GUI 以下命令可启动 <code>ddu</code> GUI。 <code>% ddu</code> 示例 2 以无提示模式启动 ddu 以下命令以无提示模式启动 <code>ddu</code>。 <code>% ddu --silent</code>
退出状态	<code>0</code> 应用程序成功退出。 <code>>0</code> 应用程序退出失败。
文件	<code>/usr/bin/ddu</code> <code>ddu</code> 的可执行文件。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	diagnostic/ddu
接口稳定性	External（外部）

其他可用性属性值包括：`diagnostic/ddu/data`、`diagnostic/ddu/library`和 `diagnostic/ddu/locale`。

另请参见

[ddu-text\(1M\)](#)、[attributes\(5\)](#)

引用名 ddu-text – 基于文本的设备驱动实用程序

用法概要 ddu-text

描述 基于文本的设备驱动实用程序 `ddu-text` 提供有关运行 Oracle Solaris 操作系统的系统上的设备信息。使用设备驱动实用程序，您可以连接到映像包管理系统 (Image Packaging System, IPS)，并为没有连接驱动程序的设备搜索设备驱动程序。这样，通过 `ddu-text`，用户就有机会安装缺少的驱动程序。

`ddu-text` 是基于文本的。它有一个 GUI 对应项 [ddu\(1M\)](#)，相关信息可参见其手册页。

`ddu-text` 没有命令行选项。

示例 示例1 调用 `ddu-text`

以下命令调用 `ddu-text`。

退出状态 % `ddu-text`

0 应用程序成功退出。

>0 应用程序退出失败。

文件 `/usr/bin/ddu-text`
`ddu-text` 的可执行文件。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	diagnostic/ddu
接口稳定性	External（外部）

其他可用性属性值包括：`diagnostic/ddu/data`、`diagnostic/ddu/library`和 `diagnostic/ddu/locale`。

另请参见 [ddu\(1M\)](#)、[attributes\(5\)](#)

引用名	devchassisd – devchassis daemon
用法概要	<code>/usr/lib/devchassis/devchassisd [-d]</code>
描述	The devchassis daemon, <code>devchassisd</code>, provides user-level services for the management of the <code>/dev/chassis</code> name space. It is a system daemon started by the Service Management Facility (see smf(5)). Its fault management resource identifier (FMRI) is: <pre>svc:/system/devchassis:daemon</pre> Note that <code>devchassisd</code> is a Consolidation Private interface. See attributes(5).
选项	The following option is supported: <code>-d, --debug</code> Run the daemon in standalone debug mode. Messages will be displayed on the controlling terminal instead of to <code>syslog</code>. And increased verbosity will be enabled to display more details about the internal operations of the daemon.
示例	示例 1 Enabling the devchassis Service The following command enables the <code>devchassis</code> service: <pre># svcadm enable svc:/system/devchassis</pre> 示例 2 Disabling the devchassis Service The following command disables the <code>devchassis</code> service: <pre># svcadm disable svc:/system/devchassis</pre>
错误	The <code>devchassisd</code> daemon uses syslog(3C) to report status and error messages. All of the messages are logged with the <code>LOG_DAEMON</code> facility. Error messages are logged with the <code>LOG_ERR</code> and <code>LOG_NOTICE</code> priorities, and informational messages are logged with the <code>LOG_INFO</code> priority. The default entries in the <code>/etc/syslog.conf</code> file log all of the <code>devchassisd</code> daemon error messages to the <code>/var/adm/messages</code> log.
文件	<code>/usr/lib/devchassis/devchassisd</code> devchassisd daemon binary.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Consolidation Private

另请参见

[svcs\(1\)](#), [svcadm\(1M\)](#), [syslog\(3C\)](#), [syslog.conf\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [devchassis\(7FS\)](#)

附注

The `devchassis` service is managed by the service management facility, [smf\(5\)](#), under the fault management resource identifier (FMRI):

```
svc:/system/devchassis
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

The `devchassis` service is enabled by default.

引用名	devfsadm, devfsadmd – administration command for /dev
用法概要	<pre>/usr/sbin/devfsadm [-C] [-c <i>device_class</i>] [-i <i>driver_name</i>] [-n] [-r <i>root_dir</i>] [-s] [-t <i>table_file</i>] [-u] [-v] /usr/lib/devfsadm/devfsadmd</pre>
描述	devfsadm maintains the /dev namespace. It replaces the previous suite of devfs administration tools including drvconfig(1M), disks(1M), tapes(1M), ports(1M), audlinks(1M), and devlinks(1M). The default operation is to attempt to load every driver in the system and attach to all possible device instances. Next, devfsadm creates logical links to device nodes in /dev and /devices and loads the device policy. devfsadmd(1M) is the daemon version of devfsadm(1M). The daemon is started during system startup and is responsible for handling both reconfiguration boot processing and updating /dev and /devices in response to dynamic reconfiguration event notifications from the kernel. For compatibility purposes, drvconfig(1M), disks(1M), tapes(1M), ports(1M), audlinks(1M), and devlinks(1M) are implemented as links to devfsadm. In addition to managing /dev, devfsadm also maintains the path_to_inst(4) database.
选项	The following options are supported: -C Cleanup mode. Prompt devfsadm to cleanup dangling /dev links that are not normally removed. If the -c option is also used, devfsadm only cleans up for the listed devices' classes. -c <i>device_class</i> Restrict operations to devices of class <i>device_class</i>. Solaris defines the following values for <i>device_class</i>: disk, tape, port, audio, and pseudo. This option might be specified more than once to specify multiple device classes. -i <i>driver_name</i> Configure only the devices for the named driver, <i>driver_name</i>. -n Do not attempt to load drivers or add new nodes to the kernel device tree. -s Suppress any changes to /dev. This is useful with the -v option for debugging. -t <i>table_file</i> Read an alternate devlink.tab file. devfsadm normally reads /etc/devlink.tab.

- u
Activate and attach devices for drivers added with [add_drv\(1M\)](#) -u. Cannot be used together with -n or -r.
- r *root_dir*
Presume that the /dev directory trees are found under *root_dir*, not directly under root (/). No other use or assumptions are made about *root_dir*.
- v
Print changes to /dev in verbose mode.

退出状态 The following exit values are returned:

- 0
Successful completion.
- 1
An error occurred.

文件

- /devices
device nodes directory
- /dev
logical symbolic links to /devices
- /usr/lib/devfsadm/devfsadmd
devfsadm daemon
- /dev/.devfsadm_dev.lock
update lock file
- /dev/.devfsadm_daemon.lock
daemon lock file
- /etc/security/device_policy
device policy file
- /etc/security/extra_privs
additional device privileges

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [svcs\(1\)](#), [add_drv\(1M\)](#), [modinfo\(1M\)](#), [modload\(1M\)](#), [modunload\(1M\)](#), [rem_drv\(1M\)](#), [svcadm\(1M\)](#), [tapes\(1M\)](#), [path_to_inst\(4\)](#), [attributes\(5\)](#), [privileges\(5\)](#), [smf\(5\)](#), [devfs\(7FS\)](#)

附注

This document does not constitute an API. The `/devices` directory might not exist or might have different contents or interpretations in a future release. The existence of this notice does not imply that any other documentation that lacks this notice constitutes an API.

`devfsadm` no longer manages the `/devices` name space. See [devfs\(7FS\)](#).

As a daemon to support hot-plug and synchronous device naming, `devfsadm` is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/system/devfsadm:default
```

The status of the service can be queried using the [svcs\(1\)](#) command.

引用名 device_allocate – enable and disable device allocation

用法概要 svc:/system/device/allocate:default

描述 The device_allocate service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/system/device/allocate:default

An administrator with the Devices Security Rights Profile can enable or disable this service to provide for or remove the device allocation functionality as described in [allocate\(1\)](#).

示例 示例 1 Enabling Device Allocation

The following command enables device allocation.

svcadm enable svc:/system/device/allocate:default

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Volatile

另请参见 [allocate\(1\)](#), [deallocate\(1\)](#), [list_devices\(1\)](#), [devfsadm\(1M\)](#), [svcadm\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

引用名	device_remap – administer the Solaris I/O remapping feature
用法概要	<code>/usr/platform/sun4v/sbin/device_remap [-v -R dir]</code>
描述	Certain multi-node sun4v platforms, such as T5440 and T5240 servers, have an integrated PCI topology that cause the I/O device paths to change in a CPU node failover condition. The device remapping script, <code>device_remap</code>, remaps the device paths in <code>/etc/path_to_inst</code> file and the symlinks under <code>/dev</code> to match the hardware.
选项	The following options are supported: -v Displays the <code>/etc/path_to_inst</code> and <code>/dev</code> symlink changes. -R <i>dir</i> Perform remapping on the <code>/etc/path_to_inst</code> and <code>/etc/path_to_inst</code> files in the root image at <i>dir</i>.
用法	The primary function of <code>device_remap</code> is to remap the device paths in the <code>/etc/path_to_inst</code> file and the symlinks under <code>/dev</code> in a CPU node failover condition to match the hardware. After adding CPU node(s) or removing CPU node(s), boot the system to the OBP prompt and use the following procedure: 1. Boot either the failsafe miniroot using: <code>boot -F failsafe</code>, or an install miniroot using <code>boot net -s</code> or similar command. 2. Mount the root disk as <code>/mnt</code>. 3. Change directory to the mounted root disk: <pre># cd /mnt</pre> 4. Run <code>device_remap</code> script: <pre># /mnt/usr/platform/sun4v/sbin/device_remap</pre> 5. Boot the system from disk. All the error messages are self-explanatory, except for the error message “missing ioaliases node” which means the firmware on the system does not support device remapping.
示例	示例 1 Displaying Changes Following Failover The following command displays the <code>path_to_inst</code> and <code>/dev</code> changes following a CPU node failover. <pre># device_remap -v</pre>

示例 2 Changing Directory Prior to Any Changes

The following command changes the directory on which the boot image is mounted prior to making any changes.

```
# device_remap -R /newroot
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/library/platform
Interface Stability	Uncommitted

另请参见

[boot\(1M\)](#), [attributes\(5\)](#)

引用名	devinfo – print device specific information
用法概要	<pre>/usr/sbin/devinfo -i device /usr/sbin/devinfo -p device</pre>
描述	The devinfo command is used to print device specific information about disk devices on standard out. The command can only be used by the superuser.
选项	-i Prints the following device information: ▪ Device name ▪ Software version (not supported and prints as 0) ▪ Drive id number (not supported and prints as 0) ▪ Device blocks per cylinder ▪ Device bytes per block ▪ Number of device partitions with a block size greater than zero -p Prints the following device partition information: ▪ Device name ▪ Device major and minor numbers (in hexadecimal) ▪ Partition start block ▪ Number of blocks allocated to the partition ▪ Partition flag ▪ Partition tag This command is used by various other commands to obtain device specific information for the making of file systems and determining partition information. If the device cannot be opened, an error message is reported.
操作数	<i>device</i> Device name.
退出状态	0 Successful operation. 2 Operation failed.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [prtvtoc\(1M\)](#), [attributes\(5\)](#)

引用名	devlinks – adds /dev entries for miscellaneous devices and pseudo-devices												
用法概要	<code>/usr/sbin/devlinks [-d] [-r rootdir] [-t table-file]</code>												
描述	devfsadm(1M) is now the preferred command for /dev and /devices and should be used instead of devlinks. devlinks creates symbolic links from the /dev directory tree to the actual block- and character-special device nodes under the /devices directory tree. The links are created according to specifications found in the <i>table-file</i> (by default /etc/devlink.tab). devlinks is called each time the system is reconfiguration-booted, and can only be run after drvconfig(1M) is run. The <i>table-file</i> (normally /etc/devlink.tab) is an ASCII file, with one line per record. Comment lines, which must contain a hash character ('#') as their first character, are allowed. Each entry must contain at least two fields, but may contain three fields. Fields are separated by single TAB characters. The fields are: <table><tr><td><i>devfs-spec</i></td><td>Specification of devinfo nodes that will have links created for them. This specification consists of one or more keyword-value pairs, where the keyword is separated from the value by an equal-sign ('='), and keyword-value pairs are separated from one another by semicolons.</td></tr><tr><td colspan="2">The possible keywords are:</td></tr><tr><td><i>type</i></td><td>The devinfo device type. Possible values are specified in ddi_create_minor_node(9F)</td></tr><tr><td><i>name</i></td><td>The name of the node. This is the portion of the /devices tree entry name that occurs before the first '@' or ':' character.</td></tr><tr><td><i>addr[n]</i></td><td>The address portion of a node name. This is the portion of a node name that occurs between the '@' and the ':' characters. It is possible that a node may have a name without an address part, which is the case for many of the pseudo-device nodes. If a number is given after the <i>addr</i> it specifies a match of a particular comma-separated subfield of the address field: <i>addr1</i> matches the first subfield, <i>addr2</i> matches the second, and so on. <i>addr0</i> is the same as <i>addr</i> and matches the whole field.</td></tr><tr><td><i>minor[n]</i></td><td>The minor portion of a node name – the portion of the name after the ':'. As with <i>addr</i> above, a number after the <i>minor</i> keyword specifies a subfield to match.</td></tr></table>	<i>devfs-spec</i>	Specification of devinfo nodes that will have links created for them. This specification consists of one or more keyword-value pairs, where the keyword is separated from the value by an equal-sign ('='), and keyword-value pairs are separated from one another by semicolons.	The possible keywords are:		<i>type</i>	The devinfo device type. Possible values are specified in ddi_create_minor_node(9F)	<i>name</i>	The name of the node. This is the portion of the /devices tree entry name that occurs before the first '@' or ':' character.	<i>addr[n]</i>	The address portion of a node name. This is the portion of a node name that occurs between the '@' and the ':' characters. It is possible that a node may have a name without an address part, which is the case for many of the pseudo-device nodes. If a number is given after the <i>addr</i> it specifies a match of a particular comma-separated subfield of the address field: <i>addr1</i> matches the first subfield, <i>addr2</i> matches the second, and so on. <i>addr0</i> is the same as <i>addr</i> and matches the whole field.	<i>minor[n]</i>	The minor portion of a node name – the portion of the name after the ':'. As with <i>addr</i> above, a number after the <i>minor</i> keyword specifies a subfield to match.
<i>devfs-spec</i>	Specification of devinfo nodes that will have links created for them. This specification consists of one or more keyword-value pairs, where the keyword is separated from the value by an equal-sign ('='), and keyword-value pairs are separated from one another by semicolons.												
The possible keywords are:													
<i>type</i>	The devinfo device type. Possible values are specified in ddi_create_minor_node(9F)												
<i>name</i>	The name of the node. This is the portion of the /devices tree entry name that occurs before the first '@' or ':' character.												
<i>addr[n]</i>	The address portion of a node name. This is the portion of a node name that occurs between the '@' and the ':' characters. It is possible that a node may have a name without an address part, which is the case for many of the pseudo-device nodes. If a number is given after the <i>addr</i> it specifies a match of a particular comma-separated subfield of the address field: <i>addr1</i> matches the first subfield, <i>addr2</i> matches the second, and so on. <i>addr0</i> is the same as <i>addr</i> and matches the whole field.												
<i>minor[n]</i>	The minor portion of a node name – the portion of the name after the ':'. As with <i>addr</i> above, a number after the <i>minor</i> keyword specifies a subfield to match.												

Of these four specifications, only the *type* specification must always be present.

name Specification of the /dev links that correspond to the devinfo nodes. This field allows devlinks to determine matching /dev names for the /devices nodes it has found. The specification of this field uses escape-sequences to allow portions of the /devices name to be included in the /dev name, or to allow a counter to be used in creating node names. If a counter is used to create a name, the portion of the name before the counter must be specified absolutely, and all names in the /dev/-subdirectory that match (up to and including the counter) are considered to be subdevices of the same device. This means that they should all point to the same directory, name and address under the /devices/-tree

The possible escape-sequences are:

- `\D` Substitute the device-name (name) portion of the corresponding devinfo node-name.
- `\An` Substitute the *n*th component of the address component of the corresponding devinfo node name. Sub-components are separated by commas, and sub-component 0 is the whole address component.
- `\Mn` Substitute the *n*th sub-component of the minor component of the corresponding devinfo node name. Sub-components are separated by commas, and sub-component 0 is the whole minor component.
- `\Nn` Substitute the value of a 'counter' starting at *n*. There can be only one counter for each dev-spec, and counter-values will be selected so they are as low as possible while not colliding with already-existing link names.

In a dev-spec the counter sequence should not be followed by a digit, either explicitly or as a result of another escape-sequence expansion. If this occurs, it would not be possible to correctly match already-existing links to their counter entries, since it would not be possible to unambiguously parse the already-existing /dev-name.

extra-dev-link Optional specification of an extra /dev link that points to the initial /dev link (specified in field 2). This field may contain a counter escape-sequence (as described for the *dev-spec* field) but may not contain any of the other escape-sequences. It provides a way to specify an alias of a particular /dev name.

选项

The following options are supported:

-d

Debugging mode – print out all devinfo nodes found, and indicate what links would be created, but do not do anything.

-r *rootdir*

Use *rootdir* as the root of the /dev and /devices directories under which the device nodes and links are created. Changing the root directory does not change the location of the /etc/devlink.tab default table, nor is the root directory applied to the filename supplied to the -t option.

-t *table-file*

Set the table file used by devlinks to specify the links that must be created. If this option is not given, /etc/devlink.tab is used. This option gives a way to instruct devlinks just to perform a particular piece of work, since just the links-types that devlinks is supposed to create can be specified in a command-file and fed to devlinks.

错误

If devlinks finds an error in a line of the *table-file* it prints a warning message on its standard output and goes on to the next line in the *table-file* without performing any of the actions specified by the erroneous rule.

If it cannot create a link for some filesystem-related reason it prints an error-message and continues with the current rule.

If it cannot read necessary data it prints an error message and continues with the next *table-file* line.

示例

示例 1 Using the /etc/devlink.tab Fields

The following are examples of the /etc/devlink.tab fields:

```
type=pseudo;name=win      win\M0
type=ddi_display  framebuffer/\M0    fb\N0
```

The first example states that all devices of type pseudo with a name component of win will be linked to /dev/win*x*, where *x* is the minor-component of the *devinfo-name* (this is always a single-digit number for the win driver).

The second example states that all devinfo nodes of type ddi_display will be linked to entries under the /dev/framebuffer directory, with names identical to the entire minor component of the /devices name. In addition an extra link will be created pointing from /dev/fb*n* to the entry under /dev/framebuffer. This entry will use a counter to end the name.

文件

/dev

entries for the miscellaneous devices for general use

/devices

device nodes

/etc/devlink.tab

the default rule-file

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[devfsadm\(1M\)](#), [attributes\(5\)](#), [devfs\(7FS\)](#), [ddi_create_minor_node\(9F\)](#)

已知问题

It is very easy to construct mutually-contradictory link specifications, or specifications that can never be matched. The program does not check for these conditions.

引用名 devnm – device name

用法概要 /usr/sbin/devnm *name* [*name*]...

描述 The devnm command identifies the special file associated with the mounted file system where the argument *name* resides. One or more *name* can be specified.

示例 示例1 Using the devnm Command

Assuming that /usr is mounted on /dev/dsk/c0t3d0s6, the following command :

/usr/sbin/devnm /usr

produces:

/dev/dsk/c0t3d0s6 /usr

文件 /dev/dsk/*

/etc/mnttab

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [mnttab\(4\)](#), [attributes\(5\)](#)

引用名	devprop – display device properties
用法概要	<code>/usr/sbin/devprop [-n <i>device-path</i>] [-vq] [-{b i l s}] [<i>property</i>...]</code>
描述	The devprop command displays named device properties from the device tree. If a device path is specified on the command line, devprop displays device properties for that device node. The output formats, while Volatile, are patterned after those of prtconf(1M).
选项	The following options are supported: -b The properties to be output are sequences of bytes (DI_PROP_TYPE_BYTES). -b, -i, -l, and -s are mutually exclusive of one another. -i The properties to be output are integers (DI_PROP_TYPE_INT). -l The properties to be output are 64-bit integers (DI_PROP_TYPE_INT64). -s The properties to be output are strings (DI_PROP_TYPE_STRING) (the default). -n <i>device-path</i> The path to a target device node for which properties are displayed. The default path is that of the root node (equivalent to specifying -n /). -q Specifies quoted output mode, in which string properties are output surrounded by double quotation marks ("). -v Specifies verbose mode, in which the name of the property is output before its value.
操作数	The following operand is supported: <i>property</i>... Name of the property to be displayed.
退出状态	0 No error occurred. non-zero An error occurred.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	See below.

The command invocation and output are both Volatile. The location of the utility is Committed.

另请参见 [prtconf\(1M\)](#), [libdevinfo\(3LIB\)](#), [attributes\(5\)](#)

引用名	<code>df</code> – displays number of free disk blocks and free files
用法概要	<pre>df [-F <i>FSType</i>] [-abeghklntPVvZ] [-o <i>FSType-specific_options</i>] [<i>block_device</i> <i>directory</i> <i>file</i> <i>resource</i> ...]</pre>
描述	The <code>df</code> utility displays the amount of disk space occupied by mounted or unmounted file systems, the amount of used and available space, and how much of the file system's total capacity has been used. The file system is specified by device, or by referring to a file or directory on the specified file system. Used without operands or options, <code>df</code> reports on all mounted file systems. <code>df</code> may not be supported for all <i>FSTypes</i>. If <code>df</code> is run on a networked mount point that the automounter has not yet mounted, the file system size will be reported as zero. As soon as the automounter mounts the file system, the sizes will be reported correctly.
选项	The following options are supported: -a Reports on all file systems including ones whose entries in <code>/etc/mnttab</code> (see mnttab(4)) have the <code>ignore</code> option set. -b Prints the total number of kilobytes free. -e Prints only the number of files free. -F <i>FSType</i> Specifies the <i>FSType</i> on which to operate. The <code>-F</code> option is intended for use with unmounted file systems. The <i>FSType</i> should be specified here or be determinable from <code>/etc/vfstab</code> (see vfstab(4)) by matching the <i>directory</i>, <i>block_device</i>, or <i>resource</i> with an entry in the table, or by consulting <code>/etc/default/fs</code>. See default_fs(4). -g Prints the entire statvfs(2) structure. This option is used only for mounted file systems. It can not be used with the <code>-o</code> option. This option overrides the <code>-b</code>, <code>-e</code>, <code>-k</code>, <code>-n</code>, <code>-P</code>, and <code>-t</code> options. -h Like <code>-k</code>, except that sizes are in a more human readable format. The output consists of one line of information for each specified file system. This information includes the file system name, the total space allocated in the file system, the amount of space allocated to existing files, the total amount of space available for the creation of new files by unprivileged users, and the percentage of normally available space that is currently allocated to all files on the file system. All sizes are scaled to a human readable format, for example, 14K, 234M, 2.7G, or 3.0T. Scaling is done by repetitively dividing by 1024.

This option overrides the `-b`, `-e`, `-g`, `-k`, `-n`, `-t`, and `-V` options. This option only works on mounted filesystems and can not be used together with `-o` option.

`-k`

Prints the allocation in kbytes. The output consists of one line of information for each specified file system. This information includes the file system name, the total space allocated in the file system, the amount of space allocated to existing files, the total amount of space available for the creation of new files by unprivileged users, and the percentage of normally available space that is currently allocated to all files on the file system. This option overrides the `-b`, `-e`, `-n`, and `-t` options and may not be used together with the `-v` option.

`-l`

Reports on local file systems only. This option is used only for mounted file systems. It can not be used with the `-o` option.

`-n`

Prints only the *FSType* name. Invoked with no operands, this option prints a list of mounted file system types. This option is used only for mounted file systems. It can not be used with the `-o` option.

`-o FSType-specific_options`

Specifies *FSType-specific* options. These options are comma-separated, with no intervening spaces. See the manual page for the *FSType-specific* command for details.

`-t`

Prints full listings with totals. This option overrides the `-b`, `-e`, and `-n` options.

`-P`

Same as `-h` except in 512-byte units.

`-V`

Echoes the complete set of file system specific command lines, but does not execute them. The command line is generated by using the options and operands provided by the user and adding to them information derived from `/etc/mnttab`, `/etc/vfstab`, or `/etc/default/fs`. This option may be used to verify and validate the command line.

`-v`

Like `-k`, except that sizes are displayed in multiples of the smallest block size supported by each specified file system.

This option may not be used with the `-k` option.

The output consists of one line of information for each file system. This one line of information includes the following:

- the file system's mount point
- the file system's name
- the total number of blocks allocated to the file system
- the number of blocks allocated to existing files

- the number of blocks available for the creation of new files by unprivileged users
- the percentage of blocks in use by files

-Z
Displays mounts in all visible zones. By default, df displays mounts located only within the current zone. This option has no effect in a non-global zone.

操作数

The df utility interprets operands according to the following precedence: *block_device*, *directory*, *file*, *resource*. The following operands are supported:

- block_device*
Represents a block special device (for example, /dev/dsk/c1d0s7).
- directory*
Represents a valid directory name. df reports on the file system that contains *directory*.
- file*
Represents a valid file name. df reports on the file system that contains *file*.
- resource*
Represents an NFS resource name.

用法

See [largefile\(5\)](#) for the description of the behavior of df when encountering files greater than or equal to 2 Gbyte (2³¹ bytes).

示例

示例 1 Executing the df command
The following example shows the df command and its output:

```
example% /usr/bin/df

/                (/dev/dsk/c0t0d0s0 ): 287530 blocks   92028 files
/system/contract (ctfs                ):      0 blocks 2147483572 files
/system/object   (objfs                ):      0 blocks 2147483511 files
/usr             (/dev/dsk/c0t0d0s6 ): 1020214 blocks 268550 files
/proc            (/proc                 ):      0 blocks    878 files
/dev/fd          (fd                   ):      0 blocks     0 files
/etc/mnttab      (mnttab                ):      0 blocks     0 files
/var/run         (swap                  ): 396016 blocks   9375 files
/tmp            (swap                  ): 396016 blocks   9375 files
/opt            (/dev/dsk/c0t0d0s5 ): 381552 blocks   96649 files
/export/home     (/dev/dsk/c0t0d0s7 ): 434364 blocks  108220 files
```

where the columns represent the mount point, device (or “filesystem”, according to df -k), free blocks, and free files, respectively. For contract file systems, /system/contract is the mount point, ctfs is the contract file system (used by SMF) with 0 free blocks and 2147483582(INTMAX-1) free files. For object file systems, /system/object is the mount point, objfs is the object file system (see [objfs\(7FS\)](#)) with 0 free blocks and 2147483511 free files.

示例 2 Writing Portable Information About the /usr File System

The following example writes portable information about the /usr file system:

```
example% /usr/bin/df -P /usr
```

示例 3 Writing Portable Information About the /usr/src file System

Assuming that /usr/src is part of the /usr file system, the following example writes portable information :

```
example% /usr/bin/df -P /usr/src
```

示例 4 Using df to Display Inode Usage

The following example displays inode usage on all ufs file systems:

```
example%/usr/bin/df -F ufs -o i
```

环境变量

When set, any header which normally displays files will now display nodes. See [environ\(5\)](#) for descriptions of the following environment variables that affect the execution of df: LANG, LC_ALL, LC_CTYPE, LC_MESSAGES, and NLSPATH.

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

文件

- /dev/dsk/* Disk devices
- /etc/default/fs Default local file system type. Default values can be set for the following flags in /etc/default/fs. For example: LOCAL=ufs, where LOCAL is the default partition for a command if no FSType is specified.
- /etc/mnttab Mount table
- /etc/vfstab List of default parameters for each file system

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Standard	See standards(5) .

另请参见 [find\(1\)](#), [df_ufs\(1M\)](#), [mount\(1M\)](#), [statvfs\(2\)](#), [default_fs\(4\)](#), [mnttab\(4\)](#), [vfstab\(4\)](#), [attributes\(5\)](#), [environ\(5\)](#), [largefile\(5\)](#), [standards\(5\)](#), [objfs\(7FS\)](#)

附注 If UFS logging is enabled on a file system, the disk space used for the log is reflected in the `df` report. The log is allocated from free blocks on the file system, and it is sized approximately 1 Mbyte per 1 Gbyte of file system, up to 256 Mbytes. The log size may be larger (up to a maximum of 512 Mbytes) depending on the number of cylinder groups present in the file system.

引用名	dfmounts – display mounted resource information								
用法概要	<pre>dfmounts [-F FSType] [-h] [-o specific_options] [restriction]...</pre>								
描述	dfmounts shows the local resources shared through a distributed file system <i>FSType</i> along with a list of clients that have the resource mounted. If <i>restriction</i> is not specified, dfmounts shows file systems that are currently shared on any NFS server. <i>specific_options</i> as well as the availability and semantics of <i>restriction</i> are specific to particular distributed file system types. If dfmounts is entered without arguments, remote resources currently mounted on the local system are displayed, regardless of file system type. However, the dfmounts command does not display the names of NFS Version 4 clients.								
dfmounts Output	The output of dfmounts consists of an optional header line (suppressed with the -h flag) followed by a list of lines containing whitespace-separated fields. For each resource, the fields are: <pre>resource server pathname clients ...</pre> where: <table><tr><td><i>resource</i></td><td>Specifies the resource name that must be given to the mount(1M) command.</td></tr><tr><td><i>server</i></td><td>Specifies the system from which the resource was mounted.</td></tr><tr><td><i>pathname</i></td><td>Specifies the pathname that must be given to the share(1M) command.</td></tr><tr><td><i>clients</i></td><td>Is a comma-separated list of systems that have mounted the resource. Clients are listed in the form <i>domain.</i>, <i>domain.system</i>, or <i>system</i>, depending on the file system type.</td></tr></table> A field can be null. Each null field is indicated by a hyphen (–) unless the remainder of the fields on the line are also null, in which case the hyphen can be omitted. Fields with whitespace are enclosed in quotation marks (" ").	<i>resource</i>	Specifies the resource name that must be given to the mount(1M) command.	<i>server</i>	Specifies the system from which the resource was mounted.	<i>pathname</i>	Specifies the pathname that must be given to the share(1M) command.	<i>clients</i>	Is a comma-separated list of systems that have mounted the resource. Clients are listed in the form <i>domain.</i> , <i>domain.system</i> , or <i>system</i> , depending on the file system type.
<i>resource</i>	Specifies the resource name that must be given to the mount(1M) command.								
<i>server</i>	Specifies the system from which the resource was mounted.								
<i>pathname</i>	Specifies the pathname that must be given to the share(1M) command.								
<i>clients</i>	Is a comma-separated list of systems that have mounted the resource. Clients are listed in the form <i>domain.</i> , <i>domain.system</i> , or <i>system</i> , depending on the file system type.								
选项	<table><tr><td>-F <i>FSType</i></td><td>Specify filesystem type. Defaults to the first entry in <code>/etc/dfs/fstypes</code>. Note: currently the only valid <i>FSType</i> is <code>nfs</code>.</td></tr><tr><td>-h</td><td>Suppress header line in output.</td></tr><tr><td>-o <i>specific_options</i></td><td>Specify options specific to the filesystem provided by the -F option. Note: currently no options are supported.</td></tr></table>	-F <i>FSType</i>	Specify filesystem type. Defaults to the first entry in <code>/etc/dfs/fstypes</code> . Note: currently the only valid <i>FSType</i> is <code>nfs</code> .	-h	Suppress header line in output.	-o <i>specific_options</i>	Specify options specific to the filesystem provided by the -F option. Note: currently no options are supported.		
-F <i>FSType</i>	Specify filesystem type. Defaults to the first entry in <code>/etc/dfs/fstypes</code> . Note: currently the only valid <i>FSType</i> is <code>nfs</code> .								
-h	Suppress header line in output.								
-o <i>specific_options</i>	Specify options specific to the filesystem provided by the -F option. Note: currently no options are supported.								
文件	<code>/etc/dfs/fstypes</code> file system types								
属性	See attributes(5) for descriptions of the following attributes:								

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [dfshares\(1M\)](#), [mount\(1M\)](#), [share\(1M\)](#), [unshare\(1M\)](#), [attributes\(5\)](#)

引用名	dfmounts_nfs – display mounted NFS resource information
用法概要	dfmounts [-F nfs] [-h] [server]...
描述	dfmounts shows the local resources shared through NFS, along with the list of clients that have mounted the resource. The -F flag may be omitted if NFS is the only file system type listed in the file /etc/dfs/fstypes. dfmounts without options, displays all remote resources mounted on the local system, regardless of file system type. The output of dfmounts consists of an optional header line (suppressed with the -h flag) followed by a list of lines containing whitespace-separated fields. For each resource, the fields are: <i>resource server pathname clients ...</i> where <i>resource</i> Does not apply to NFS. Printed as a hyphen (-). <i>server</i> Specifies the system from which the resource was mounted. <i>pathname</i> Specifies the pathname that must be given to the share(1M) command. <i>clients</i> Is a comma-separated list of systems that have mounted the resource.
选项	-F nfs Specifies the nfs-FSType. -h Suppress header line in output. <i>server</i> Displays information about the resources mounted from each server, where <i>server</i> can be any system on the network. If no server is specified, the <i>server</i> is assumed to be the local system.
文件	/etc/dfs/fstypes
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/nfs

另请参见 [mount\(1M\)](#), [share\(1M\)](#), [unshare\(1M\)](#), [attributes\(5\)](#)

引用名	dfshares – list available resources from remote or local systems
用法概要	dfshares [-F <i>FSType</i>] [-h] [-o <i>specific_options</i>] [<i>server</i>] . . .
描述	dfshares provides information about resources available to the host through a distributed file system of type <i>FSType</i>. <i>specific_options</i> as well as the semantics of <i>server</i> are specific to particular distributed file systems. If dfshares is entered without arguments, all resources currently shared on the local system are displayed, regardless of file system type. The output of dfshares consists of an optional header line (suppressed with the -h flag) followed by a list of lines containing whitespace-separated fields. For each resource, the fields are: <i>resource server access transport</i> where <i>resource</i> Specifies the resource name that must be given to the mount(1M) command. <i>server</i> Specifies the name of the system that is making the resource available. <i>access</i> Specifies the access permissions granted to the client systems, either ro (for read-only) or rw (for read/write). If dfshares cannot determine access permissions, a hyphen (–) is displayed. <i>transport</i> Specifies the transport provider over which the resource is shared. A field may be null. Each null field is indicated by a hyphen (–) unless the remainder of the fields on the line are also null; in which case, the hyphen may be omitted.
选项	-F <i>FSType</i> Specify filesystem type. Defaults to the first entry in /etc/dfs/fstypes. -h Suppress header line in output. -o <i>specific_options</i> Specify options specific to the filesystem provided by the -F option.
文件	/etc/dfs/fstypes
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [dfmounts\(1M\)](#), [mount\(1M\)](#), [share\(1M\)](#), [unshare\(1M\)](#), [attributes\(5\)](#)

引用名

dfshares_nfs – list available NFS resources from remote systems

用法概要

dfshares [-F nfs] [-h] [server]...

描述

dfshares provides information about resources available to the host through NFS. The -F flag may be omitted if NFS is the first file system type listed in the file /etc/dfs/fstypes.

The query may be restricted to the output of resources available from one or more servers.

dfshares without arguments displays all resources shared on the local system, regardless of file system type.

Specifying *server* displays information about the resources shared by each server. *Server* can be any system on the network. If no server is specified, then *server* is assumed to be the local system.

The output of dfshares consists of an optional header line (suppressed with the -h flag) followed by a list of lines containing whitespace-separated fields. For each resource, the fields are:

```
resource server access transport
```

where

<i>resource</i>	Specifies the resource name that must be given to the mount(1M) command.
<i>server</i>	Specifies the system that is making the resource available.
<i>access</i>	Specifies the access permissions granted to the client systems; however, dfshares cannot determine this information for an NFS resource and populates the field with a hyphen (-).
<i>transport</i>	Specifies the transport provider over which the <i>resource</i> is shared; however, dfshares cannot determine this information for an NFS resource and populates the field with a hyphen (-).

A field may be null. Each null field is indicated by a hyphen (-) unless the remainder of the fields on the line are also null; in which case, the hyphen may be omitted.

选项

-F nfs

Specify the NFS file system type

-h

Suppress header line in output.

文件

/etc/dfs/fstypes

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/nfs

另请参见

[mount\(1M\)](#), [share\(1M\)](#), [unshare\(1M\)](#), [attributes\(5\)](#)

引用名	df_ufs – report free disk space on ufs file systems						
用法概要	df -F ufs [<i>generic_options</i>] [-o <i>i</i>] [<i>directory</i> <i>special</i>]						
描述	df displays the amount of disk space occupied by ufs file systems, the amount of used and available space, and how much of the file system's total capacity has been used. The amount of space reported as used and available is less than the amount of space in the file system; this is because the system reserves a fraction of the space in the file system to allow its file system allocation routines to work well. The amount reserved is typically about 10%; this can be adjusted using tunefs(1M). When all the space on the file system except for this reserve is in use, only the superuser can allocate new files and data blocks to existing files. When the file system is overallocated in this way, df might report that the file system is more than 100% utilized. If neither <i>directory</i> nor <i>special</i> is specified, df displays information for all mounted ufs file systems.						
选项	The following options are supported: <table><tr><td><i>generic_options</i></td><td>Options supported by the generic df command. See df(1M) for a description of these options.</td></tr><tr><td>-o</td><td>Specify ufs file system specific options. The available option is:</td></tr><tr><td> i</td><td>Report the number of used and free inodes. This option can not be used with <i>generic_options</i>.</td></tr></table>	<i>generic_options</i>	Options supported by the generic df command. See df(1M) for a description of these options.	-o	Specify ufs file system specific options. The available option is:	i	Report the number of used and free inodes. This option can not be used with <i>generic_options</i> .
<i>generic_options</i>	Options supported by the generic df command. See df(1M) for a description of these options.						
-o	Specify ufs file system specific options. The available option is:						
i	Report the number of used and free inodes. This option can not be used with <i>generic_options</i> .						
文件	/etc/mnttab list of file systems currently mounted						
属性	See attributes(5) for descriptions of the following attributes:						

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os, system/xopen/xcu4

另请参见	df(1M) , fsck(1M) , fstyp(1M) , tunefs(1M) , mnttab(4) , attributes(5) , ufs(7FS) ,
附注	df calculates its results differently for mounted and unmounted file systems. For unmounted systems, the numbers reflect the 10% reservation. This reservation is not reflected in df output for mounted file systems. For this reason, the available space reported by the generic command can differ from the available space reported by this module. df might report remaining capacity even though syslog warns filesystem full. This issue can occur because df only uses the available fragment count to calculate available space, but the file system requires contiguous sets of fragments for most allocations. If you suspect that you have exhausted contiguous fragments on your file system, you can use the fstyp(1M) utility with the -v option. In the fstyp output, look at the nb free (number of blocks free) and nf free (number of fragments free) fields. On unmounted filesystems, you can use fsck(1M) and observe the last line of output, which reports, among other items, the number of fragments and the degree of fragmentation. See fsck(1M).

引用名	dhcpagent – Dynamic Host Configuration Protocol (DHCP) client daemon
用法概要	dhcpagent [-a] [-d <i>n</i>] [-f] [-v]
描述	dhcpagent implements the client half of the Dynamic Host Configuration Protocol (DHCP) for machines running Solaris software. The dhcpagent daemon obtains configuration parameters for the client (local) machine's network interfaces from a DHCP server. These parameters may include a lease on an IP address, which gives the client machine use of the address for the period of the lease, which may be infinite. If the client wishes to use the IP address for a period longer than the lease, it must negotiate an extension using DHCP. For this reason, dhcpagent must run as a daemon, terminating only when the client machine powers down. For IPv4, the dhcpagent daemon is controlled through ifconfig(1M). dhcpagent can be invoked as a user process, albeit one requiring root privileges, but this is not necessary, as ifconfig(1M) will start it automatically. For IPv6, the dhcpagent daemon is invoked automatically by in.ndpd(1M). It can also be controlled through ifconfig(1M), if necessary. When invoked, dhcpagent enters a passive state while it awaits instructions from ifconfig(1M) or in.ndpd(1M). When it receives a command to configure an interface, it brings up the interface (if necessary) and starts DHCP. Once DHCP is complete, dhcpagent can be queried for the values of the various network parameters. In addition, if DHCP was used to obtain a lease on an address for an interface, it configures the address for use. When a lease is obtained, it is automatically renewed as necessary. If the lease cannot be renewed, dhcpagent will unconfigure the address, but the interface will be left up and dhcpagent will attempt to acquire a new address lease. dhcpagent monitors system suspend/resume events and will validate any non-permanent leases with the DHCP server upon resume. Similarly, dhcpagent monitors link up/down events and will validate any non-permanent leases with the DHCP server when the downed link is brought back up. The lease validation mechanism will restart DHCP if the server indicates that the existing lease is no longer valid. If the server cannot be contacted, then the existing lease will continue. This behavior can be modified with the VERIFIED_LEASE_ONLY parameter in the /etc/default/dhcpagent file. See the description of this parameter below. For IPv4, if the configured interface is found to be unplumbed, or to have a different IP address, subnet mask, or broadcast address from those obtained from DHCP, the interface is abandoned from DHCP control. For IPv6, dhcpagent automatically plumbs and unplumbs logical interfaces as necessary for the IPv6 addresses supplied by the server. The IPv6 prefix length (netmask) is not set by the DHCPv6 protocol, but is instead set by in.ndpd(1M) using prefix information obtained by Router Advertisements. If any of the logical interfaces created by dhcpagent is unplumbed, or

configured with a different IP address, it will be abandoned from DHCP control. If the link-local interface is unplumbed, then all addresses configured by DHCP on that physical interface will be removed.

In addition to DHCP, dhcpgent also supports BOOTP (IPv4 only). See RFC 951, Bootstrap Protocol. Configuration parameters obtained from a BOOTP server are treated identically to those received from a DHCP server, except that the IP address received from a BOOTP server always has an infinite lease.

DHCP also acts as a mechanism to configure other information needed by the client, for example, the domain name and addresses of routers. Aside from the IP address, and for IPv4 alone, the netmask, broadcast address, and default router, the agent does not directly configure the workstation, but instead acts as a database which may be interrogated by other programs, and in particular by [dhcplib\(1\)](#).

On clients with a single interface, this is quite straightforward. Clients with multiple interfaces may present difficulties, as it is possible that some information arriving on different interfaces may need to be merged, or may be inconsistent. Furthermore, the configuration of the interfaces is asynchronous, so requests may arrive while some or all of the interfaces are still unconfigured. To handle these cases, one interface may be designated as primary, which makes it the authoritative source for the values of DHCP parameters in the case where no specific interface is requested. See [dhcplib\(1\)](#) and [ifconfig\(1M\)](#) for details.

All DHCP packets sent by dhcpgent include a vendor class identifier (RFC 2132, option code 60; RFC 3315, option code 16). This identifier is the same as the platform name returned by the `uname -i` command, except:

- Any commas in the platform name are changed to periods.
- If the name does not start with a stock symbol and a comma, it is automatically prefixed with `SUNW`.

Messages

The dhcpgent daemon writes information and error messages in five categories:

critical

Critical messages indicate severe conditions that prevent proper operation.

errors

Error messages are important, sometimes unrecoverable events due to resource exhaustion and other unexpected failure of system calls; ignoring errors may lead to degraded functionality.

warnings

Warnings indicate less severe problems, and in most cases, describe unusual or incorrect datagrams received from servers, or requests for service that cannot be provided.

informational

Informational messages provide key pieces of information that can be useful to debugging a DHCP configuration at a site. Informational messages are generally controlled by the `-v` option. However, certain critical pieces of information, such as the IP address obtained, are always provided.

debug

Debugging messages, which may be generated at two different levels of verbosity, are chiefly of benefit to persons having access to source code, but may be useful as well in debugging difficult DHCP configuration problems. Debugging messages are only generated when using the `-d` option.

When `dhcpcagent` is run without the `-f` option, all messages are sent to the system logger `syslog(3C)` at the appropriate matching priority and with a facility identifier `LOG_DAEMON`. When `dhcpcagent` is run with the `-f` option, all messages are directed to standard error.

DHCP Events and User-Defined Actions

If an executable (binary or script) is placed at `/etc/dhcp/eventhook`, the `dhcpcagent` daemon will automatically run that program when any of the following events occur:

BOUND and BOUND6

These events occur during interface configuration. The event program is invoked when `dhcpcagent` receives the DHCPv4 ACK or DHCPv6 Reply message from the DHCP server for the lease request of an address, indicating successful initial configuration of the interface. (See also the `INFORM` and `INFORM6` events, which occur when configuration parameters are obtained without address leases.)

EXTEND and EXTEND6

These events occur during lease extension. The event program is invoked just after `dhcpcagent` receives the DHCPv4 ACK or DHCPv6 Reply from the DHCP server for the DHCPv4 REQUEST (renew) message or the DHCPv6 Renew or Rebind message.

Note that with DHCPv6, the server might choose to remove some addresses, add new address leases, and ignore (allow to expire) still other addresses in a given Reply message. The `EXTEND6` event occurs when a Reply is received that leaves one or more address leases still valid, even if the Reply message does not extend the lease for any address. The event program is invoked just before any addresses are removed, but just after any new addresses are added. Those to be removed will be marked with the `IFF_DEPRECATED` flag.

EXPIRE and EXPIRE6

These events occur during lease expiration. For DHCPv4, the event program is invoked just before the leased address is removed from an interface. For DHCPv6, the event program is invoked just before the last remaining leased addresses are removed from the interface.

DROP and DROP6

These events occur during the period when an interface is dropped. The event program is invoked just before the interface is removed from DHCP control. If the interface has been

abandoned due the user unplugging the interface, then this event will occur after the user's action has taken place. The interface might not be present.

INFORM and INFORM6

These events occur when an interface acquires new or updated configuration information from a DHCP server by means of the DHCPv4 INFORM or the DHCPv6 Information-Request message. These messages are sent using an `ifconfig(1M)` `dhcp inform` command or when the DHCPv6 Router Advertisement 0 (letter 0) bit is set and the M bit is not set. Thus, these events occur when the DHCP client does not obtain an IP address lease from the server, and instead obtains only configuration parameters.

LOSS6

This event occurs during lease expiration when one or more valid leases still remain. The event program is invoked just before expired addresses are removed. Those being removed will be marked with the `IFF_DEPRECATED` flag.

Note that this event is not associated with the receipt of the Reply message, which occurs only when one or more valid leases remain, and occurs only with DHCPv6. If all leases have expired, then the EXPIRE6 event occurs instead.

RELEASE and RELEASE6

This event occurs during the period when a leased address is released. The event program is invoked just before `dhcpgent` relinquishes the address on an interface and sends the DHCPv4 RELEASE or DHCPv6 Release packet to the DHCP server.

The system does not provide a default event program. The file `/etc/dhcp/eventhook` is expected to be owned by root and have a mode of 755.

The event program will be passed two arguments, the interface name and the event name, respectively. For DHCPv6, the interface name is the name of the physical interface.

The event program can use the `dhcpinf(1)` utility to fetch additional information about the interface. While the event program is invoked on every event defined above, it can ignore those events in which it is not interested. The event program runs with the same privileges and environment as `dhcpgent` itself, except that `stdin`, `stdout`, and `stderr` are redirected to `/dev/null`. Note that this means that the event program runs with root privileges.

If an invocation of the event program does not exit after 55 seconds, it is sent a `SIGTERM` signal. If does not exit within the next three seconds, it is terminated by a `SIGKILL` signal.

See EXAMPLES for an example event program.

选项

The following options are supported:

-a

Adopt a configured IPv4 interface. This option is for use with diskless DHCP clients. In the case of diskless DHCP, DHCP has already been performed on the network interface

providing the operating system image prior to running `dhcpagent`. This option instructs the agent to take over control of the interface. It is intended primarily for use in boot scripts.

The effect of this option depends on whether the interface is being adopted.

If the interface is being adopted, the following conditions apply:

`dhcpagent` uses the client id specified in `/chosen:<client_id>`, as published by the PROM or as specified on a `boot(1M)` command line. If this value is not present, the client id is undefined. The DHCP server then determines what to use as a client id. It is an error condition if the interface is an Infiniband interface and the PROM value is not present.

If the interface is not being adopted:

`dhcpagent` uses the value stored in `/etc/default/dhcpagent`. If this value is not present, the client id is undefined. If the interface is Infiniband and there is no value in `/etc/default/dhcpagent`, a client id is generated as described by the draft document on DHCP over Infiniband, available at:

<http://www.ietf.org>

`-d n`

Set debug level to *n*. Two levels of debugging are currently available, 1 and 2; the latter is more verbose.

`-f`

Run in the foreground instead of as a daemon process. When this option is used, messages are sent to standard error instead of to `syslog(3C)`.

`-v`

Provide verbose output useful for debugging site configuration problems.

示例

示例 1 Example Event Program

The following script is stored in the file `/etc/dhcp/eventhook`, owned by root with a mode of 755. It is invoked upon the occurrence of the events listed in the file.

```
#!/bin/sh

(
echo "Interface name: " $1
echo "Event: " $2

case $2 in
"BOUND")
    echo "Address acquired from server \"
        '/usr/sbin/dhccpinfo -i $1 ServerID'
    ;;
"BOUND6")
```

示例1 Example Event Program (续)

```

        echo "Addresses acquired from server " \
            '/usr/sbin/dhccpinfo -v6 -i $1 ServerID'
        ;;
    "EXTEND")
        echo "Lease extended for " \
            '/usr/sbin/dhccpinfo -i $1 LeaseTime' seconds"
        ;;
    "EXTEND6")
        echo "New lease information obtained on $i"
        ;;
    "EXPIRE" | "DROP" | "RELEASE")
        ;;
esac
) >/var/run/dhcp_eventhook_output 2>&1

```

Note the redirection of stdout and stderr to a file.

文件

/etc/dhcp/if.dhc

/etc/dhcp/if.dh6

Contains the configuration for interface. The mere existence of this file does not imply that the configuration is correct, since the lease might have expired. On start-up, dhcpageant confirms the validity of the address using REQUEST (for DHCPv4) or Confirm (DHCPv6).

/etc/dhcp/duid

/etc/dhcp/iaid

Contains persistent storage for DUID (DHCP Unique Identifier) and IAID (Identity Association Identifier) values. The format of these files is undocumented, and applications should not read from or write to them.

/etc/default/dhcpageant

Contains default values for tunable parameters. All values may be qualified with the interface they apply to by prepending the interface name and a period (".") to the interface parameter name. The parameters include: the interface parameter name.

To configure IPv6 parameters, place the string .v6 between the interface name (if any) and the parameter name. For example, to set the global IPv6 parameter request list, use .v6.PARAM_REQUEST_LIST. To set the CLIENT_ID (DUID) on hme0, use hme0.v6.CLIENT_ID.

The parameters include:

VERIFIED_LEASE_ONLY

Indicates that a RELEASE rather than a DROP should be performed on managed interfaces when the agent terminates. Release causes the client to discard the lease, and the server to make the address available again. Drop causes the client to record the lease in

`/etc/dhcp/interface.dhc` or `/etc/dhcp/interface.dh6` for later use. In addition, when the link status changes to up or when the system is resumed after a suspend, the client will verify the lease with the server. If the server is unreachable for verification, then the old lease will be discarded (even if it has time remaining) and a new one obtained.

Enabling this option is often desirable on mobile systems, such as laptops, to allow the system to recover quickly from moves.

OFFER_WAIT

Indicates how long to wait between checking for valid OFFERS after sending a DISCOVER. For DHCPv6, sets the time to wait between checking for valid Advertisements after sending a Solicit.

CLIENT_ID

Indicates the value that should be used to uniquely identify the client to the server. This value can take one of three basic forms:

decimal, data...

0xHHHHH...

"string..."

The first form is an RFC 3315 DUID. This is legal for both IPv4 DHCP and DHCPv6. For IPv4, an RFC 4361 Client ID is constructed from this value. In this first form, the format of *data...* depends on the decimal value. The following formats are defined for this first form:

1,*hwtype,time,lla*

Type 1, DUID-LLT. The *hwtype* value is an integer in the range 0-65535, and indicates the type of hardware. The *time* value is the number of seconds since midnight, January 1st, 2000 UTC, and can be omitted to use the current system time. The *lla* value is either a colon-separated MAC address or the name of a physical interface. If the name of an interface is used, the *hwtype* value can be omitted. For example: 1, , , hme0

2,*enterprise,hex...*

Type 2, DUID-EN. The *enterprise* value is an integer in the range 0-4294967295 and represents the SMI Enterprise number for an organization. The *hex* string is an even-length sequence of hexadecimal digits.

3,*hwtype,lla*

Type 3, DUID-LL. This is the same as DUID-LLT (type 1), except that a time stamp is not used.

**,hex*

Any other type value (0 or 4-65535) can be used with an even-length hexadecimal string.

The second and third forms of CLIENT_ID are legal for IPv4 only. These both represent raw Client ID (without RFC 4361), in hex, or NVT ASCII string format. Thus, “Sun” and 0x53756E are equivalent.

PARAM_REQUEST_LIST
Specifies a list of comma-separated integer values of options for which the client would like values, or symbolic Site or Option option names. Symbolic option names for IPv4 are resolved through /etc/dhcp/ini^{tt}ab. Option names for IPv6 are resolved by means of /etc/dhcp/ini^{tt}ab6.

PARAM_IGNORE_LIST
Specifies a list of options (constructed in the same manner as PARAM_REQUEST_LIST) that the DHCP client will ignore. Ignored options are treated as though the server did not return the options specified. Ignored options are not visible using [dhcpginfo\(1\)](#) or acted on by the client. This parameter can be used, for example, to disable an unwanted client name or default router.

/etc/dhcp/eventhook
Location of a DHCP event program.

属性
See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见
[dhcpginfo\(1\)](#), [ifconfig\(1M\)](#), [init\(1M\)](#), [in.mpathd\(1M\)](#), [in.ndpd\(1M\)](#), [syslog\(3C\)](#), [attributes\(5\)](#), [dhcp\(5\)](#)

《Oracle Solaris 管理：IP 服务》

Croft, B. and Gilmore, J., Bootstrap Protocol (BOOTP) RFC 951, Network Working Group, September 1985.

Droms, R., Dynamic Host Configuration Protocol, RFC 2131, Network Working Group, March 1997.

Lemon, T. and B. Sommerfeld. RFC 4361, Node-specific Client Identifiers for Dynamic Host Configuration Protocol Version Four (DHCPv4). Nominum and Sun Microsystems. February 2006.

Droms, R. RFC 3315, Dynamic Host Configuration Protocol for IPv6 (DHCPv6). Cisco Systems. July 2003.

附注

The `dhcpgent` daemon can be used on IPv4 logical interfaces, just as with physical interfaces. When used on a logical interface, the daemon automatically constructs a Client ID value based on the DUID and IAID values, according to RFC 4361. The `/etc/default/dhcpgent CLIENT_ID` value, if any, overrides this automatic identifier.

Unlike physical IPv4 interfaces, `dhcpgent` does not add or remove default routes associated with logical interfaces.

DHCP can be performed on IPMP IP interfaces to acquire and maintain IPMP data addresses. Because an IPMP IP interface has no hardware address, the daemon automatically constructs a Client ID using the same approach described above for IPv4 logical interfaces. In addition, the lack of a hardware address means the daemon must set the “broadcast” flag in all DISCOVER and REQUEST messages on IPMP IP interfaces. Some DHCP servers may refuse such requests.

DHCP can be performed on IP interfaces that are part of an IPMP group (to acquire and maintain test addresses). The daemon will automatically set the NOFAILOVER and DEPRECATED flags on each test address. Additionally, the daemon will not add or remove default routes in this case. Note that the actual DHCP packet exchange may be performed over any active IP interface in the IPMP group. It is strongly recommended that test addresses have infinite leases. Otherwise, an extended network outage detectable only by probes may cause test address leases to expire, causing `in.mpathd(1M)` to revert to link-based failure detection and trigger an erroneous repair.

引用名

dhcpcfg – DHCP service configuration utility

用法概要

```
dhcpcfg -D -r resource -p path [-u uninterpreted]
    [-l lease_length] [-n ] [-d DNS_domain]
    [-a DNS_server_addresses] [-h hosts_resource]
    [-y hosts_domain]

dhcpcfg -R server_addresses

dhcpcfg -U [-f] [-x] [-h]

dhcpcfg -N network_address [-m subnet_mask] [-b ]
    [-t router_addresses] [-y NIS-domain]
    [-a NIS_server_addresses] [-g]

dhcpcfg -C -r resource -p path [-f] [-k]
    [-u uninterpreted]

dhcpcfg -X filename [-m macro_list] [-o option_list]
    [-a network_addresses] [-f] [-x] [-g]

dhcpcfg -I filename [-f] [-g]

dhcpcfg -P [parameter[=value],]...

dhcpcfg -S [-f] [-e | -d | -r | -q]
```

描述

The `dhcpcfg` command is used to configure and manage the Dynamic Host Configuration Protocol (DHCP) service or BOOTP relay services. It is intended for use by experienced Solaris system administrators and is designed for ease of use in scripts. The `dhcpgmr` utility is recommended for less experienced administrators or those preferring a graphical utility to configure and manage the DHCP service or BOOTP relay service.

The `dhcpcfg` command can be run by root, or by other users assigned to the DHCP Management profile. See [rbac\(5\)](#) and [user_attr\(4\)](#).

`dhcpcfg` requires one of the following function flags: `-D`, `-R`, `-U`, `-N`, `-C`, `-X`, `-I`, `-P` or `-S`.

The `dhcpcfg` menu driven mode is supported in Solaris 8 and previous versions of Solaris.

Where `dhcpcfg` Obtains Configuration Information

`dhcpcfg` scans various configuration files on your Solaris machine for information it can use to assign values to options contained in macros it adds to the `dhcptab` configuration table. The following table lists information `dhcpcfg` needs, the source used, and how the information is used:

Information	Source	Where Used
Timezone	System date, timezone settings	Locale macro
DNS parameters	nsswitch.conf, /etc/resolv.conf	Server macro

NIS parameters	System domainname, nsswitch.conf, NIS	Network macros
Subnetmask	Network interface, netmasks table in nameservice	Network macros

If you have not set these parameters on your server machine, you should do so before configuring the DHCP server with dhcpconfig. Note that if you specify options with the dhcpconfig -D command line, the values you supply override the values obtained from the system files.

The dhcpconfig utility is obsolete and is subject to removal in a future release of Oracle Solaris.

选项

The following options are supported:

- C
Convert to using a new data store, recreating the DHCP data tables in a format appropriate to the new data store, and setting up the DHCP server to use the new data store.

The following sub-options are required:

- p *path_to_data*
The paths for SUNWfiles and SUNWbinfiles must be absolute UNIX pathnames. See [dhcp_modules\(5\)](#).
- r *data_resource*
New data store resource. One of the following must be specified: SUNWfiles or SUNWbinfiles. See [dhcp_modules\(5\)](#).

The following sub-options are optional:

- f
Do not prompt for confirmation. If -f is not used, a warning and confirmation prompt are issued before the conversion starts.
- k
Keep the old DHCP data tables after successful conversion. If any problem occurs during conversion, tables are not deleted even if -k sub-option is not specified.
- u *uninterpreted*
Data which is ignored by dhcpconfig, but passed on to the datastore for interpretation. The private layer provides for module-specific configuration information through the use of the RESOURCE_CONFIG keyword. Uninterpreted data is stored within RESOURCE_CONFIG keyword of [dhcpsvc.conf\(4\)](#). The -u sub-option is not used with the SUNWfiles and SUNWbinfiles data stores. See [dhcp_modules\(5\)](#).
- D
Configure the DHCP service.

The following sub-options are required:

-r *data_resource*

One of the following must be specified: `SUNWfiles` or `SUNWbinfiles`. Other data stores may be available. See [dhcp_modules\(5\)](#).

-p *path*

The paths for `SUNWfiles` and `SUNWbinfiles` must be absolute UNIX pathnames. See [dhcp_modules\(5\)](#).

The following sub-options are optional:

-a *DNS_servers*

IP addresses of DNS servers, separated with commas.

-d *DNS_domain*

DNS domain name.

-h *hosts_resource*

Resource in which to place hosts data. Usually, the name service in use on the server. Valid values are `files` or `dns`.

-l *seconds*

Lease length used for addresses not having a specified lease length, in seconds.

-n

Non-negotiable leases

-y *hosts_domain*

DNS domain name to be used for hosts data. Valid only if `dns` is specified for `-h` sub-option.

-u *uninterpreted*

Data which is ignored by `dhcpconfig`, but passed on to the datastore for interpretation. The private layer provides for module-specific configuration information through the use of the `RESOURCE_CONFIG` keyword. Uninterpreted data is stored within `RESOURCE_CONFIG` keyword of [dhcpsvc.conf\(4\)](#). The `-u` sub-option is not used with the `SUNWfiles` and `SUNWbinfiles` data stores. See [dhcp_modules\(5\)](#).

-I *filename*

Import data from *filename*, containing data previously exported from a Solaris DHCP server. Note that after importing, you may have to edit macros to specify the correct domain names, and edit network tables to change the owning server of addresses in imported networks. Use `dhtadm` and `pntadm` to do this.

The following sub-options are supported:

-f

Replace any conflicting data with the data being imported.

- g
Signal the daemon to reload the `dhcptab` once the import has been completed.
- N *net_address*
Configure an additional network for DHCP service.

The following sub-options are supported:
 - a *NIS_server_addresses*
List of IP addresses of NIS servers.
 - b
Network is a point-to-point (PPP) network, therefore no broadcast address should be configured. If -b is not used, the network is assumed to be a LAN, and the broadcast address is determined using the network address and subnet mask.
 - g
Signal the daemon to reload the `dhcptab`.
 - m *xxx.xxx.xxx.xxx*
Subnet mask for the network; if -m is not used, subnet mask is obtained from `netmasks`.
 - t *router_addresses*
List of router IP addresses; if not specified, router discovery flag is set.
 - y *NIS_domain_name*
If NIS is used on this network, specify the NIS domain name.
- P
Configure the DHCP service parameters. Each parameter and value are specified by the following pattern:
parameter[=*value*] , . . .

Where *parameter* and *value* are:

parameter
One of the DHCP service parameters listed in [dhcpsvc.conf\(4\)](#). If the corresponding *value* is not specified, the current parameter value is displayed. If *parameter* is not specified, all parameters and current values are displayed.

value
Optional string to set the servers parameter to if the value is acceptable. If the value is missing or is empty (""), the parameter and its current value are deleted.

After a parameter has changed the DHCP server requires re-starting before you can use new parameter values.
- R *server_addresses*
Configure the BOOTP relay service. BOOTP or DHCP requests are forwarded to the list of servers specified.

server_addresses is a comma separated list of hostnames and/or IP addresses.

-S

Control the DHCP service.

The following sub-options are supported:

-d

Disable and stop the DHCP service.

-e

Enable and start the DHCP service.

-q

Display the state of the DHCP service. The state is encoded into the exit status.

0	DHCP service disabled and stopped
1	DHCP service enabled and stopped
2	DHCP service disabled and running
3	DHCP service enabled and running

-r

Enable and restart the DHCP service.

-U

Unconfigure the DHCP service or BOOTP relay service.

The following sub-options are supported:

-f

Do not prompt for confirmation. If -f is not used, a warning and confirmation prompt is issued.

-h

Delete hosts entries from name service.

-x

Delete the dhcptab and network tables.

-X*filename*

Export data from the DHCP data tables, saving to *filename*, to move the data to another Solaris DHCP server.

The following sub-options are optional:

-a *networks_to_export*

List of networks whose addresses should be exported, or the keyword ALL to specify all networks. If -a is not specified, no networks are exported.

-g

Signal the daemon to reload the dhcptab after the export has been completed.

`-m macros_to_export`

List of macros to export, or the keyword ALL to specify all macros. If `-m` is not specified, no macros are exported.

`-o options_to_export`

List of options to export, or the keyword ALL to specify all options. If `-o` is not specified, no options are exported.

`-x`

Delete the data from this server after it is exported. If `-x` is not specified you are in effect copying the data.

示例

示例 1 Configuring DHCP Service with Binary Files Data Store

The following command configures DHCP service, using the binary files data store, in the DNS domain `acme.eng`, with a lease time of 28800 seconds (8 hours),

```
example# dhcpconfig -D -r SUNWbinfiles -p /var/dhcp -l 28800\
-d acme.eng -a 120.30.33.4 -h dns -y acme.eng
```

示例 2 Configuring BOOTP Relay Agent

The following command configures the DHCP daemon as a BOOTP relay agent, which forwards BOOTP and DHCP requests to the servers having the IP addresses 120.30.33.7 and 120.30.42.132:

```
example# dhcpconfig -R 120.30.33.7,120.30.42.132
```

示例 3 Unconfiguring DHCP Service

The following command unconfigures the DHCP service, with confirmation, and deletes the DHCP data tables and host table entries:

```
example# dhcpconfig -U -x -h
```

示例 4 Configuring a Network for DHCP Service

The following command configures an additional LAN network for DHCP service, specifying that clients should use router discovery and providing the NIS domain name and NIS server address:

```
example# dhcpconfig -N 120.30.171.0 -y east.acme.eng.com\
-a 120.30.33.4
```

示例 5 Exporting a Network, Macros, and Options from a DHCP Server

The following command exports one network (120.30.171.0) and its addresses, the macro 120.30.171.0, and the options motd and PSptr from a DHCP server, saves the exported data in file `/export/var/120301710_data`, and deletes the exported data from the server.

```
example# dhcpconfig -X /var/dhcp/120301710_export
-a 120.30.171.0 -m 120.30.171.0 -o motd,PSptr
```

示例 6 Importing Data on a DHCP Server

The following command imports DHCP data from a file, /net/golduck/export/var/120301710_data, containing data previously exported from a Solaris DHCP server, overwrites any conflicting data on the importing server, and signals the daemon to reload the dhcptab once the import is completed:

```
example# dhcpcfg -I /net/golduck/export/var/120301710_data -f -g
```

示例 7 Setting DHCP Server Parameters

The following command sets the number of minutes that the DHCP server waits before timing out when updating DNS information on DHCP clients to five minutes.

```
example# example# dhcpcfg -P UPDATE_TIMEOUT=5
```

示例 8 Re-starting the DHCP server

The following command stops and re-starts the DHCP server.

```
example# example# dhcpcfg -S -r
DHCP server stopped
DHCP server started
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTE VALUE
Availability	service/network/dhcp
Interface Stability	Obsolete

另请参见

[dhcpcmgr\(1M\)](#), [dhtadm\(1M\)](#), [in.dhcpd\(1M\)](#), [pntadm\(1M\)](#), [dhcp_network\(4\)](#), [dhcptab\(4\)](#), [dhcpsvc.conf\(4\)](#), [nsswitch.conf\(4\)](#), [resolv.conf\(4\)](#), [user_attr\(4\)](#), [attributes\(5\)](#), [dhcp\(5\)](#), [dhcp_modules\(5\)](#), [rbac\(5\)](#)

《Oracle Solaris 管理：IP 服务》

引用名

dhcpcmgr – graphical interface for managing DHCP service

用法概要

/usr/sadm/admin/bin/dhcpcmgr

描述

dhcpcmgr is a graphical user interface which enables you to manage the Dynamic Host Configuration Protocol (DHCP) service on the local system. It performs the functions of the dhcpcfig, dhtadm, and pntadm command line utilities. You must be root to use dhcpcmgr. The dhcpcmgr Help, available from the Help menu, contains detailed information about using the tool.

The dhcpcmgr utility is obsolete and is subject to removal in a future release of Oracle Solaris.

用法

You can perform the following tasks using dhcpcmgr:

Configure DHCP service	Use dhcpcmgr to configure the DHCP daemon as a DHCP server, and select the data store to use for storing network configuration tables..
Configure BOOTP relay service	Use dhcpcmgr to configure the DHCP daemon as a BOOTP relay.
Manage DHCP or BOOTP relay service	Use dhcpcmgr to start, stop, enable, disable or unconfigure the DHCP service or BOOTP relay service, or change DHCP server parameters.
Manage DHCP addresses	Use dhcpcmgr to add, modify, or delete IP addresses leased by the DHCP service.
Manage DHCP macros	Use dhcpcmgr to add, modify or delete macros used to supply configuration parameters to DHCP clients.
Manage DHCP options	Use dhcpcmgr to add, modify or delete options used to define parameters deliverable through DHCP.
Convert to a new DHCP data store	Use dhcpcmgr to configure the DHCP server to use a different data store, and convert the DHCP data to the format used by the new data store.
Move DHCP data to another server	Use dhcpcmgr to export data from one Solaris DHCP server and import data onto another Solaris DHCP server.

退出状态

The following exit values are returned:

0	Successful completion.
non-zero	An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	network/dhcp/dhcpgmgr
Interface Stability	Obsolete

另请参见

[dhcpconfig\(1M\)](#), [dhtadm\(1M\)](#), [pntadm\(1M\)](#), [in.dhcpd\(1M\)](#), [dhcpsvc.conf\(4\)](#),
[dhcp_network\(4\)](#), [dhcptab\(4\)](#), [attributes\(5\)](#), [dhcp\(5\)](#), [dhcp_modules\(5\)](#)

《Oracle Solaris DHCP Service Developer’s Guide》

《Oracle Solaris 管理：IP 服务》

引用名

dhtadm – DHCP configuration table management utility

用法概要

```
dhtadm -C [-r resource] [-p path] [-u uninterpreted] [-g]
dhtadm -A -s symbol_name -d definition [-r resource]
    [-p path] [-u uninterpreted] [-g]
dhtadm -A -m macro_name -d definition [-r resource]
    [-p path] [-u uninterpreted] [-g]
dhtadm -M -s symbol_name -d definition [-r resource]
    [-p path] [-u uninterpreted] [-g]
dhtadm -M -s symbol_name -n new_name [-r resource]
    [-p path] [-u uninterpreted] [-g]
dhtadm -M -m macro_name -n new_name [-r resource] [-p path]
    [-u uninterpreted] [-g]
dhtadm -M -m macro_name -d definition [-r resource]
    [-p path] [-u uninterpreted] [-g]
dhtadm -M -m macro_name -e symbol=value [-r resource]
    [-p path] [-u uninterpreted] [-g]
dhtadm -D -s symbol_name [-r resource] [-p path]
    [-u uninterpreted] [-g]
dhtadm -D -m macro_name [-r resource] [-p path]
    [-u uninterpreted] [-g]
dhtadm -P [-r resource] [-p path] [-u uninterpreted] [-g]
dhtadm -R [-r resource] [-p path] [-u uninterpreted] [-g]
dhtadm -B [-v] [batchfile] [-g]
```

描述

dhtadm manages the Dynamic Host Configuration Protocol (DHCP) service configuration table, `dhcptab`. You can use it to add, delete, or modify DHCP configuration macros or options or view the table. For a description of the table format, see [dhcptab\(4\)](#).

The dhtadm command can be run by root, or by other users assigned to the DHCP Management profile. See [rbac\(5\)](#) and [user_attr\(4\)](#).

After you make changes with dhtadm, you should issue a SIGHUP to the DHCP server, causing it to read the `dhcptab` and pick up the changes. Do this using the `-g` option.

The dhtadm utility is obsolete and is subject to removal in a future release of Oracle Solaris.

选项

One of the following function flags must be specified with the dhtadm command: `-A`, `-B`, `-C`, `-D`, `-M`, `-P` or `-R`.

The following options are supported:

-A

Add a symbol or macro definition to the `dhcptab` table.

The following sub-options are required:

-d *definition*

Specify a macro or symbol definition.

definition must be enclosed in single quotation marks. For macros, use the form `-d ' :symbol=value:symbol=value: '`. Enclose a *value* that contains colons in double quotation marks. For symbols, the definition is a series of fields that define a symbol's characteristics. The fields are separated by commas. Use the form `-d 'context,code,type,granularity,maximum'`. See [dhcptab\(4\)](#) for information about these fields.

-m *macro_name*

Specify the name of the macro to be added.

The `-d` option must be used with the `-m` option. The `-s` option cannot be used with the `-m` option.

-s *symbol_name*

Specify the name of the symbol to be added.

The `-d` option must be used with the `-s` option. The `-m` option cannot be used with the `-s` option.

-B

Batch process `dhtadm` commands. `dhtadm` reads from the specified file or from standard input a series of `dhtadm` commands and execute them within the same process. Processing many `dhtadm` commands using this method is much faster than running an executable batchfile itself. Batch mode is recommended for using `dhtadm` in scripts.

The following sub-option is optional:

-v

Display commands to standard output as they are processed.

-C

Create the DHCP service configuration table, `dhcptab`.

-D

Delete a symbol or macro definition.

The following sub-options are required:

-m *macro_name*

Delete the specified macro.

-s *symbol_name*

Delete the specified symbol.

-g
Signal the DHCP daemon to reload the `dhcptab` after successful completion of the operation.

-M
Modify an existing symbol or macro definition.

The following sub-options are required:

-d *definition*
Specify a macro or symbol definition to modify.

The definition must be enclosed in single quotation marks. For macros, use the form -d '*:symbol=value:symbol=value:*'. Enclose a *value* that contains colons in double quotation marks. For symbols, the definition is a series of fields that define a symbol's characteristics. The fields are separated by commas. Use the form -d '*context,code,type,granularity,maximum*'. See [dhcptab\(4\)](#) for information about these fields.

-e
This sub-option uses the *symbol=value* argument. Use it to edit a *symbol/value* pair within a macro. To add a symbol which does not have an associate value, enter:

```
symbol=_NULL_VALUE_
```

To delete a symbol definition from a macro, enter:

```
symbol=
```

-m
This sub-option uses the *macro_name* argument. The -n, -d, or -e sub-options are legal companions for this sub-option..

-n
This sub-option uses the *new_name* argument and modifies the name of the object specified by the -m or -s sub-option. It is not limited to macros. . Use it to specify a new macro name or symbol name.

-s
This sub-option uses the *symbol_name* argument. Use it to specify a symbol. The -d sub-option is a legal companion.

-p *path*
Override the [dhcpsvc.conf\(4\)](#) configuration value for `PATH=` with *path*. See [dhcpsvc.conf\(4\)](#) for more details regarding *path*. See [dhcp_modules\(5\)](#) for information regarding data storage modules for the DHCP service.

-P
Print (display) the `dhcptab` table.

-r *data_store_resource*

Override the `dhcpsvc.conf(4)` configuration value for `RESOURCE=` with the *data_store_resource* specified. See [dhcpsvc.conf\(4\)](#) for more details on resource type. See [《Oracle Solaris DHCP Service Developer's Guide》](#) for more information about adding support for other data stores. See [dhcp_modules\(5\)](#) for information regarding data storage modules for the DHCP service.

-R

Remove the `dhcptab` table.

-u *uninterpreted*

Data which is ignored by `dhtadm`, but passed to currently configured public module, to be interpreted by the data store. The private layer provides for module-specific configuration information through the use of the `RESOURCE_CONFIG` keyword. Uninterpreted data is stored within `RESOURCE_CONFIG` keyword of [dhcpsvc.conf\(4\)](#). See [dhcp_modules\(5\)](#) for information regarding data storage modules for the DHCP service.

示例**示例 1** Creating the DHCP Service Configuration Table

The following command creates the DHCP service configuration table, `dhcptab`:

```
# dhtadm -C
```

示例 2 Adding a Symbol Definition

The following command adds a Vendor option symbol definition for a new symbol called `MySym` to the `dhcptab` table in the `SUNWfiles` resource in the `/var/mydhcp` directory:

```
# dhtadm -A -s MySym
  -d 'Vendor=SUNW.PCW.LAN,20,IP,1,0'
  -r SUNWfiles -p /var/mydhcp
```

示例 3 Adding a Macro Definition

The following command adds the `aruba` macro definition to the `dhcptab` table. Note that symbol/value pairs are bracketed with colons (:).

```
# dhtadm -A -m aruba \
  -d ':Timeserv=10.0.0.10 10.0.0.11:DNSServ=10.0.0.1:'
```

示例 4 Modifying a Macro Definition

The following command modifies the `Locale` macro definition, setting the value of the `UTCOffset` symbol to 18000 seconds. Note that any macro definition which includes the definition of the `Locale` macro inherits this change.

```
# dhtadm -M -m Locale -e 'UTCOffset=18000'
```

示例 5 Deleting a Symbol

The following command deletes the Timeserv symbol from the aruba macro. Any macro definition which includes the definition of the aruba macro inherits this change.

```
# dhtadm -M -m aruba -e 'Timeserv='
```

示例 6 Adding a Symbol to a Macro

The following command adds the Hostname symbol to the aruba macro. Note that the Hostname symbol takes no value, and thus requires the special value `_NULL_VALUE_`. Note also that any macro definition which includes the definition of the aruba macro inherits this change.

```
# dhtadm -M -m aruba -e 'Hostname=_NULL_VALUE_'
```

示例 7 Renaming a Macro

The following command renames the Locale macro to MyLocale. Note that any Include statements in macro definitions which include the Locale macro also need to be changed.

```
# dhtadm -M -m Locale -n MyLocale
```

示例 8 Deleting a Symbol Definition

The following command deletes the MySym symbol definition. Note that any macro definitions which use MySym needs to be modified.

```
# dhtadm -D -s MySym
```

示例 9 Printing a dhcptab

The following command prints to standard output the contents of the dhcptab that is located in the data store and path indicated in the dhcpsvc.conf file:.

```
# dhtadm -P
```

示例 10 Executing dhtadm in Batch Mode

The following command runs a series of dhtadm commands contained in a batch file and signals the daemon to reload the dhcptab once the commands have been executed: :

```
# dhtadm -B addmacros -g
```

退出状态

- 0 Successful completion.
- 1 Object already exists.
- 2 Object does not exist.

3
Non-critical error.

4
Critical error.

文件
/etc/inet/dhcpsvc.conf
contains service configuration parameters for the DHCP service

属性
See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/dhcp
Interface Stability	Obsolete

另请参见
[dhcpconfig\(1M\)](#), [dhcpgmgr\(1M\)](#), [in.dhcpd\(1M\)](#), [dhcpsvc.conf\(4\)](#), [dhcp_network\(4\)](#), [dhcptab\(4\)](#), [hosts\(4\)](#), [user_attr\(4\)](#), [attributes\(5\)](#), [dhcp\(5\)](#), [dhcp_modules\(5\)](#)[rbac\(5\)](#)

《Oracle Solaris DHCP Service Developer’s Guide》

《Oracle Solaris 管理：IP 服务》

Alexander, S., and R. Droms, DHCP Options and BOOTP Vendor Extensions, RFC 1533, Lachman Technology, Inc., Bucknell University, October 1993.

Droms, R., Interoperation Between DHCP and BOOTP, RFC 1534, Bucknell University, October 1993.

Droms, R., Dynamic Host Configuration Protocol, RFC 1541, Bucknell University, October 1993.

Wimer, W., Clarifications and Extensions for the Bootstrap Protocol, RFC 1542, Carnegie Mellon University, October 1993.

引用名 dig – DNS 查找实用程序

用法概要

```
dig [@server] [-b address] [-c class] [-f filename]
    [-k filename] [-m] [-p port#] [-q name] [-t type] [-x addr]
    [-y [hmac:]name:key] [-4] [-6] [name] [type] [class] [queryopt]...

dig [-h]

dig [global-queryopt...] [query...]
```

描述

dig 实用程序（域信息探测器）是一种灵活的工具，用于查询 DNS 名称服务器。它可以执行 DNS 查找，并显示从查询的名称服务器返回的答案。由于该实用程序灵活、易用并且输出明确，因此大多数 DNS 管理员都使用 **dig** 来解决 DNS 问题。其他查找工具的功能通常比 **dig** 少。

虽然 **dig** 通常与命令行参数一起使用，但它还具有从文件中读取查找请求的批处理操作模式。如果指定了 **-h** 选项，将输出该实用程序的命令行参数以及选项的概述。与早期版本不同，**dig** 的 BIND 9 实现允许从命令行发出多个查找。

除非收到查询特定名称服务器的指令，否则 **dig** 会尝试查询 `/etc/resolv.conf` 中列出的每个服务器。

如果未指定任何命令行参数或选项，**dig** 会执行“.”（根）的 NS 查询。

可以使用 `${HOME}/.digrc` 为 **dig** 设置每用户缺省值。系统会读取此文件，且在命令行参数之前应用其中的任何选项。

IN 和 **CH** 类名与 **IN** 和 **CH** 顶级域名重叠。因此，可 **-t** 和 **-c** 选项来指定类型和类，或使用 **"IN."** 和 **"CH."** 查找这些顶级域。

简单用法

下面是 **dig** 的典型调用：

```
dig @server name type
```

其中：

server

要查询的名称服务器的名称或 IP 地址。可以是采用点分十进制记法的 IPv4 地址或采用冒号分隔记法的 IPv6 地址。如果提供的 *server* 参数为主机名，**dig** 会在查询该名称服务器之前解析该名称。如果未提供 *server* 参数，**dig** 会查找 `/etc/resolv.conf` 并查询其中列出的名称服务器。将会显示来自做出响应的名称服务器的回复。

name

要查找的资源记录的名称。

type

指示所需的查询类型（**ANY**、**A**、**MX**、**SIG** 等）。*type* 可以是任何有效的查询类型。如果未提供 *type* 参数，**dig** 会执行 **A** 记录查找。

选项

所支持的选项如下：

-4

仅使用 IPv4 传输。缺省情况下，IPv4 和 IPv6 传输均可使用。选项 **-4** 和 **-6** 相互排斥。

-6

仅使用 IPv6 传输。缺省情况下，IPv4 和 IPv6 传输均可使用。选项 **-4** 和 **-6** 相互排斥。

-b address

将查询的源 IP 地址设置为 *address*。该地址必须是主机的网络接口之一上的有效地址、**0.0.0.0** 或 **::**。通过附加 **#<port>** 可以指定可选端口。

-c class

覆盖缺省的查询类（IN 表示 Internet）。*class* 参数可以是任何有效类，例如，HS 表示 Hesiod 记录，CH 表示 CHAOSNET 记录。

-f filename

通过从文件 *filename* 中读取要处理的查找请求列表，以批处理模式运行。此文件包含多个查询（每行一个）。文件中每个项的组织方式应与使用命令行界面查询 **dig** 时显示的方式相同。

-h

输出命令行参数与选项的概述。

-k filename

指定事务签名 (TSIG) 密钥文件以使用 TSIG 签署由 **dig** 发送的 DNS 查询及其响应。

-m

启用内存使用情况调试。

-p port#

查询非标准端口号。*port#* 参数是 **dig** 发送其查询所用的端口号，而不是标准 DNS 端口号 53。此选项会测试配置为侦听非标准端口号上的查询的名称服务器。

-q name

将查询名称设置为 *name*。这样可以轻松将查询名称与其他参数区分开来。

-T type

将查询类型设置为 *type*，可以是 BIND9 中支持的任意有效查询类型。缺省的查询类型为“A”，除非提供了 **-x** 选项以指示反向查找。通过指定 AXFR 类型可请求区域传输。需要进行增量区域传输 (IXFR) 时，将 *type* 设置为 **ixfr=N**。由于区域的 SOA 记录中的序列号为 *N*，增量区域传输将包含对该区域所做的更改。

-x addr

简化反向查找（将地址映射至名称）。*addr* 参数是采用点分十进制记法的 IPv4 地址或采用冒号分隔记法的 IPv6 地址。如果使用此选项，则无需提供 *name*、*class* 和 *type* 参数。**dig** 实用程序会自动执行名称查找，例如

11.12.13.10.in-addr.arpa，并将查询类型和类分别设置为 PTR 和 IN。缺省情况

下，在 IP6.ARPA 域下使用半字节格式查找 IPv6 地址。要使用采用 IP6.INT 域的早期 RFC1886 方法，请指定 `-i` 选项。位字符串标签 (RFC 2874) 目前处于试验阶段，不会尝试使用。

`-y [hmac:]name:key`

在命令行上指定一个事务签名 (TSIG) 密钥。这样可以签署 `dig` 发送的 DNS 查询及其响应。您还可以在命令行上使用 `-y` 选项指定 TSIG 密钥本身。可选的 `hmac` 是 TSIG 的类型；缺省值为 `HMAC-MD5`。`name` 参数是 TSIG 密钥的名称，`key` 参数是实际密钥。此密钥为 base-64 编码字符串，通常由 `dnssec-keygen(1M)` 生成。

在多用户系统上使用 `-y` 选项时应谨慎，因为该密钥会显示在来自 `ps(1)` 的输出中或 `shell` 的历史文件中。将 TSIG 验证与 `dig` 结合使用时，所查询的名称服务器需要知道正在使用的密钥和算法。在 BIND 中，这一点可通过在 `named.conf` 中提供适当的 `key` 和 `server` 语句实现。

查询选项

`dig` 实用程序提供了一些会影响查找实施以及显示结果的方式的查询选项。其中的一些选项用于设置或复位查询标头中的标志位，一些选项用于确定输出答案的哪些部分，其他选项用于确定超时和重试策略。

每个查询选项都通过一个带加号 (+) 前缀的关键字标识。一些关键字可设置或复位选项。这些关键字可能会具有 `no` 字符串前缀，以否定该关键字的含义。其他关键字用于分配选项的值，例如超时时间间隔。这些关键字的格式为 `+keyword=value`。查询选项包括：

`+[no]tcp`

查询名称服务器时使用 [不使用] TCP。缺省的行为是使用 UDP，除非请求了 AXFR 或 IXFR 查询，在这种情况下将使用 TCP 连接。

`+[no]vc`

查询名称服务器时使用 [不使用] TCP。这是 `+[no]tcp` 的另一种等效语法，之所以提供此语法是为了实现向后兼容性。“vc”表示 virtual circuit（虚拟线路）。

`+[no]ignore`

忽略 UDP 响应中的截断，而不是使用 TCP 重试。缺省情况下，执行 TCP 重试。

`+domain=somename`

设置搜索列表，使其包含单个域 `somename`，如同在 `/etc/resolv.conf` 的 `domain` 指令中指定一样，并且启用搜索列表处理，如同指定了 `+search` 选项一样。

`+[no]search`

使用 [不使用] `resolv.conf`（如果有）中的 `searchlist` 或 `domain` 指令定义的搜索列表。缺省情况下，不使用此搜索列表。

`+[no]showsearch`

执行 [不执行] 显示中间结果的搜索。

`+[no]defname`

已过时，视为 `+[no]search` 的同义语法。

- +`[no]aaonly`**
设置查询中的 `aa` 标志。
- +`[no]aaflag`**
+`[no]aaonly` 的同义语法。
- +`[no]adflag`**
设置 [不设置] 查询中的 AD (authentic data, 可信数据) 位。这要求服务器返回值, 无论是否根据服务器的安全策略将所有答案和授权部分均验证为安全。设置为 `AD=1` 指示所有记录都已验证为安全, 而且答案不是来自 `OPT-OUT` 范围。`AD=0` 指示答案的某些部分不安全或未经验证。
- +`[no]cdflag`**
设置 [不设置] 查询中的 CD (checking disabled, 禁用检查) 位。这要求服务器不对响应执行 DNSSEC 验证。
- +`[no]cl`**
输出记录时显示 [不显示] 类。
- +`[no]ttlid`**
输出记录时显示 [不显示] TTL。
- +`[no]recurse`**
切换查询中 RD (recursion desired, 需要递归) 位的设置。缺省情况下会设置此位, 这意味着 `dig` 通常会发送递归查询。使用 `+nssearch` 或 `+trace` 查询选项时, 会自动禁用递归。
- +`[no]nssearch`**
设置此选项时, `dig` 会尝试查找区域的包含待查找名称的权威名称服务器, 并显示每台名称服务器中对应于该区域的 SOA 记录。
- +`[no]trace`**
对要查找的名称, 从根名称服务器切换委托路径跟踪。缺省情况下, 禁用跟踪。启用跟踪时, `dig` 会执行迭代查询以解析要查找的名称。将按照来自根服务器的引用, 显示来自每个服务器的用于解析查找的答案。
- +`[no]cmd`**
切换是否在输出中显示标识 `dig` 的版本以及所应用的查询选项的初始注释。缺省情况下, 输出此注释。
- +`[no]short`**
提供扼要答案。缺省情况下, 以详细模式输出答案。
- +`[no]identify`**
启用 `+short` 选项时, 显示 [或不显示] 提供答案的 IP 地址和端口号。如果要求提供简洁格式的答案, 缺省情况下, 不显示提供答案的服务器的源地址和端口号。
- +`[no]comments`**
切换输出中注释行的显示。缺省情况下显示注释。

+`[no]stats`

切换统计信息的输出：执行查询的时间、回复大小等。缺省行为是输出查询统计信息。

+`[no]qr`

发送时输出 [不输出] 查询。缺省情况下不输出查询。

+`[no]question`

返回答案时输出 [不输出] 查询的问题部分。缺省情况下，以注释形式输出问题部分。

+`[no]answer`

显示 [不显示] 回复的答案部分。缺省情况下显示该部分。

+`[no]authority`

显示 [不显示] 回复的授权部分。缺省情况下显示该部分。

+`[no]additional`

显示 [不显示] 回复的附加部分。缺省情况下显示该部分。

+`[no]all`

设置或清除所有显示标志。

+`time=T`

将查询的超时时间设置为 T 秒。缺省超时时间为 5 秒。尝试将 T 设置为小于 1 的值将会应用 1 秒的查询超时。

+`tries=T`

将 UDP 尝试的最大次数设置为 T 。缺省值为 3（1 次初始尝试以及随后 2 次重试）。如果 T 小于等于 0，会将重试次数向上舍入为 1 次而不显示相关提示。

+`retry=T`

将 UDP 重试次数设置为 T 。缺省为 2。

+`ndots=D`

将需要显示在 *name* 中的点数设置为 D ，以便将其视为绝对名。缺省值是使用 `/etc/resolv.conf` 中的 `ndots` 语句定义的值，如果 `ndots` 语句不存在，则缺省值为 1。点数较少的名称将解释为相对名称，并在 `/etc/resolv.conf` 中的 `search` 或 `domain` 指令中列出的域中进行搜索。

+`bufsize=B`

将使用 EDNS0 通告的 UDP 消息缓冲区大小设置为 B 字节。此缓冲区的最大和最小大小分别为 65535 和 0 字节。超出此范围的值将相应地进行向上或向下舍入。

+`edns=#`

指定用于查询的 EDNS 版本。有效值为 0 至 255。设置 EDNS 版本会导致发送一条 EDNS 查询。`+noedns` 将清除记住的 EDNS 版本。

+`[no]multiline`

以详细的多行格式输出包含用户可读注释的记录，例如 SOA 记录。缺省情况下，每行输出一条记录，以便于计算机分析 `dig` 输出。

+`[no]fail`

如果收到 `SERVFAIL`，请勿尝试下一个服务器。缺省情况下，不会尝试下一个服务器，这与常规的桩模块解析器行为相反。

+`[no]besteffort`

尝试显示格式错误的消息的内容。缺省情况下，不显示格式错误的答案。

+`[no]dnssec`

通过在查询的附加部分的 OPT 记录中设置 DNSSEC OK 位 (DO)，请求发送 DNSSEC 记录。

+`[no]sigchase`

追踪 DNSSEC 签名链。需要使用 `-DDIG_SIGCHASE` 来编译 `dig`。

+`trusted-key=####`

指定包含与 `+sigchase` 结合使用的可信密钥的文件。每条 `DNSKEY` 记录必须占据一行。

如果未指定，`dig` 会查找 `/etc/trusted-key.key`，然后在当前目录中查找 `trusted-key.key`。

需要使用 `-DDIG_SIGCHASE` 来编译 `dig`。

+`[no]topdown`

追踪 DNSSEC 签名链时，执行自上而下的验证。需要使用 `-DDIG_SIGCHASE` 来编译 `dig`。

+`[no]nsid`

发送查询时，包括 EDNS 名称服务器 ID 请求。

多个查询

除了支持 `-f` 批处理文件选项以外，`dig` 的 BIND 9 实现还支持在命令行上指定多个查询。对于每条查询，可提供其自己的标志、选项以及查询选项集。

这种情况下，在上述命令行语法中，每个 *query* 参数表示一条查询。每条查询由任一标准选项和标志、要查找的名称、可选的查询类型、类以及应当应用于该查询的任意查询选项组成。

还可以提供应当应用于所有查询的全局查询选项集。这些全局查询选项必须居于命令行上提供的第一组名称、类、类型、选项、标志和查询选项之前。任何全局查询选项（`+[no]cmd` 选项除外）均可由查询特定的查询选项集覆盖。例如：

```
dig +qr www.isc.org any -x 127.0.0.1 isc.org ns +noqr
```

以上语法显示如何在命令行中使用 `dig` 来执行三个查找：针对 `www.isc.org` 的 ANY 查询、反向查找 `127.0.0.1` 以及查询 `isc.org` 的 NS 记录。应用全局查询选项 `+qr`，因此 `dig` 可以显示针对每次查找所做的初始查询。最后的查询具有本地查询选项 `+noqr`，表示 `dig` 在查找 `isc.org` 的 NS 记录时不会输出初始查询。

文件 `/etc/resolv.conf`
 解析器配置文件

`${HOME}/.digrc`
 用户定义的配置文件

属性 有关以下属性的说明，请参见 `attributes(5)`：

属性类型	属性值
可用性	network/dns/bind
接口稳定性	Volatile（可变）

另请参见 `dnssec-keygen(1M)`、`host(1M)`、`named(1M)`、`nslookup(1M)`、`attributes(5)`

 RFC1035

 请参见 BIND 9 Administrator's Reference Manual。从本手册页发布之日起，将在 <https://www.isc.org/software/bind/documentation> 上提供该文档。

已知问题 查询选项可能过多。

附注 `nslookup(1M)` 和 `dig` 目前将“Not Implemented”（未实施）报告为 `NOTIMP` 而不是 `NOTIMPL`。这会影响查找 `NOTIMPL` 的脚本。

引用名 directoryserver – front end for the Directory Server (DS)

用法概要

```
/usr/sbin/directoryserver
    { setup [-f configuration_file] | uninstall}

/usr/sbin/directoryserver
    {start-admin | stop-admin | restart-admin | startconsole}

/usr/sbin/directoryserver [{-s | -server} server-instance ]
    {start | stop | restart}

/usr/sbin/directoryserver { -s |-server } server-instance
    { monitor | saveconfig | restoreconfig | db2index-task |
      ldif2db-task | ldif2db | ldif2ldap | vlvindex | db2ldif |
      db2ldif-task | db2bak | db2bak-task | bak2db | bak2db-task |
      suffix2instance | account-status | account-activate |
      account-inactivate }
    {...}

/usr/sbin/directoryserver nativetoascii | admin_ip | ldif |
    pdhash | idsktune | mmlldif | keyupg
    {...}

/usr/sbin/directoryserver { magt | sagt } {...}

/usr/sbin/directoryserver help [subcommand]
```

描述 The directoryserver command is a comprehensive, front end to the utility programs provided by the Solaris Directory Server (DS).

Options for the directoryserver command itself must appear before the subcommand. Arguments for a subcommand must appear after the subcommand. Subcommands have specific arguments. See SUBCOMMANDS.

Subcommands The following subcommands are supported:

account-inactivate <i>args</i>	Inactivates and locks an entry or group of entries.
	The account-inactivate subcommand supports the following arguments:
[-D <i>rootdn</i>]	Directory Server <i>userDN</i> with root permissions, such as Directory Manager.
[-h <i>host</i>]	Host name of Directory Server. The default value is the full hostname of the machine where Directory Server is installed.
-I <i>DN</i>	Entry DN or role DN to activate.

- j *file* Password associated with the user DN. This option allows the password to be stored in clear text in the named file for scripting.

This is considered insecure. Use with extreme caution.
- [-p *port*] Directory Server port. The default value is the LDAP port of Directory Server specified at installation time.
- w *password* Password associated with the user DN. Supplying the password on the command line is visible using the /bin/ps command. This is considered insecure. Use with extreme caution.

The value - can be used in place the password. The program prompts the user for a password to be entered from the terminal.

account-activate *args*

Activates an entry or group of entries.

- The `account-activate` subcommand supports the following arguments
- D *rootdn* Directory Server userDN with root permissions, such as Directory Manager.
 - h *host* Host name of Directory Server. The default value is the full hostname of the machine where Directory Server is installed.
 - I *DN* Entry DN or role DN to activate.
 - j *file* Password associated with the user DN. This option allows the password to be stored in clear text in the named file for scripting.

This is considered insecure. Use with extreme caution.

-p port Directory Server port. The default value is the LDAP port of Directory Server specified at installation time.

-w password Password associated with the user DN. Supplying the password on the command line is visible using the `/bin/ps` command. This is considered insecure. Use with extreme caution.

The value `-` can be used in place the password. The program prompts the user for a password to be entered from the terminal.

account-status *args*

Provides account status information to establish whether an entry or group of entries is inactivated or not.

The `account-status` subcommand supports the following arguments:

-D rootdn

-h host Host name of Directory Server. The default value is the full hostname of the machine where Directory Server is installed.

-I DN Entry DN or role DN whose status is required.

-j file Password associated with the user DN. This option allows the password to be stored in clear text in the named file for scripting.

This is considered insecure. Use with extreme caution.

-p port Directory Server port. The default value is the LDAP port of Directory Server specified at installation time.

-w password Password associated with the rootDN. Supplying the password on the command line is visible using the

/bin/ps command. This is considered insecure. Use with extreme caution.

The value - can be used in place of the password. The program prompts the user for a password to be entered from the terminal.

admin_ip args Change the IP address of the administrative server in the configuration.

The admin_ip subcommand supports the following arguments:

- dir_mgr_DN Directory Manager's DN.
- dir_mgr_password Directory Manager's password.
- old_ip Old IP.
- new_ip New IP.
- port_# Port number.

bak2db backup_directory Restore the database from the most recent archived backup.

Specify backup_directory as the backup directory.

bak2db-task args Restore the data to the database.

The bak2db-task subcommand supports the following arguments:

- a directory] Directory where the backup files are stored. By default it is under /var/ds5/slapd-serverID/bak
- D rootDN User DN with root permissions, such as Directory Manager. The default is the DN of the directory manager which is read from the nsslapd-root attribute under cn=config.

-j *file* Password associated with the user DN. This option allows the password to be stored in clear text in the named file for scripting.

This is considered insecure. Use with extreme caution.

[-t *database_type*] Database type. The only possible database type is `ldbm`.

[-v] Verbose mode.

-w *password* Password associated with the user DN. Supplying the password on the command line is visible using the `/bin/ps` command. This is considered insecure. Use with extreme caution.

The value `-` can be used in place the password. The program prompts the user for a password to be entered from the terminal.

`db2bak-task args` Back up the contents of the database. It creates an entry in the directory that launches this dynamic task. An entry is generated based upon the values provided for each option.

The `db2bak-task` subcommand supports the following arguments:

[-a *directory*] Directory where the backup files are stored. By default it is under `/var/ds5/slapd-serverID/bak`. The backup file is named according to the year-month-day-hour format (`YYYY_MM_DD_hhmmss`).

-D *rootDN* User DN with root permissions, such as Directory

	Manager. The default is the DN of the directory manager which is read from the <code>nsslapd-root</code> attribute under <code>cn=config</code> .
<code>-j file</code>	Password associated with the user DN. This option allows the password to be stored in clear text in the named file for scripting.
	This is considered insecure. Use with extreme caution.
<code>-t database_type</code>	Database type. The only possible database type is <code>ldbm</code> .
<code>[-v]</code>	Verbose mode.
<code>-w password</code>	Password associated with the user DN. Supplying the password on the command line is visible using the <code>/bin/ps</code> command. This is considered insecure. Use with extreme caution.
	The value <code>-</code> can be used in place the password. The program prompts the user for a password to be entered from the terminal.
<code>db2bak [backup_directory]</code>	Create a backup of the current database contents. The server must be stopped to run this subcommand. The default is <code>/var/ds5/slapd-serverID/bak</code> . The backup file is named according to the year-month-day-hour format (<code>YYYY_MM_DD_hhmmss</code>).
<code>db2index-text args</code>	Create and generate the new set of indexes to be maintained following the modification of indexing entries in the <code>cn=config</code> configuration file. The <code>db2index-text</code> subcommand supports the following arguments:

- D *rootdn* User DN with root permissions, such as Directory Manager.
 - j *file* Password associated with the user DN. This option allows the password to be stored in clear text in the named file for scripting. This is considered insecure. Use with extreme caution.
 - n *backend_instance* Instance to be indexed.
 - [-t *attributeName*] Name of the attribute to be indexed. If omitted, all indexes defined for that instance are generated.
 - [-v] Verbose mode.
 - w *password* Password associated with the user DN. Supplying the password on the command line is visible using the /bin/ps command. This is considered insecure. Use with extreme caution.
- The value - can be used in place the password. The program prompts the user for a password to be entered from the terminal.

db2ldif-task *args*

Exports the contents of the database to LDIF. It creates an entry in the directory that launches this dynamic task. The entry is generated based upon the values you provide for each option. To run this subcommand the server must be running and either -n *backend_instance* or -s *include* suffix is required.

The db2ldif-task subcommand supports the following arguments:

- [-a *outputfile*] File name of the output LDIF file.

-C	Only the main db file is used.
-D <i>rootDN</i>	User DN with root permissions, such as Directory Manager.
-j <i>file</i>	Password associated with the user DN. This option allows the password to be stored in clear text in the named file for scripting. This is considered insecure. Use with extreme caution.
[-M]	Output LDIF is stored in multiple files.
[-m]	Minimal base 64 encoding.
{-n <i>backend_instance</i> }*	Instance to be exported.
[-N]	Minimal base 64 encoding.
[-o]	Output LDIF to be stored in one file by default with each instance stored in <i>instance_file</i> name.
[-r]	Export replica.
[-s] <i>includesuffix</i> *	Suffix(es) to be included or to specify the subtrees to be included if -n has been used.
[-u]	Request that the unique ID is not exported.
[-U]	Request that the output LDIF is not folded.
-w <i>password</i>	Password associated with the user DN. Supplying the password on the command line is visible

		using the <code>/bin/ps</code> command. This is considered insecure. Use with extreme caution.
		The value <code>-</code> can be used in place the password. The program prompts the user for a password to be entered from the terminal.
	<code>{ -x <i>excludesuffix</i> }*</code>	Suffixes to be excluded.
	<code>[-1]</code>	Delete, for reasons of backward compatibility the first line of the LDIF file that gives the version of the LDIF standard.
<code>db2ldif</code>	<code>args</code>	Export the contents of the database to LDIF. You must specify either the <code>-n</code> or the <code>-s</code> option or both.
		The <code>db2ldif</code> subcommand supports the following options:
	<code>[-a <i>outputfile</i>]</code>	File name of the output LDIF file.
	<code>[-C]</code>	Only use the main db file.
	<code>[-m]</code>	Minimal base64 encoding.
	<code>[-M]</code>	Use of several files for storing the output LDIF with each instance stored in <i>instance_file name</i> (where file name is the file name specified for <code>-a</code> option).
	<code>{ -n <i>baclemd_instance</i> }*</code>	Instance to be exported.
	<code>[-N]</code>	Specify that the entry IDs are not to be included in the LDIF output. The entry IDs are necessary

	only if the db2ldif output is to be used as input to db2index-text.						
[-r]	Export replica.						
{ -s includesuffix }*	Suffixes to be included or to specify the subtrees to be included if -n has been used.						
[{ -x excludesuffix }]*	Suffixes to be excluded.						
[-u]	Request that the unique id is not exported.						
[-U]	Request that the output LDIF is not folded.						
[-1]	Delete, for reasons of backward compatibility, the first line of the LDIF file which gives the version of the LDIF standard.						
help [subcommand]	Display directoryserver usage message or subcommand specific usage message.						
idsktune args	Provide an easy and reliable way of checking the patch levels and kernel parameter settings for your system. You must install the Directory Server before you can run idsktune. It gathers information about the operating system, kernel, and TCP stack to make tuning recommendations. The idsktune subcommand supports the following arguments: <table><tr><td>[-c]</td><td>Client-specific tuning: the output only includes tuning recommendations for running a directory client application.</td></tr><tr><td>[-D]</td><td>Debug mode: the output includes the commands it runs internally, preceded by DEBUG heading.</td></tr><tr><td>[-i installldir]</td><td>The install directory.</td></tr></table>	[-c]	Client-specific tuning: the output only includes tuning recommendations for running a directory client application.	[-D]	Debug mode: the output includes the commands it runs internally, preceded by DEBUG heading.	[-i installldir]	The install directory.
[-c]	Client-specific tuning: the output only includes tuning recommendations for running a directory client application.						
[-D]	Debug mode: the output includes the commands it runs internally, preceded by DEBUG heading.						
[-i installldir]	The install directory.						

	[-q] Quiet mode. Output only includes tuning recommendations. OS version statements are omitted.
	[-v] Version. Gives the build date identifying the version of the toll.
keyupg args	Upgrade the key from Lite to normal (only one way). The keyupg subcommand supports the following arguments: -kkey The key to be upgraded. -f key_file_path The key file path.
ldif2db-task args	Import data to the directory. It create an entry in the directory that launches this dynamic task. The entry is generated based upon the values you provide for each option. The server must be running when you run this subcommand. The ldif2sb-task subcommand supports the following arguments: [-c] Request that only the core db is created without attribute indexes. -D rootDN User DN with root permissions, such as Directory Manager. [-g string] Generation of a unique ID. Enter none for no unique ID to be generated and deterministic for the generated unique ID to be name-based. Generates a time based unique ID by default. If you use the deterministic generation to have a name-based unique ID, you can also

specify the namespace you want the server to use as follows:

`-g deterministic namespace_id`

where `namespace_id` is a string of characters in the following format

`00-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx`

Use this option if you want to import the same LDIF file into two different directory servers, and you want the contents of both directories to have the same set of unique IDs. If unique IDs already exist in the LDIF file you are importing, then the existing IDs are imported to the server regardless of the options you have specified.

`[ -G namespace_id ]`

Generate a namespace ID as a name-based unique ID. This is the same as specifying `-g deterministic`.

`{ -i filename }*`

File name of the input LDIF files. When you import multiple files, they are imported in the order in which you specify them on the command line.

`-j file`

Password associated with the user DN. This option allows the password to be stored in clear text in the named file for scripting. This is considered insecure. Use with extreme caution.

`-n backend_instance`

Instance to be imported.

	[-O]	Request that only the core db is created without attribute indexes.
	{ -s <i>includesuffix</i> }*	Suffixes to be included. This argument can also be used to specify the subtrees to be included with -n.
	-w <i>password</i>	Password associated with the user DN. Supplying the password on the command line is visible using the /bin/ps command. This is considered insecure. Use with extreme caution.
		The value - can be used in place the password. The program prompts the user for a password to be entered from the terminal.
	[{-x <i>excludesuffix</i> }*]	
	[-v]	Verbose mode.
ldif <i>args</i>		Format LDIF files, and create base 64 encoded attribute values. With Base 64 Encoding you can represent binary data, such as a JPEG image, in LDIF by using base 64 encoding. You identify base 64 encoded data by using the :: symbol. The <code>ldif</code> subcommand takes any input and formats it with the correct line continuation and appropriate attribute information. The subcommand also senses whether the input requires base 64 encoding.
		The <code>ldif</code> subcommand supports the following arguments
	[-b]	Interpret the entire input as a single binary value. If -b is not present, each line is considered to be a separate input value.
	[<i>attrtype</i>]	If -b is specified, the output is <code>attrtype:: <base 64 encoded</code>

	value.
ldif2db args	Import the data to the directory. To run this subcommand the server must be stopped. Note that ldif2db supports LDIF version 1 specifications. You can load an attribute using the URL specifier notation, for example: jpegphoto:file:///tmp/myphoto.jpg [-c] Merge chunk size. [-g <i>string</i>] Generation of a unique ID. Type none for no unique ID to be generated and deterministic for the generated unique ID to be name-based. By default a time based unique ID is generated. If you use the deterministic generation to have a name-based unique ID, you can also specify the namespace you want the server to use as follows: where <i>namespace_id</i> is a string of characters in the following format: 00-xxxxxxxx-xxxxxxxx-xxxxxxxx-xxxxxxxx Use this option if you want to import the same LDIF file into two different directory servers, and you want the contents of both directories to have the same set of unique IDs. If unique IDs already exist in the LDIF file you are importing, then the existing IDs are imported to the server regardless of the options you have specified.
-g deterministic namespace_id	

	<i>[-G naemspace_id]</i>	Generate a namespace ID as a name-based unique ID. This is the same as specifying the -g deterministic option.
	<i>{- filename}*</i>	File name of the input LDIF file(s). When you import multiple files, they are imported in the order in which you specify them on the command line.
	<i>-n backend_instance</i>	Instance to be imported.
	<i>[-O]</i>	Request that only the core db is created without attribute indexes.
	<i>{-s includesuffix}*</i>	Suffixes to be included or to specify the subtrees to be included if -n has been used.
	<i>[{-x excludesuffix}*]</i>	Suffixes to be excluded
<i>ldif2ldap rootDN password filename</i>		Perform an import operation over LDAP to the Directory Server. To run this subcommand the server must be running.
		The <i>ldif2ldap</i> subcommand supports the following arguments:
	<i>rootdn</i>	User DN with root permissions, such as Directory Manager.
	<i>password</i>	Password associated with the user DN.
	<i>filename</i>	File name of the file to be imported. When you import multiple files, they are imported in the order in which you specify them on the command line.
<i>magt CONFIG INIT</i>		Start SNMP master agent. The Config and INIT files are in <i>/usr/iplanet/ds5/plugins/snmp/magt</i> . For more information, see the «iPlanet Directory Server 5.1 Administrator's Guide» .
		The <i>magt</i> subcommand supports the following options:

	CONFIG	The CONFIG file defines the community and the manager that master agent works with. Specify the manager value as a valid system name or an IP address.						
	INIT	The INIT file is a nonvolatile file that contains information from the MIB-II system group, including system location and contact information. If INIT doesn't already exist, starting the master agent for the first time creates it. An invalid manager name in the CONFIG file causes the master agent start-up to fail.						
monitor		Retrieves performance monitoring information using the <code>ldapsearch</code> command-line utility.						
mmldif args		Combine multiple LDIF files into a single authoritative set of entries. Typically each LDIF file is from a master server cooperating in a multi master replication agreement.[e.g. masters that refuse to sync up for whatever reason]. Optionally, it can generate LDIF change files that could be applied to original to bring it up to date with authoritative. At least two input files must be specified. The <code>mmldif</code> subcommand supports the following arguments: <table><tr><td>[-c <i>inputfile</i> ...]</td><td>Write a change file (.delta) for each input file. Specify <i>inputfile</i> as the input LDIF files.</td></tr><tr><td>[-D]</td><td>Print debugging information.</td></tr><tr><td>[-o <i>out.ldif</i>]</td><td>Write authoritative data to this file.</td></tr></table>	[-c <i>inputfile</i> ...]	Write a change file (.delta) for each input file. Specify <i>inputfile</i> as the input LDIF files.	[-D]	Print debugging information.	[-o <i>out.ldif</i>]	Write authoritative data to this file.
[-c <i>inputfile</i> ...]	Write a change file (.delta) for each input file. Specify <i>inputfile</i> as the input LDIF files.							
[-D]	Print debugging information.							
[-o <i>out.ldif</i>]	Write authoritative data to this file.							
nativetoascii args		Convert one language encoding to another. For example, convert a native language to UTF-8 format. The <code>nativetoascii</code> subcommand supports the following options: <table><tr><td>-d <i>Encodings Directory</i></td><td>Path to the</td></tr></table>	-d <i>Encodings Directory</i>	Path to the				
-d <i>Encodings Directory</i>	Path to the							

	directory which contains the conv directory
<code>[-i <i>input_filename</i> -o <i>output_filename</i>]</code>	The input file name and output file name.
<code>-l</code>	List supported encodings
<code>-r</code>	Replace existing files.
<code>-s <i>suffix</i></code>	Suffix to be mapped to the backend.
<code>-s <i>SourceEncoding</i></code>	Source Encoding of input stream.
<code>-t <i>TargetEncoding</i></code>	Target Encoding of output stream.
<code>-v</code>	Verbose output.
<code>pwdhash <i>args</i></code>	Print the encrypted form of a password using one of the server's encryption algorithms. If a user cannot log in, you can use this script to compare the user's password to the password stored in the directory. The <code>pwdhash</code> subcommand supports the following arguments:

	-c <i>comparepwd</i> -s <i>scheme</i>	The available schemes are SSHA, SHA, CRYPT and CLEARE. It generates the encrypted passwords according to scheme's algorithm. The -c specifies the encrypted password to be compared with. The result of comparison is either OK or doesn't match.
	-D <i>instance-dir</i>	The instance directory.
	[-H]	The passwords are hex-encoded.
	<i>password ...</i>	The clear passwords to generate encrypted form from or to be compared with.
restart	Restarts the directory server. When the -s option is not specified, restarts all instances of servers. When the -s option is specified, restarts the server specified by -s. restart-admin Restarts the administration server.	
restoreconfig	Restores the most recently saved Administration Server configuration information to the NetscapeRoot partition under /var/ds5/slapd-serverID/confbak.	
sagt -c CONFIG	Start proxy SNMP agent. For more information, see the 《iPlanet Directory Server 5.1 Administrator's Guide》. The sagt subcommand supports the following options:	

	<i>-c configfile</i> The CONFIG file includes the port that the SNMP daemon listens to. It also needs to include the MIB trees and traps that the proxy SNMP agent forwards. Edit the CONFIG file located in <code>/usr/iplanet/ds5/plugins/snmp/sagt.</code>
saveconfig	Saves the administration server configuration information to the <code>/var/ds5/slapd-serverID/confbak</code> directory.
setup [-f <i>configuration_file</i>]	Configures an instance of the directory server or administration server. Creates a basic configuration for the directory server and the administrative server that is used to manage the directory. The setup subcommand has two modes of operation. You can invoke it with a curses-based interaction to gather input. Alternatively, you can provide input in a configuration file using the <i>-f</i> option. The setup subcommand supports the following option: <i>-f configuration_file</i> Specifies the configuration file for silent installation.
start	Starts the directory server. When the <i>-s</i> option is not specified, starts servers of all instances. When the <i>-s</i> option is specified, starts the server instance specified by <i>-s</i>.
start-admin	Starts the directory server. When the <i>-s</i> option is not specified, restarts all instances of servers. When the <i>-s</i> option is specified, restarts the server specified by <i>-s</i>.
startconsole	Starts the directory console..
stop	Stops the directory server. When the <i>-s</i> option is not specified, restarts all instances of servers. When the <i>-s</i> option is specified, restarts the server specified by <i>-s</i>.
stop-admin	Stop the administration server.

suffix2instance { -s <i>suffix</i> }	Map a suffix to a backend name. Specify -s <i>suffix</i> as the suffix to be mapped to the backend.								
uninstall	Uninstalls the directory server and the administration server. This subcommand stops servers of all instances and removes all the changes created by setup.								
vlvindex <i>args</i>	Create virtual list view (VLV) indexes, known in the Directory Server Console as Browsing Indexes. The server must be stopped beforehand. The vlvindex subcommand supports the following arguments: <table><tr><td>-d <i>debug_level</i></td><td>Specify the debug level to use during index creation. Debug levels are defined in <code>nsslapd-errorlog-level</code> (error Log Level). See the «iPlanet Directory Server 5.1 Configuration, Command, and File Reference» .</td></tr><tr><td>-n <i>backend_instance</i></td><td>Name of the database containing the entries to index.</td></tr><tr><td>-s <i>suffix</i></td><td>Name of the suffix containing the entries to index.</td></tr><tr><td>-T <i>VLVTag</i></td><td>Name of the database containing the entries to index.</td></tr></table>	-d <i>debug_level</i>	Specify the debug level to use during index creation. Debug levels are defined in <code>nsslapd-errorlog-level</code> (error Log Level). See the «iPlanet Directory Server 5.1 Configuration, Command, and File Reference» .	-n <i>backend_instance</i>	Name of the database containing the entries to index.	-s <i>suffix</i>	Name of the suffix containing the entries to index.	-T <i>VLVTag</i>	Name of the database containing the entries to index.
-d <i>debug_level</i>	Specify the debug level to use during index creation. Debug levels are defined in <code>nsslapd-errorlog-level</code> (error Log Level). See the «iPlanet Directory Server 5.1 Configuration, Command, and File Reference» .								
-n <i>backend_instance</i>	Name of the database containing the entries to index.								
-s <i>suffix</i>	Name of the suffix containing the entries to index.								
-T <i>VLVTag</i>	Name of the database containing the entries to index.								

选项

Options for the `directoryserver` command itself must appear before the subcommand argument.

The following options are supported:

-s <i>server-instance</i>	
-server <i>server-instance</i>	The server instance name. Specify the directory server instance to process the command against. For some of the listed

subcommands the server instance is optional and for other sub commands it is a required option.

示例

示例 1 Starting All Instances of the Directory Servers

The following command starts all the instances of the directory servers:

```
example% directoryserver start
```

示例 2 Starting the Instances of myhost of the Directory Server

The following command starts the instances myhost of the directory server.

```
example% directoryserver -s myhost start
```

示例 3 Running the Monitor Tool and Outputting the Current Status

The following command runs the monitor tool and output the current status of the ephesus directory instance.

```
example% directoryserver -s ephesus monitor
```

示例 4 Running the idsktune Tool and Outputting Performance Tuning Information

The following command runs the idsktune tool and outputs performance tuning information:

```
example% directoryserver idsktune
```

退出状态

The following exit values are returned:

- 0 Successful completion.
- non-zero An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	IPLTdsr, IPLTdsu

另请参见

- 《iPlanet Directory Server 5.1 Administrator’s Guide》
- 《iPlanet Directory Server 5.1 Configuration, Command, and File Reference》

引用名	<code>disks</code> – creates <code>/dev</code> entries for hard disks attached to the system
用法概要	<code>/usr/sbin/disks [-C] [-r rootdir]</code>
描述	<code>devfsadm(1M)</code> is now the preferred command for <code>/dev</code> and should be used instead of <code>disks</code>. <code>disks</code> creates symbolic links in the <code>/dev/dsk</code> and <code>/dev/rdisk</code> directories pointing to the actual disk device special files under the <code>/devices</code> directory tree. It performs the following steps: 1. <code>disks</code> searches the kernel device tree to see what hard disks are attached to the system. It notes the <code>/devices</code> pathnames for the slices on the drive and determines the physical component of the corresponding <code>/dev/dsk</code> or <code>/dev/rdisk</code> name. 2. The <code>/dev/dsk</code> and <code>/dev/rdisk</code> directories are checked for disk entries – that is, symbolic links with names of the form <code>cN[tN]dNsN</code>, or <code>cN[tN]dNpN</code>, where <code>N</code> represents a decimal number. <code>cN</code> is the logical controller number, an arbitrary number assigned by this program to designate a particular disk controller. The first controller found on the first occasion this program is run on a system, is assigned number 0. <code>tN</code> is the bus-address number of a subsidiary controller attached to a peripheral bus such as SCSI or IPI (the target number for SCSI, and the <code>facility</code> number for IPI controllers). <code>dN</code> is the number of the disk attached to the controller. <code>sN</code> is the <i>slice</i> number on the disk. <code>pN</code> is the FDISK partition number used by <code>fdisk(1M)</code>. (x86 Only) 3. If only some of the disk entries are found in <code>/dev/dsk</code> for a disk that has been found under the <code>/devices</code> directory tree, <code>disks</code> creates the missing symbolic links. If none of the entries for a particular disk are found in <code>/dev/dsk</code>, <code>disks</code> checks to see if any entries exist for other disks attached to the same controller, and if so, creates new entries using the same controller number as used for other disks on the same controller. If no other <code>/dev/dsk</code> entries are found for slices of disks belonging to the same physical controller as the current disk, <code>disks</code> assigns the lowest-unused controller number and creates entries for the disk slices using this newly-assigned controller number. <code>disks</code> is run automatically each time a reconfiguration-boot is performed or when <code>add_drv(1M)</code> is executed. When invoking <code>disks</code> manually, first run <code>drvconfig(1M)</code> to ensure <code>/devices</code> is consistent with the current device configuration. Notice to Driver Writers <code>disks</code> considers all devices with a node type of <code>DDI_NT_BLOCK</code>, <code>DDI_NT_BLOCK_CHAN</code>, <code>DDI_NT_CD</code>, <code>DDI_NT_BLOCK_WWN</code> or <code>DDI_NT_CD_CHAN</code> to be disk devices. <code>disks</code> requires the minor name of disk devices obey the following format conventions. The minor name for block interfaces consists of a single lowercase ASCII character, <code>a</code> through <code>u</code>, representing the slices and the primary partitions. The minor name for logical drive block interfaces consists of the strings <code>p5</code> through <code>p36</code>. The minor name for character (raw) interfaces consists of a single lowercase ASCII character, <code>a</code> through <code>a</code>, followed by the string <code>,raw</code>, representing the slices and the primary partitions. The minor name for logical drive character (raw) interfaces consists of the string <code>p5</code> through <code>p36</code> followed by <code>,raw</code>. <code>disks</code> performs the following translations:

- a through p to s0 through s15
- q through u to p0 through p4
- p5 through p36 to p5 through p36

SPARC drivers should only use the first eight slices: a through h, while x86 drivers can use a through u, with q through u corresponding to [fdisk\(1M\)](#) primary partitions. q represents the entire disk, while r, s, t, and u represent up to four additional primary partitions. For logical drives, p5 to p36 correspond to the 32 logical drives that are supported. The device nodes for logical drives change dynamically as and when they are created or deleted.

To prevent disks from attempting to automatically generate links for a device, drivers must specify a private node type and refrain from using a node type: `DDI_NT_BLOCK`, `DDI_NT_BLOCK_CHAN`, `DDI_NT_CD`, or `DDI_NT_CD_CHAN` when calling [ddi_create_minor_node\(9F\)](#).

选项

The following options are supported:

- C Causes disks to remove any invalid links after adding any new entries to `/dev/dsk` and `/dev/rsk`. Invalid links are links which refer to non-existent disk nodes that have been removed, powered off, or are otherwise inaccessible.
- r *rootdir* Causes disks to presume that the `/dev/dsk`, `/dev/rsk` and `/devices` directory trees are found under *rootdir*, not directly under `/`.

错误

If disks finds entries of a particular logical controller linked to different physical controllers, it prints an error message and exits without making any changes to the `/dev` directory, since it cannot determine which of the two alternative logical-to-physical mappings is correct. The links should be manually corrected or removed before another reconfiguration-boot is performed.

示例

示例 1 Creating Block and Character Minor Devices

The following example demonstrates creating the block and character minor devices from within the `xkdisk` driver's [attach\(9E\)](#) function.

```
#include    <sys/dkio.h>
/*
 * Create the minor number by combining the instance number
 * with the slice number.
 */
#define MINOR_NUM(i, s)    ((i) << 4 | (s))

int
xkdiskattach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    int instance, slice;
    char name[8];
```

示例1 Creating Block and Character Minor Devices (续)

```

/* other stuff in attach... */

instance = ddi_get_instance(dip);
for (slice = 0; slice < V_NUMPAR; slice++) {
/*
 * create block device interface
 */
sprintf(name, "%c", slice + 'a');
ddi_create_minor_node(dip, name, S_IFBLK,
    MINOR_NUM(instance, slice), DDI_NT_BLOCK_CHAN, 0);

/*
 * create the raw (character) device interface
 */
sprintf(name, "%c,raw", slice + 'a');
ddi_create_minor_node(dip, name, S_IFCHR,
    MINOR_NUM(instance, slice), DDI_NT_BLOCK_CHAN, 0);
}
}

```

Installing the xkdisk disk driver on a Sun Fire 4800, with the driver controlling a SCSI disk (target 3 attached to an [isp\(7D\)](#) SCSI HBA) and performing a reconfiguration-boot (causing disks to be run) creates the following special files in /devices.

```

# ls -l /devices/ssm@0,0/pci@18,700000/pci@1/SUNW,isp@two@4/
brw-r----- 1 root sys 32, 16 Aug 29 00:02 xkdisk@3,0:a
crw-r----- 1 root sys 32, 16 Aug 29 00:02 xkdisk@3,0:a,raw
brw-r----- 1 root sys 32, 17 Aug 29 00:02 xkdisk@3,0:b
crw-r----- 1 root sys 32, 17 Aug 29 00:02 xkdisk@3,0:b,raw
brw-r----- 1 root sys 32, 18 Aug 29 00:02 xkdisk@3,0:c
crw-r----- 1 root sys 32, 18 Aug 29 00:02 xkdisk@3,0:c,raw
brw-r----- 1 root sys 32, 19 Aug 29 00:02 xkdisk@3,0:d
crw-r----- 1 root sys 32, 19 Aug 29 00:02 xkdisk@3,0:d,raw
brw-r----- 1 root sys 32, 20 Aug 29 00:02 xkdisk@3,0:e
crw-r----- 1 root sys 32, 20 Aug 29 00:02 xkdisk@3,0:e,raw
brw-r----- 1 root sys 32, 21 Aug 29 00:02 xkdisk@3,0:f
crw-r----- 1 root sys 32, 21 Aug 29 00:02 xkdisk@3,0:f,raw
brw-r----- 1 root sys 32, 22 Aug 29 00:02 xkdisk@3,0:g
crw-r----- 1 root sys 32, 22 Aug 29 00:02 xkdisk@3,0:g,raw
brw-r----- 1 root sys 32, 23 Aug 29 00:02 xkdisk@3,0:h
crw-r----- 1 root sys 32, 23 Aug 29 00:02 xkdisk@3,0:h,raw

```

/dev/dsk will contain the disk entries to the block device nodes in /devices

```

# ls -l /dev/dsk
/dev/dsk/c0t3d0s0 -> ../../devices/[...]/xkdisk@3,0:a

```

示例 1 Creating Block and Character Minor Devices (续)

```
/dev/dsk/c0t3d0s1 -> ../../devices/[...]/xkdisk@3,0:b
/dev/dsk/c0t3d0s2 -> ../../devices/[...]/xkdisk@3,0:c
/dev/dsk/c0t3d0s3 -> ../../devices/[...]/xkdisk@3,0:d
/dev/dsk/c0t3d0s4 -> ../../devices/[...]/xkdisk@3,0:e
/dev/dsk/c0t3d0s5 -> ../../devices/[...]/xkdisk@3,0:f
/dev/dsk/c0t3d0s6 -> ../../devices/[...]/xkdisk@3,0:g
/dev/dsk/c0t3d0s7 -> ../../devices/[...]/xkdisk@3,0:h
```

and /dev/rdsk will contain the disk entries for the character device nodes in /devices

```
# ls -l /dev/rdsk
/dev/rdsk/c0t3d0s0 -> ../../devices/[...]/xkdisk@3,0:a,raw
/dev/rdsk/c0t3d0s1 -> ../../devices/[...]/xkdisk@3,0:b,raw
/dev/rdsk/c0t3d0s2 -> ../../devices/[...]/xkdisk@3,0:c,raw
/dev/rdsk/c0t3d0s3 -> ../../devices/[...]/xkdisk@3,0:d,raw
/dev/rdsk/c0t3d0s4 -> ../../devices/[...]/xkdisk@3,0:e,raw
/dev/rdsk/c0t3d0s5 -> ../../devices/[...]/xkdisk@3,0:f,raw
/dev/rdsk/c0t3d0s6 -> ../../devices/[...]/xkdisk@3,0:g,raw
/dev/rdsk/c0t3d0s7 -> ../../devices/[...]/xkdisk@3,0:h,raw
```

- 文件 /dev/dsk/* Disk entries (block device interface)
- /dev/rdsk/* Disk entries (character device interface)
- /devices/* Device special files (minor device nodes)

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [add_drv\(1M\)](#), [devfsadm\(1M\)](#), [fdisk\(1M\)](#), [attributes\(5\)](#), [isp\(7D\)](#), [devfs\(7FS\)](#), [dkio\(7I\)](#), [attach\(9E\)](#), [ddi_create_minor_node\(9F\)](#)

《Writing Device Drivers》

已知问题 disks silently ignores malformed minor device names.

引用名	diskscan – perform surface analysis
用法概要	diskscan [-W] [-n] [-y] <i>raw_device</i>
描述	diskscan is used by the system administrator to perform surface analysis on a portion of a hard disk. The disk portion may be a raw partition or slice; it is identified using its raw device name. By default, the specified portion of the disk is read (non-destructive) and errors reported on standard error. In addition, a progress report is printed on standard out. The list of bad blocks should be saved in a file and later fed into addbadsec(1M), which will remap them.
选项	The following options are supported: -n Causes diskscan to suppress linefeeds when printing progress information on standard out. -W Causes diskscan to perform write and read surface analysis. This type of surface analysis is destructive and should be invoked with caution. -y Causes diskscan to suppress the warning regarding destruction of existing data that is issued when -W is used.
操作数	The following operands are supported: <i>raw_device</i> The address of the disk drive (see FILES).
文件	The raw device should be /dev/rdisk/c?[t?][d?][ps]?. See disks(1M) for an explanation of SCSI and IDE device naming conventions.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86
Availability	system/core-os

另请参见 [addbadsec\(1M\)](#), [disks\(1M\)](#), [fdisk\(1M\)](#), [fmthard\(1M\)](#), [format\(1M\)](#), [attributes\(5\)](#)

附注 The [format\(1M\)](#) utility is available to format, label, analyze, and repair SCSI disks. This utility is included with the diskscan, [addbadsec\(1M\)](#), [fdisk\(1M\)](#), and [fmthard\(1M\)](#) commands available for x86. To format an IDE disk, use the DOS format utility; however, to label, analyze, or repair IDE disks on x86 systems, use the Solaris [format\(1M\)](#) utility.

引用名	dispadmin – process scheduler administration
用法概要	<pre>dispadmin -l</pre> <pre>dispadmin -c class {-g [-r res] -s file}</pre> <pre>dispadmin -d [class]</pre>
描述	The <code>dispadmin</code> command displays or changes process scheduler parameters while the system is running. <code>dispadmin</code> does limited checking on the values supplied in <i>file</i> to verify that they are within their required bounds. The checking, however, does not attempt to analyze the effect that the new values have on the performance of the system. Inappropriate values can have a negative effect on system performance. (See 《Oracle Solaris 管理：常见任务》.)
选项	The following options are supported: -c <i>class</i> Specifies the class whose parameters are to be displayed or changed. Valid <i>class</i> values are: RT for the real-time class, TS for the time-sharing class, IA for the inter-active class, FSS for the fair-share class, and FX for the fixed-priority class. The time-sharing and inter-active classes share the same scheduler, so changes to the scheduling parameters of one will change those of the other. -d [<i>class</i>] Sets or displays the name of the default scheduling class to be used on reboot when starting <code>svc:/system/scheduler:default</code>. If class name is not specified, the name and description of the current default scheduling class is displayed. If class name is specified and is a valid scheduling class name, then it is saved in <code>dispadmin</code>'s private configuration file <code>/etc/dispadmin.conf</code>. Only super-users can set the default scheduling class. -g Gets the parameters for the specified class and writes them to the standard output. Parameters for the real-time class are described in rt_dptbl(4). Parameters for the time-sharing and inter-active classes are described in ts_dptbl(4). Parameters for the fair-share class are described in FSS(7). Parameters for the fixed-priority class are described in fx_dptbl(4). The -g and -s options are mutually exclusive: you may not retrieve the table at the same time you are overwriting it. -l Lists the scheduler classes currently configured in the system. -r <i>res</i> When using the -g option you may also use the -r option to specify a resolution to be used for outputting the time quantum values. If no resolution is specified, time quantum values are in milliseconds. If <i>res</i> is specified it must be a positive integer between 1 and 1000000000 inclusive, and the resolution used is the reciprocal of <i>res</i> in seconds. For

example, a *res* value of 10 yields time quantum values expressed in tenths of a second; a *res* value of 1000000 yields time quantum values expressed in microseconds. If the time quantum cannot be expressed as an integer in the specified resolution, it is rounded up to the next integral multiple of the specified resolution.

-s *file*

Sets scheduler parameters for the specified class using the values in *file*. These values overwrite the current values in memory—they become the parameters that control scheduling of processes in the specified class. The values in *file* must be in the format output by the *-g* option. Moreover, the values must describe a table that is the same size (has same number of priority levels) as the table being overwritten. Super-user privileges are required in order to use the *-s* option.

Specify time quantum values for scheduling classes in system clock ticks, and not in constant-time units. Time quantum values are based on the value of the kernel's *hz* variable. If kernel variable *hires_tick* is set to 1 to get higher resolution clock behavior, the actual time quanta will be reduced by the order of 10.

The *-g* and *-s* options are mutually exclusive: you may not retrieve the table at the same time you are overwriting it.

示例

示例 1 Retrieving the Current Scheduler Parameters for the real-time class

The following command retrieves the current scheduler parameters for the real-time class from kernel memory and writes them to the standard output. Time quantum values are in microseconds.

```
dispadmin -c RT -g -r 1000000
```

示例 2 Overwriting the Current Scheduler Parameters for the Real-time Class

The following command overwrites the current scheduler parameters for the real-time class with the values specified in *rt.config*.

```
dispadmin -c RT -s rt.config
```

示例 3 Retrieving the Current Scheduler Parameters for the Time-sharing Class

The following command retrieves the current scheduler parameters for the time-sharing class from kernel memory and writes them to the standard output. Time quantum values are in nanoseconds.

```
dispadmin -c TS -g -r 1000000000
```

示例 4 Overwriting the Current Scheduler Parameters for the Time-sharing Class

The following command overwrites the current scheduler parameters for the time-sharing class with the values specified in *ts.config*.

```
dispadmin -c TS -s ts.config
```

文件 /etc/dispadmin.conf
Possible location for argument to -s option.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [prioctl\(1\)](#), [svcs\(1\)](#), [svcadm\(1M\)](#), [prioctl\(2\)](#), [fx_dptbl\(4\)](#), [rt_dptbl\(4\)](#), [ts_dptbl\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [FSS\(7\)](#)

 《Oracle Solaris 管理：常见任务》 《Programming Interfaces Guide》

诊断 dispadmin prints an appropriate diagnostic message if it fails to overwrite the current scheduler parameters due to lack of required permissions or a problem with the specified input file.

附注 The default scheduling class setting facility is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/system/scheduler:default

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Note that disabling the service while it is running will not change anything. The service's status can be queried using the [svcs\(1\)](#) command.

引用名

dladm – 管理数据链路

用法概要

```

dladm show-link [-PZ] [-s [-i interval]] [[-p] -o field[,...]]
    [-z zone[,...]] [link]
dladm rename-link [-R root-dir] link new-link

dladm delete-phys phys-link
dladm show-phys [-PZ] [-Lm] [[-p] -o field[,...]] [-H]
    [-z zone[,...]] [phys-link]

dladm create-aggr [-t] [-R root-dir] [-P policy] [-L mode]
    [-T time] [-u address] -l ether-link1 [-l ether-link2...] aggr-link
dladm modify-aggr [-t] [-R root-dir] [-P policy] [-L mode]
    [-T time] [-u address] aggr-link
dladm delete-aggr [-t] [-R root-dir] aggr-link
dladm add-aggr [-t] [-R root-dir] -l ether-link1 [-l ether-link2...]
    aggr-link
dladm remove-aggr [-t] [-R root-dir] -l ether-link1 [-l ether-link2...]
    aggr-link
dladm show-aggr [-PLxZ] [-s [-i interval]] [[-p] -o field[,...]]
    [-z zone[,...]] [aggr-link]

dladm create-bridge [-P protect] [-R root-dir] [-p priority]
    [-m max-age] [-h hello-time] [-d forward-delay] [-f force-protocol]
    [-l link...] bridge-name
dladm modify-bridge [-P protect] [-R root-dir] [-p priority]
    [-m max-age] [-h hello-time] [-d forward-delay] [-f force-protocol]
    bridge-name
dladm delete-bridge [-R root-dir] bridge-name
dladm add-bridge [-R root-dir] -l link [-l link...] bridge-name
dladm remove-bridge [-R root-dir] -l link [-l link...] bridge-name
dladm show-bridge [-flt] [-s [-i interval]] [[-p] -o field,...]
    [bridge-name]

dladm create-vlan [-ft] [-R root-dir] -l ether-link -v vid [vlan-link]
dladm delete-vlan [-t] [-R root-dir] vlan-link
dladm show-vlan [-PZ] [[-p] -o field[,...]] [-z zone[,...]] [vlan-link]

dladm scan-wifi [[-p] -o field[,...]] [wifi-link]
dladm connect-wifi [-e ssid] [-i bssid] [-k key,...]
    [-s none | wep | wpa] [-a open | shared] [-b bss | ibss] [-c]
    [-m a | b | g | n] [-T time] [wifi-link]
dladm disconnect-wifi [-a] [wifi-link]
dladm show-wifi [-Z] [[-p] -o field[,...]] [-z zone[,...]] [wifi-link]

dladm show-ether [-xZ] [[-p] -o field[,...]] [-z zone[,...]] [ether-link]

dladm set-linkprop [-t] [-R root-dir] -p prop=value[,...] link
dladm reset-linkprop [-t] [-R root-dir] [-p prop[,...]] link
dladm show-linkprop [-PZ] [[-c] -o field[,...]] [-p prop[,...]]
    [-z zone[,...]] [link]

```

```

dladm create-secobj [-t] [-R root-dir] [-f file] -c class secobj
dladm delete-secobj [-t] [-R root-dir] secobj[,...]
dladm show-secobj [-P] [[-p] -o field[,...]] [secobj,...]

dladm create-vnic [-t] -l link [-R root-dir] [-m value | auto |
    {factory [-n slot-identifier]} | {vrrp -A {inet | inet6} -V vrid}
    | {random [-r prefix]}] [-v vlan-id] [-p prop=value[,...]] vnic-link
dladm delete-vnic [-t] [-R root-dir] vnic-link
dladm show-vnic [-pPZ] [-s [-i interval]] [-o field[,...]]
    [-l link] [-z zone[,...]] [vnic-link]

dladm create-etherstub [-t] [-R root-dir] etherstub
dladm delete-etherstub [-t] [-R root-dir] etherstub
dladm show-etherstub [-Z] [-z zone[,...]] [etherstub]

dladm create-iptun [-t] [-R root-dir] -T type [-a {local|remote}=addr,...]
    iptun-link
dladm modify-iptun [-t] [-R root-dir] -a {local|remote}=addr,...
    iptun-link
dladm delete-iptun [-t] [-R root-dir] iptun-link
dladm show-iptun [-PZ] [[-p] -o field[,...]] [-z zone[,...]] [iptun-link]

dladm create-part [-t] [-f] -l ib-link [-R root-dir] -P pkey
    [-p prop=value[,...]] part-link
dladm delete-part [-t] [-R root-dir] part-link
dladm show-part [-pP] [-o field[,...]] [-l ib-link] [part-link]

dladm show-ib [-pP] [-o field[,...]] [ib-link]

dladm show-usage [-a] -f filename [-p plotfile -F format] [-s time]
    [-e time] [link]

```

描述

dladm 命令用于管理数据链路。数据链路在系统中表示为 STREAMS DLPI (v2) 接口，可在 TCP/IP 等协议栈下检测到此接口。各数据链路依赖于一台网络设备或设备聚合来向网络发送包或从网络接收包。

每个 **dladm** 子命令都在以下某个对象中运行：

link

数据链路，通过名称标识。一般说来，此名称可使用任何字母数字字符（或者下划线_或句点.），但必须以字母字符开头，并且以数字结尾。数据链路名称最多可包含 31 个字符，并且结尾的数字必须介于 0 到 4294967294（含 0 和 4294967294）之间。结尾的数字不能以零开头。建议采用包含 3 至 8 个字符的数据链路名称。

某些子命令仅在特定类型或种类的数据链路上运行。在这些情况下，请使用以下对象名称：

aggr-link

聚合数据链路（或键；请参见“说明”）。

ether-link

物理以太网数据链路。

iptun-link
IP 隧道链路。

part-link
InfiniBand (IB) 分区数据链路。

phys-link
物理数据链路。

vlan-link
VLAN 数据链路。

vnic-link
在链路或 **etherstub** 上创建的虚拟网络接口。伪设备可以计算机上的网络接口卡的方式处理。

wifi-link
WiFi 数据链路。

bridge
网桥实例，通过以管理方式选择的名称标识。此名称可使用任何字母数字字符或下划线 `_`，但必须以字母字符开头和结尾。网桥名称最多可包含 31 个字符。保留名称 **default**，同时保留所有以 **SUNW** 开头的名称。

请注意，在网桥名称后附加零 (**0**) 可生成一个用于监测的有效链路名称。

另请注意，随下面的 **dladm** 子命令一起介绍的与网桥相关的子命令要求安装 **pkg://solaris/network/bridging** 软件包。

dev
网络设备，通过串联驱动程序名称和实例编号标识。

etherstub
创建 VNIC 时，可以使用以太网桩模块（而非物理 NIC）。基于 **etherstub** 创建的 VNIC 将显示为通过虚拟交换机进行连接，并允许在没有物理硬件的情况下构建完整的虚拟网络。

part
在 IB 物理链路上创建的 IB 分区链路。

secobj
安全对象，通过以管理方式选择的名称标识。此名称可使用任何字母数字字符、下划线 (`_`)、句点 (`.`) 和连字符 (`-`)。安全对象名称最多可包含 32 个字符。

选项

每个 **dladm** 子命令都有各自的选项集。但是，许多子命令都具有下列共同选项：

-R root-dir、**--root-dir=root-dir**
指定应当应用创建、删除或重命名等操作的备用根目录。

子命令

支持以下子命令：

`dladm show-link [-PZ] [-s [-i interval]] [[-p] -o field[,...]] [-z zone[,...]] [link]`

显示所有数据链路或指定链路的链路配置信息。缺省情况下，对于每个已知的网络设备，系统都会配置一个数据链路。输出链路统计信息的选项已移至 `dlstat(1M)`。

`-o field[,...]`、`--output=field[,...]`

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。如果并非使用 `-s` 选项（如下文所述）修改，字段名称必须是下面列出的某个字段，或者是可以显示所有字段的特殊值 `all`。缺省情况下（无 `-o`），`show-link` 显示所有字段。

LINK

数据链路的名称。

ZONE

数据链路的当前区域。

CLASS

数据链路的类。`dladm` 可区分下列类：

aggr

IEEE 802.3ad 链路聚合。`show-aggr` 子命令显示此类数据链路的更多详细信息。

bridge

网桥实例，通过以管理方式选择的名称标识。

etherstub

`etherstub` 的实例。创建 VNIC 时，可以使用以太网桩模块（而非物理 NIC）。基于 `etherstub` 创建的 VNIC 将显示为通过虚拟交换机进行连接，并允许在没有物理硬件的情况下构建完整的虚拟网络。

iptun

IP 隧道链路的实例。

part

IP-over-IB 接口。`show-part` 子命令显示此类数据链路的更多详细信息。

phys

物理数据链路。`show-phys` 子命令显示此类数据链路的更多详细信息。

vlan

VLAN 数据链路。`show-vlan` 子命令显示此类数据链路的更多详细信息。

vnic

虚拟网络接口。`show-vnic` 子命令显示此类数据链路的更多详细信息。

MTU

显示的数据链路的最大传输单元大小。

STATE

数据链路的链路状态。此状态可以为 `up`、`down` 或 `unknown`。

BRIDGE

分配了此链路的网桥的名称（如果有）。

OVER

运行此数据链路所基于的物理数据链路。这适用于数据链路的 `aggr`、`bridge`、`vlan` 和 `part` 分区类。VLAN 或 IB 分区在一个物理数据链路上创建，网桥具有多个附加链路，而聚合则包括一个或多个物理数据链路。

当配合使用 `-o` 选项和 `-s` 选项（用于显示链路统计信息）时，字段名称必须是下面列出的某个字段，或者是可以显示所有字段的特殊值 `all`。

LINK

数据链路的名称。

IPACKETS

在此链路上收到的包的数目。

RBYTES

在此链路上收到的字节的数目。

IERRORS

输入错误数。

OPACKETS

在此链路上发送的包的数目。

OBYTES

在此链路上发送的字节数目。

OERRORS

输出错误数。

-p、**--parseable**

使用稳定的可供计算机解析的格式显示。`-o` 选项需要与 `-p` 一起使用。请参见下面的“可解析的输出格式”。

-P、**--persistent**

显示持久性链路配置。

-s、**--statistics**

显示链路统计信息。在 `dlstat(1M)` 中，此选项已过时。

-i interval、**--interval=interval**

与 `-s` 选项配合使用，用于指定应显示统计信息的时间间隔（以秒为单位）。在 `dlstat(1M)` 中，此选项已过时。

-Z

在输出中显示 `ZONE` 列。

-z zone[,...]

显示指定区域的链路。缺省情况下，当 `dladm` 从全局区域运行时，将显示所有区域中的链路。其他区域中的链路的显示方式是：以相应区域名称作为其前缀，后跟斜杠 (/) 分隔符。例如，`zone1/net0`

从非全局区域运行时，该子命令仅显示该区域中的链路。非全局区域无法看到其他区域中的链路。

dladm rename-link [-R root-dir] link new-link

将 *link* 重命名为 *new-link*。此选项用于为链路指定一个有意义的名称，或者将现有链路配置（例如，已删除设备的链路属性）与新设备相关联。有关如何使用此子命令的具体示例，请参见“示例”部分。

-R root-dir、**--root-dir=root-dir**

请参见上面的“选项”。

dladm delete-phys phys-link

此命令用于删除与物理硬件（已从系统中删除）相关联的链路的持久性配置。请参见“示例”部分。

dladm show-phys [-PZ] [-Lm] [[-p] -o field[,...]] [-H] [-z zone[,...]] [phys-link]

显示所有物理链路或指定物理链路的物理设备和属性。如果没有 **-P**，则仅显示正在运行的系统上的可用物理链路。

-H

显示 NIC 驱动程序返回的硬件资源使用情况。**-H** 输出显示下列元素：

LINK

与 NIC 驱动程序对应的物理设备。

RINGTYPE

环类型（RX 或 TX）。

RINGS

环索引。环是通常映射到 DMA 通道的可针对特殊用途进行编程的硬件资源。例如，可对 RX 环进行编程，使之仅接收隶属于特定 MAC 地址的包。

CLIENTS

使用环的 MAC 客户机。

-L

显示物理设备/链路的位置信息。按照位置顺序显示输出（即，板载设备位于扩展槽前面），如果可用，则提供位置信息（例如，PCIexp 插槽 2，MB）。**-L** 的输出支持以下元素：

LINK

与 NIC 驱动程序对应的物理设备。

DEVICE

此链路下的物理设备的名称。

LOC

物理位置描述字符串（如果可用）。

-m

显示出厂 MAC 地址及其插槽标识符和可用性的列表。

-o *field*、--output=*field*

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名称必须是下面列出的某个字段，或者是可以显示所有字段的特殊值 `all`。对于每个链路，可显示下列字段：

LINK

数据链路的名称。

MEDIA

物理数据链路提供的介质类型。

STATE

链路的状态。此状态可以为 `up`、`down` 或 `unknown`。

SPEED

链路的当前速度（以兆位/秒为单位）。

DUPLEX

对于以太网链路，如果链路状态为 `up`，则会显示此链路的全/半双工状态。在其他所有情况下，双工显示为 `unknown`。

DEVICE

此链路下的物理设备的名称。

-p、--parseable

使用稳定的可供计算机解析的格式显示。`-o` 选项需要与 `-p` 一起使用。请参见下面的“可解析的输出格式”。

-P、--persistent

此选项显示所有链路（包括已从系统中删除的链路）的持久性配置。此输出提供了一个 **FLAGS** 列，在此列中，`r` 标志表示与物理链路相关联的物理设备已移除。对于此类链路，使用 `delete-phys` 可从系统中清除链路的配置。

-Z

在输出中显示 **ZONE** 列。

-z *zone*[,...]

请参见上面 `dladm show-link` 下 `-z` 选项的描述。

```
dladm create-aggr [-t] [-R root-dir] [-P policy] [-L mode] [-T time] [-u address] -l
ether-link1 [-l ether-link2...] aggr-link
```

将一组链路组合为一个名为 *aggr-link* 的 IEEE 802.3ad 链路聚合。此外，还支持使用整数 *key* 生成聚合链路名称，以便实现向后兼容。下面的许多 `*-aggr` 子命令还支持使用 *key* 引用给定聚合，但建议使用聚合链路名称。有关键的更多信息，请参见“说明”部分。

dladm 支持针对端口聚合的多种端口选择策略。（请参见下面的 -P 选项的说明。）如果未指定策略，create-aggr 将使用 -P 选项下介绍的缺省 L4 策略。

-l *ether-link*、--link=*ether-link*

使用 -l 选项（后跟要包含在聚合中的链路的名称），可以指定聚合中的每个以太网链路（或端口）。通过指定多个 -l 选项，聚合可以包括多个链路。为了与 Solaris 的早期版本实现向后兼容，dladm 命令还支持使用带有设备名称的 -d 选项（或 --dev），根据链路的底层设备名称来指定链路。采用 -l 选项的其他 *-aggr 子命令还接受 -d。

-t、--temporary

指定聚合为临时聚合。临时聚合将持续到下一次重新引导。

-R *root-dir*、--root-dir=*root-dir*

请参见上面的“选项”。

-P *policy*、--policy=*policy*

指定用于外发通信的负荷分配的端口选择策略。此策略指定用于发送包的 *dev* 对象。策略是以逗号分隔的一个或多个层说明符的列表。层说明符为下列项之一：

L2

根据包的源和目标 MAC 地址选择外发设备。

L3

根据包的源和目标 IP 地址选择外发设备。

L4

根据包中包含的上层协议信息选择外发设备。对于 TCP 和 UDP，这包括源和目标端口。对于 IPsec，这包括 SPI（Security Parameters Index，安全参数索引）。

例如，要使用上层协议信息，可使用以下策略：

-P L4

请注意，策略 L4 为缺省值。

要使用源和目标 MAC 地址以及源和目标 IP 地址，可使用以下策略：

-P L2,L3

-L *mode*、--lACP-mode=*mode*

指定是否应使用 LACP，以及使用的话应采用何种运行模式。支持的值包括 off、active 或 passive。

-T *time*、--lACP-timer=*time*

指定 LACP 计时器值。支持的值包括 short 或 long。

-u *address*、--unicast=*address*

指定用于聚合的固定单点传送硬件地址。如果未指定此选项，则会从组件设备的地址集中自动选择一个地址。

`dladm modify-aggr [-t] [-R root-dir] [-P policy] [-L mode] [-T time] [-u address] aggr-link`
 修改指定聚合的参数。

`-t`、`--temporary`

指定修改为临时修改。临时聚合将持续到下一次重新引导。

`-R root-dir`、`--root-dir=root-dir`

请参见上面的“选项”。

`-P policy`、`--policy=policy`

指定用于外发通信的负荷分配的端口选择策略。有关有效策略值的说明，请参见 `dladm create-aggr`。

`-L mode`、`--lACP-mode=mode`

指定是否应使用 LACP，以及使用的话应采用何种运行模式。支持的值包括 `off`、`active` 或 `passive`。

`-T time`、`--lACP-timer=time`

指定 LACP 计时器值。支持的值包括 `short` 或 `long`。

`-u address`、`--unicast=address`

指定用于聚合的固定单点传送硬件地址。如果未指定此选项，则会从组件设备的地址集中自动选择一个地址。

（请注意，修改固定的单点传送硬件地址时，将覆盖以前为聚合定义的所有已定义 `mac-address` 链路属性。请参见“常规链路属性”。）

`dladm delete-aggr [-t] [-R root-dir] aggr-link`

删除指定的聚合。

`-t`、`--temporary`

指定删除为临时删除。临时删除将持续到下一次重新引导。

`-R root-dir`、`--root-dir=root-dir`

请参见上面的“选项”。

`dladm add-aggr [-t] [-R root-dir] -l ether-link1 [--link=ether-link2...] aggr-link`

向指定聚合添加链路。

`-l ether-link`、`--link=ether-link`

指定要添加到聚合的以太网链路。通过提供多个 `-l` 选项，可以添加多个链路。

`-t`、`--temporary`

指定添加为临时添加。临时添加将持续到下一次重新引导。

`-R root-dir`、`--root-dir=root-dir`

请参见上面的“选项”。

`dladm remove-aggr [-t] [-R root-dir] -l ether-link1 [--l=ether-link2...] aggr-link`

从指定聚合中删除链路。

-l ether-link、**--link=ether-link**

指定要从聚合中删除的以太网链路。通过提供多个 **-l** 选项，可以添加多个链路。

-t、**--temporary**

指定删除为临时删除。临时删除将持续到下一次重新引导。

-R root-dir、**--root-dir=root-dir**

请参见上面的“选项”。

dladm show-aggr [-PLxZ] [-s [-i interval]] [[-p] -o field[,...]] [-z zone[,...]] [aggr-link]

显示所有聚合或指定聚合的聚合配置（缺省值）或 LACP 信息。

缺省情况下（无任何选项），可显示下列字段：

LINK

聚合链路的名称。

POLICY

聚合的 LACP 策略。有关可能的值的说明，请参见 **create-aggr -P** 选项。

ADDRPOLICY

如果将聚合配置为自动配置其单点传送 MAC 地址（未使用 **-u** 选项创建或修改聚合时的缺省值），则为 **auto**；如果已使用 **-u** 设置固定 MAC 地址，则为 **fixed**。

LACPACTIVITY

聚合的 LACP 模式。可能的值包括 **off**、**active** 或 **passive**，这些值使用 **create-aggr** 或 **modify-aggr** 的 **-l** 选项进行设置。

LACPTIMER

使用 **create-aggr** 或 **modify-aggr** 的 **-T** 选项设置的聚合的 LACP 计时器值。

FLAGS

与聚合相关联的一组状态标志。可能的唯一标志为 **f**，该标志在管理员强制使用 **create-aggr** 的 **-f** 选项创建聚合时显示。将来可定义其他标志。

show-aggr 命令接受下列选项：

-L、**--laccp**

显示聚合链路和各底层端口的详细 LACP 信息。此选项显示的大多数状态信息都根据 IEEE 802.3 定义。使用此选项，可显示下列字段：

LINK

聚合链路的名称。

PORT

某个底层聚合端口的名称。

AGGREGATABLE

是否可以将端口添加到聚合中。

SYNC

如果为 **yes**，则系统会考虑要同步的端口及部分聚合。

COLL

如果为 **yes**，将支持在关联的端口上收集传入帧。

DIST

如果为 **yes**，将支持在关联的端口上分发传出帧。

DEFAULTED

如果为 **yes**，则表示端口使用的是缺省合作伙伴信息（即，尚未从 LACP 合作伙伴处收到 LACP 数据）。

EXPIRED

如果为 **yes**，则表示端口的接收状态为 **EXPIRED** 状态。

-x、--extended

显示其他聚合信息，其中包括有关各底层端口的详细信息。使用 **-x**，可以显示下列字段：

LINK

聚合链路的名称。

PORT

某个底层聚合端口的名称。

SPEED

链路或端口的速度（以兆位/秒为单位）。

DUPLEX

如果链路状态为 **up**，则会显示链路或端口的全/半双工状态。在其他所有情况下，双工状态显示为 **unknown**。

STATE

链路状态。此状态可以为 **up**、**down** 或 **unknown**。

ADDRESS

链路或端口的 MAC 地址。

PORTSTATE

此选项指示各聚合端口是处于 **standby** 状态还是 **attached** 状态。

-o field[,...]、--output=field[,...]

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名称必须是上面列出的某个字段，或者是可以显示所有字段的特殊值 **all**。适用于 **-o** 选项的字段仅限于在每种输出模式下列出的字段。例如，如果使用 **-L**，只有上面的 **-L** 下列出的字段才能与 **-o** 配合使用。

-p、--parseable

使用稳定的可供计算机解析的格式显示。**-o** 选项需要与 **-p** 一起使用。请参见下面的“可解析的输出格式”。

-P、**--persistent**

显示持久性聚合配置（而非正在运行的系统的状态）。

-s、**--statistics**

显示聚合统计信息。在 **dlstat(1M)** 中，此选项已过时。

-i interval、**--interval=interval**

与 **-s** 选项配合使用，用于指定应显示统计信息的时间间隔（以秒为单位）。在 **dlstat(1M)** 中，此选项已过时。

-Z

在输出中显示 ZONE 列。

-z zone[,...]

请参见上面 **dladm show-link** 下 **-z** 选项的描述。

dladm create-bridge [**-P protect**] [**-R root-dir**] [**-p priority**] [**-m max-age**] [**-h hello-time**] [**-d forward-delay**] [**-f force-protocol**] [**-l link...**] **bridge-name**

创建一个 802.1D 网桥实例，并将一个或多个网络链路分配给新网桥（可选）。缺省情况下，系统上不存在任何网桥实例。

为了在链路之间进行桥接，必须至少创建一个网桥实例。每个网桥实例都是独立的，并且网桥之间没有任何转发连接。

请注意，与网桥相关的子命令（其中包括 **create-bridge**）要求安装 **pkg://solaris/network/bridging** 软件包。

-P protect、**--protect=protect**

指定保护方法。定义的保护方法包括表示跨树协议的 **stp** 和表示 TRILL 的 **trill**（在 RBridge 上使用）。缺省值为 **stp**。

-R root-dir、**--root-dir=root-dir**

请参见上面的“选项”。

-p priority、**--priority=priority**

指定网桥优先级。这会设置用于确定网络中的根网桥节点的 IEEE STP 优先级值。缺省值为 **32768**。有效值介于 **0**（最高优先级）和 **61440**（最低优先级）之间，并以 **4096** 为增量。

如果使用了不能被 4096 除尽的值，系统会以无提示方式向下舍入到可被 4096 除尽的下一较小值。

-m max-age、**--max-age=max-age**

指定配置信息的最长使用寿命（以秒为单位）。这会设置 STP 网桥最长使用寿命参数。如果此节点为根网桥，则会将此值用于网络中的所有节点。将放弃在此时间之前的网桥链路信息。此值缺省为 20 秒。有效值介于 6 至 40 秒之间。有关其他约束，请参见 **-d forward-delay** 参数。

-h *hello-time*、--hello-time=*hello-time*

指定 STP 网桥 Hello 时间参数。如果此节点为根节点，则会按此时间间隔在整个网络中发送配置 BPDU。缺省值为 2 秒。有效值介于 1 至 10 秒之间。有关其他约束，请参见 **-d *forward-delay*** 参数。

-d *forward-delay*、--forward-delay=*forward-delay*

指定 STP 网桥转发延迟参数。如果此节点为根节点，当启用端口时，网络中的所有网桥都会使用此计时器确定链路状态的顺序。缺省值为 15 秒。有效值介于 4 至 30 秒之间。

网桥必须遵循下列两个约束：

$$2 * (\text{forward-delay} - 1.0) \geq \text{max-age}$$

$$\text{max-age} \geq 2 * (\text{hello-time} + 1.0)$$

任何违反这些约束的参数设置都将视为错误，并且会导致命令失败，同时出现诊断消息。此消息提供了所提供的值的有效备用项。

-f *force-protocol*、--force-protocol=*force-protocol*

指定 MSTP 强制最大支持协议。缺省值为 3（第 1 部分）。有效值为非负整数。当前实现不支持 RSTP 或 MSTP，因此目前这不起作用。但是，为了防止在将来使用 MSTP，此参数可设置为 0（仅限 STP）或 2（适用于 STP 和 RSTP）。

-l *link*、--link=*link*

指定要添加到新创建的网桥的一个或多个链路。这与使用 **add-bridge** 子命令创建网桥并添加一个或多个链路相似。但是，如果无法添加任何链路，整个命令将失败，并且无法创建新网桥自身。要在同一命令行中添加多个链路，请对每个链路重复此选项。允许创建没有链路的网桥。有关链路分配的更多信息，请参见 **add-bridge** 子命令。

创建网桥和分配链路需要 **PRIV_SYS_DL_CONFIG** 特权。如果未在系统上安装可选网桥功能，网桥创建可能失败。

dladm modify-bridge [-P *protect*] [-R *root-dir*] [-p *priority*] [-m *max-age*] [-h *hello-time*] [-d *forward-delay*] [-f *force-protocol*] [-l *link...*] *bridge-name*

修改现有网桥的操作参数。这些选项与 **create-bridge** 子命令的选项相同，但不允许 **-l** 选项。要向现有网桥添加链路，请使用 **add-bridge** 子命令。

修改网桥参数需要 **PRIV_SYS_DL_CONFIG** 特权。

dladm delete-bridge [-R *root-dir*] *bridge-name*

删除网桥实例。删除的网桥不能有任何附加链路。在删除网桥之前，请使用 **remove-bridge** 子命令停用链路。

删除网桥需要 **PRIV_SYS_DL_CONFIG** 特权。

-R (--root-dir) 选项与 **create-bridge** 子命令的选项相同。

```
dladm add-bridge [-R root-dir] -l link [-l link...] bridge-name
```

向现有网桥添加一个或多个链路。如果指定了多个链路，并且添加其中任何一个链路导致错误，该命令将失败，并且不会对系统进行任何更改。

向网桥添加链路需要 `PRIV_SYS_DL_CONFIG` 特权。

链路最多只能是一个网桥的成员。尝试添加已属于另一个网桥的链路时将出错。要将链路从一个网桥实例移至另一个网桥实例，请将其从当前网桥中删除，然后再添加到新的网桥实例中。

此外，分配给网桥的链路还不能是 VLAN、VNIC 或隧道。只允许将物理以太网数据链路、聚合数据链路和以太网桩模块分配给网桥。

分配给网桥的所有链路必须具有相同的 MTU。当分配链路时，将对此进行检查。如果链路不是网桥中的第一个链路，并且具有不同的 MTU，则会以停用形式将此链路添加到网桥中。

请注意，使用网桥的系统不能将 `eeprom(1M) local-mac-address?` 变量设置为 `false`。

这些选项与 `create-bridge` 子命令的选项相同。

```
dladm remove-bridge [-R root-dir] -l link [-l link...] bridge-name
```

从网桥实例中删除一个或多个链路。如果指定了多个链路，并且删除其中的任何一个链路会导致错误，该命令将失败，并且不会删除任何链路。

从网桥删除链路需要 `PRIV_SYS_DL_CONFIG` 特权。

这些选项与 `create-bridge` 子命令的选项相同。

```
dladm show-bridge [-flt] [-s [-i interval]] [[-p] -o field,...] [bridge-name]
```

显示网桥的运行状态和配置、其附加链路、识别的转发条目以及 TRILL 昵称数据库。显示整个网桥状态和配置时，可以省略网桥名称以显示所有网桥。其他形式需要指定的网桥。

`show-bridge` 子命令接受下列选项：

`-i interval`、`--interval=interval`

与 `-s` 选项配合使用，用于指定应显示统计信息的时间间隔（以秒为单位）。如果未指定此选项，则仅显示一次统计信息。

`-s`、`--statistics`

显示指定网桥或给定网桥的附加链路的统计信息。此选项不能与 `-f` 选项以及 `-t` 选项一起使用。

`-p`、`--parseable`

使用稳定的可供计算机解析的格式显示。请参见下面的“可解析的输出格式”。

`-o field[,...]、--output=field[,...]`

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。下面介绍的是字段名称。特殊值 `all` 将显示所有字段。如果未指定 `-o`，每组字段都有各自要显示的缺省集。

缺省情况下，`show-bridge` 子命令显示网桥配置。可显示下列字段：

BRIDGE

网桥的名称。

ADDRESS

网桥唯一标识符值（MAC 地址）。

PRIORITY

配置的优先级值；使用 `-p` 以及 `create-bridge` 和 `modify-bridge` 进行设置。

BMAXAGE

配置的网桥最长使用寿命；使用 `-m` 以及 `create-bridge` 和 `modify-bridge` 进行设置。

BHELLOTIME

配置的网桥招呼时间；使用 `-h` 以及 `create-bridge` 和 `modify-bridge` 进行设置。

BFWDDELAY

配置的转发延迟；使用 `-d` 以及 `create-bridge` 和 `modify-bridge` 进行设置。

FORCEPROTO

配置的强制最大协议；使用 `-f` 以及 `create-bridge` 和 `modify-bridge` 进行设置。

TCTIME

自上次更改拓扑以来所经过的时间（以秒为单位）。

TCCOUNT

拓扑更改的计数。

TCHANGE

此字段表示已检测到拓扑更改。

DESRROOT

根节点的网桥标识符。

ROOTCOST

到根节点的路径的成本。

ROOTPORT

用于访问根节点的端口号。

MAXAGE

根节点的最长使用寿命值。

HELLOTIME

根节点的招呼时间值。

FWDELAY

根节点的转发延迟值。

HOLDTIME

最长 BPDU 时间间隔。

缺省情况下，如果未指定 **-o** 选项，则仅显示 **BRIDGE**、**ADDRESS**、**PRIORITY** 和 **DESR00T** 字段。

如果指定了 **-s** 选项，**show-bridge** 子命令将显示网桥统计信息。可显示下列字段：

BRIDGE

网桥名称。

DROPS

因资源问题而删除的包的数目。

FORWARDS

从一个链路转发到另一个链路的包的数目。

MBCAST

网桥处理的多播和广播包的数目。

RECV

在所有附加链路上收到的包的数目。

SENT

在所有附加链路上发送的包的数目。

UNKNOWN

处理的具有未知目标的包的数目。此类包将发送到所有链路。

缺省情况下，如果未指定 **-o** 选项，则仅显示 **BRIDGE**、**DROPS** 和 **FORWARDS** 字段。

show-bridge 子命令还接受下列选项：

-l、--link

显示附加到某个网桥实例的所有链路的链路相关状态和统计信息。通过使用此选项而不使用 **-s** 选项，将显示每个链路的下列字段：

LINK

链路名称。

INDEX

网桥上的端口（链路）索引编号。

STATE

链路的状态。此状态可以为 **disabled**、**discarding**、**learning**、**forwarding**、**non-stp** 或 **bad-mtu**。

UPTIME

自上次重置或初始化以来所经过的秒数。

OPERCOST

实际使用成本 (1-65535)。

OPERP2P

此选项指示是否已检测到点对点 (P2P) 模式。

OPEREDGE

此选项指示是否已检测到边缘模式。

DESR00T

此端口上显示的根网桥标识符。

DESCOST

通过指定端口到达网络根节点的路径成本。

DESBRIIDGE

此端口的网桥标识符。

DESPORT

端口的 ID 和优先级，用于传送该端口的配置消息。

TCACK

此选项指示是否显示拓扑更改确认。

如果指定了 **-l** 选项，而未指定 **-o** 选项，则仅显示 **LINK**、**STATE**、**UPTIME** 和 **DESR00T** 字段。

如果指定了 **-l** 选项，**-s** 选项可用于显示各链路的下列字段：

LINK

链路名称。

CFGBPDU

收到的配置 BPDU 的数目。

TCNBPDU

收到的拓扑更改 BPDU 的数目。

RSTPBPDU

收到的快速跨树 BPDU 的数目。

TXBPDU

传送的 BPDU 的数目。

DROPS

因资源问题而删除的包的数目。

RECV

网桥接收的包的数目。

XMIT

网桥发送的包的数目。

如果未指定 `-o` 选项，则仅显示 `LINK`、`DROPS`、`RECV` 和 `XMIT` 字段。

`-f`、`--forwarding`

显示某个网桥实例的转发条目。使用此选项，可针对每个转发条目显示下列字段：

`DEST`

目标 MAC 地址。

`AGE`

条目使用寿命（以秒和毫秒为单位）。对于本地条目，则省略此字段。

`FLAGS`

如果 MAC 地址属于附加链路或某个附加链路上的 VNIC，则会显示 `L`（本地）标志。

`OUTPUT`

对于本地条目，这是具有 MAC 地址的附加链路的名称。否则，对于使用跨树协议的网桥，此字段即为输出接口名称。对于 `RBridge`，此字段为输出 `TRILL` 昵称。

如果未指定 `-o` 选项，则会显示 `DEST`、`AGE`、`FLAGS` 和 `OUTPUT` 字段。

`-t`、`--trill`

显示某个网桥实例的 `TRILL` 昵称条目。使用此选项，可针对每个 `TRILL` 昵称条目显示下列字段：

`NICK`

此 `RBridge` 的 `TRILL` 昵称，该昵称是介于 1 至 65535 之间的某个数字。

`FLAGS`

如果昵称标识本地系统，则会显示 `L` 标志。

`LINK`

将消息发送到此 `RBridge` 时要输出的链路名称。

`NEXTHOP`

用于访问具有此昵称的 `RBridge` 的下一中继 `RBridge` 的 MAC 地址。

如果未指定 `-o` 选项，则会显示 `NICK`、`FLAGS`、`LINK` 和 `NEXTHOP` 字段。

`dladm create-vlan [-ft] [-R root-dir] -l ether-link -v vid [vlan-link]`

使用 ID `vid` 通过以太网链路 `ether-link` 创建带有标记的 VLAN 链路。此 VLAN 链路的名称可指定为 `vlan-link`。如果未指定此名称，则会自动生成名称（假定 `ether-link` 为 `namePPA`）：

`<name><1000 * vlan-tag + PPA>`

例如，如果 `ether-link` 为 `bge1`，并且 `vid` 为 2，则生成的名称为 `bge2001`。

-f、**--force**

强制创建 VLAN 链路。某些设备允许的帧不过大，以致无法包括 VLAN 头。当通过此类设备创建 VLAN 链路时，需要使用 **-f** 选项，并且生成的 VLAN 上的 IP 接口的 MTU 必须设置为 1496，而不是 1500。

-l ether-link

指定创建 VLAN 所基于的以太网链路。

-t、**--temporary**

指定 VLAN 链路为临时链路。临时 VLAN 链路将持续到下一次重新引导。

-R root-dir、**--root-dir=root-dir**

请参见上面的“选项”。

dladm delete-vlan [-t] [-R root-dir] vlan-link

删除指定的 VLAN 链路。

delete-vlan 子命令接受下列选项：

-t、**--temporary**

指定删除为临时删除。临时删除将持续到下一次重新引导。

-R root-dir、**--root-dir=root-dir**

请参见上面的“选项”。

dladm show-vlan [-PZ] [[-p] -o field[,...]] [-z zone[,...]] [vlan-link]

显示所有 VLAN 链路或指定 VLAN 链路的 VLAN 配置。

show-vlan 子命令接受下列选项：

-o field[,...]、**--output=field[,...]**

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名称必须是下面列出的某个字段，或者是可以显示所有字段的特殊值 **all**。对于每个 VLAN 链路，可显示下列字段：

LINK

VLAN 链路的名称。

VID

与 VLAN 关联的 ID。

OVER

配置此 VLAN 所基于的物理链路的名称。

FLAGS

一组与 VLAN 链路相关联的标志。可能的标志包括：

f

VLAN 是使用 **create-vlan** 的 **-f** 选项创建的。

i

VLAN 是在打开 DLPI 链路时隐式创建的。在最终关闭 DLPI 链路时（例如，当未检测到与 VLAN 链路关联的 IP 接口时），这些 VLAN 链路将自动删除。

将来可定义其他标志。

-p、**--parseable**

使用稳定的可供计算机解析的格式显示。**-o** 选项需要与 **-p** 一起使用。请参见下面的“可解析的输出格式”。

-P、**--persistent**

显示持久性 VLAN 配置（而非正在运行的系统的状态）。

-Z

在输出中显示 ZONE 列。

-z zone[,...]

请参见上面 `dladm show-link` 下 **-z** 选项的描述。

`dladm scan-wifi [[-p] -o field[,...]] [wifi-link]`

在所有 WiFi 链路上或仅在指定 *wifi-link* 上扫描 WiFi 网络。

缺省情况下，当前将显示除 BSSTYPE 以外的所有字段。

-o field[,...]、**--output=field[,...]**

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名称必须是下面列出的某个字段，或者是可以显示所有字段的特殊值 **all**。对于找到的每个 WiFi 网络，可显示下列字段：

LINK

WiFi 网络所在的链路的名称。

ESSID

WiFi 网络的 ESSID（名称）。

BSSID

此字段的值为 WiFi 网络访问点的硬件地址（适用于 BSS 网络）或 WiFi 网络随机生成的唯一标记（适用于 IBSS 网络）。

SEC

此字段的值为 **none**（适用于未采用任何安全机制的 WiFi 网络）、**wep**（适用于需要 WEP（Wired Equivalent Privacy，有线对等保密）的 WiFi 网络）或 **wpa**（适用于需要 WPA（Wi-Fi Protected Access，Wi-Fi 保护访问）的 WiFi 网络）。

MODE

支持的连接模式：**a**、**b**、**g** 或 **n** 中的一个或多个。

STRENGTH

信号强度：**excellent**、**very good**、**good**、**weak** 或 **very weak** 之一。

SPEED

WiFi 网络的最大速度（以兆位/秒为单位）。

BSSTYPE

此字段的值为 **bss**（适用于 BSS（基础结构）网络）或 **ibss**（适用于 IBSS（即席）网络）。

-p、--parseable

使用稳定的可供计算机解析的格式显示。**-o** 选项需要与 **-p** 一起使用。请参见下面的“可解析的输出格式”。

```
dladm connect-wifi [-e essid] [-i bssid] [-k key,...] [-s none|wep|wpa] [-a open|shared]
[-b bss|ibss] [-c] [-m a|b|g|n] [-T time] [wifi-link]
```

连接到 WiFi 网络。这包括四个步骤：**搜索、过滤、区分优先级和关联**。但是，要启用到非广播 WiFi 网络的连接并提高性能，如果使用 **-e** 或 **-i** 选项指定 BSSID 或 ESSID，则会跳过前三个步骤，并且 **connect-wifi** 将立即尝试与符合提供的其余参数的 BSSID 或 ESSID 关联。如果此关联失败，但可能存在与指定条件相符的其他网络，则会按照如下所述开始传统的搜索过程。

搜索步骤可用于查找指定 WiFi 链路上的所有可用 WiFi 网络，这些网络必须处于未连接状态。为了便于管理，如果系统中只有一个 WiFi 链路，则可以省略 *wifi-link*。

搜索完成后，将根据下列选项的值过滤网络列表：

-e *essid*、--essid=*essid*

系统将过滤出没有相同 *essid* 的网络。

-b *bss|ibss*、--bsstype=*bss|ibss*

系统将过滤出没有相同 *bsstype* 的网络。

-m *a|b|g*、--mode=*a|b|g|n*

系统将过滤出不适用于指定 802.11 模式的网络。

-k *key*,...、--key=*key*,...

使用按键名指定的 *secobj* 连接到网络。系统将过滤出不适用于指定键的网络。

-s none|wep|wpa、--sec=none|wep|wpa

系统将过滤出不适用于指定安全模式的网络。

接着，系统先后根据信号强度和最大速度区分剩余网络的优先级。最后，依次尝试与列表中的每个网络关联，直到某个关联成功或没有其他剩余网络为止。

除上面介绍的选项以外，下列选项也可以用于控制 **connect-wifi** 的行为：

-a open|shared、--auth=open|shared

使用指定验证模式进行连接。缺省情况下，将依次尝试执行 **open** 和 **shared**。

-c、--create-ibss

如果找不到与指定 ESSID 匹配的网络，则配合使用此选项和 **-b ibss** 以创建新的即席网络。如果未指定 ESSID，**-c -b ibss** 将始终触发系统创建新的即席网络。

-T *time* \ --timeout=*time*

指定关联成功需要等待的秒数。如果 *time* 为 **forever**，关联将需要无限等待。当前缺省值为 10 秒，但可能会在以后进行更改。超时短于缺省值时，可能无法确保成功。

-k *key*,... \ --key=*key*,...

除上面介绍的过滤操作以外，还将使用指定键来确保关联的安全。使用的安全模式将基于键类；如果明确指定了安全模式，该模式必须与键类兼容。所有键必须属于同一类。

对于支持多个键插槽的安全模式，放置键的插槽将由后跟有索引的冒号指定。因此，**-k mykey:3** 会将 *mykey* 放置到插槽 3 中。缺省情况下，将采用插槽 1。对于支持多个键的安全模式，可以指定以逗号分隔的列表，并将第一个键作为活动键。

dladm disconnect-wifi [-a] [*wifi-link*]

断开与一个或多个 WiFi 网络的连接。如果 *wifi-link* 指定了已连接的 WiFi 链路，则会断开与此链路的连接。为了便于管理，如果仅连接一个 WiFi 链路，则可以省略 *wifi-link*。

-a \ --all-links

断开与连接的所有链路的连接。这主要供脚本使用。

dladm show-wifi [-Z] [[-p] -o *field*,...] [-z *zone*[,...]] [*wifi-link*]

显示所有 WiFi 链路或指定链路 *wifi-link* 的 WiFi 配置信息。

-o *field*,... \ --output=*field*

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名称必须是下面列出的某个字段，或者是可以显示所有字段的特殊值 **all**。对于每个 WiFi 链路，可显示下列字段：

LINK

显示的链路的名称。

STATUS

如果链路已连接，则为 **connected**，如果未连接，则为 **disconnected**。如果链路已断开连接，所有其余字段的值均为 **--**。

ESSID

连接的 WiFi 网络的 ESSID（名称）。

BSSID

此字段的值为 WiFi 网络访问点的硬件地址（适用于 BSS 网络）或 WiFi 网络随机生成的唯一标记（适用于 IBSS 网络）。

SEC

此字段的值为 **none**（适用于未采用任何安全机制的 WiFi 网络）、**wep**（适用于需要 WEP 的 WiFi 网络）或 **wpa**（适用于需要 WPA 的 WiFi 网络）。

MODE

支持的连接模式：`a`、`b`、`g` 或 `n` 中的一个或多个。

STRENGTH

连接强度：`excellent`、`very good`、`good`、`weak` 或 `very weak` 之一。

SPEED

连接速度（以兆位/秒为单位）。

AUTH

此字段的值为 `open` 或 `shared`（请参见 `connect-wifi`）。

BSSTYPE

此字段的值为 `bss`（适用于 BSS（基础结构）网络）或 `ibss`（适用于 IBSS（即席）网络）。

缺省情况下，当前将显示除 `AUTH`、`BSSID` 和 `BSSTYPE` 以外的所有字段。

-p, --parseable

使用稳定的可供计算机解析的格式显示。`-o` 选项需要与 `-p` 一起使用。请参见下面的“可解析的输出格式”。

-Z

在输出中显示 `ZONE` 列。

-z zone[,...]

请参见上面 `dladm show-link` 下 `-z` 选项的描述。

dladm show-ether [-xZ] [[-p] -o *field*,...] [-z *zone*[,...]] [*ether-link*]

显示所有物理以太网链路或指定物理以太网链路的状态信息。

`show-ether` 子命令接受下列选项：

-o *field*,... , --output=*field*

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名称必须是下面列出的某个字段，或者是可以显示所有字段的特殊值 `all`。对于每个链路，可显示下列字段：

LINK

显示的链路的名称。

PTYPE

参数类型，其中 `current` 表示链路的协商状态，`capable` 表示受设备支持的功能，`adv` 表示声明的功能，`peeradv` 表示链路合作伙伴声明的功能。

STATE

链路的状态。

AUTO

值为 `yes/no`，指示是否声明自动协商。

SPEED-DUPLEX

可用的速度和双工值的组合。速度单位使用尾随后缀 **G** (GB/s) 或 **M** (Mb/s) 进行编码。双工值以 **f** (全双工) 或 **h** (半双工) 的形式进行编码。

PAUSE

流量控制信息。可以为 **no** (表示流量控制不可用)、**tx** (表示端点可传送暂停帧, 但忽略收到的任何暂停帧)、**rx** (表示端点接收暂停帧, 并根据收到的暂停帧进行操作) 或 **bi** (表示双向流量控制)。

REM_FAULT

故障检测信息。有效值为 **none** 或 **fault**。

缺省情况下, 将针对“current”PTYPE 显示除 **REM_FAULT** 以外的所有字段。

-p, --parseable

使用稳定的可供计算机解析的格式显示。**-o** 选项需要与 **-p** 一起使用。请参见下面的“可解析的输出格式”。

-x, --extended

针对 PTYPE 值 **current**、**capable**、**adv** 和 **peeradv** 显示扩展输出。

-Z

在输出中显示 **ZONE** 列。

-z zone[,...]

请参见上面 **dladm show-link** 下 **-z** 选项的描述。

dladm set-linkprop [-t] [-R root-dir] -p prop=value[,...] link

在指定链路上设置一个或多个属性的值。属性及其可能的值的列表取决于链路类型、网络设备驱动程序和联网硬件。这些属性可使用 **show-linkprop** 进行检索。

-t, --temporary

指定更改为临时更改。临时更改将持续到下一次重新引导。

-R root-dir, --root-dir=root-dir

请参见上面的“选项”。

-p prop=value[,...]、--prop prop=value[,...]

要设置为指定值的属性的逗号分隔列表。

请注意, 如果设置了持久性值, 临时值将更改为相同值。

dladm reset-linkprop [-t] [-R root-dir] [-p prop,...] link

在指定链路上将一个或多个属性重置为其相应值。这些属性将重置为启动时所设置的值。如果未指定属性, 则重置所有属性。有关属性说明, 请参见 **show-linkprop**。

-t, --temporary

指定重置为临时重置。将值重置为缺省值。临时重置将持续到下一次重新引导。

-R *root-dir* \ --root-dir=*root-dir*

请参见上面的“选项”。

-p *prop*, ... \ --prop=*prop*, ...

要重置的属性的逗号分隔列表。

请注意，如果重置了持久性值，临时值将更改为相同值。

dladm show-linkprop [-PZ] [[-c] -o *field*[,...]] [-p *prop*[,...]] [-z *zone*[,...]] [*link*]

显示所有数据链路或指定链路的一个或多个属性的当前值或持久性值。缺省情况下，将显示当前值。如果未指定任何属性，则显示所有可用的链路属性。对于每个属性，将显示下列字段：

-o *field*[,...] \ --output=*field*

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名称必须是下面列出的某个字段，或者是可以显示所有字段的特殊值 **all**。对于每个链路，可显示下列字段：

LINK

数据链路的名称。

PROPERTY (属性)

属性的名称。

PERM

属性的读/写权限。显示的值为 **ro** 或 **rw** 之一。

VALUE (值)

当前（或持久性）属性值。如果未设置此值，此值将显示为 **--**。在未知情况下，此值将显示为 **?**。尚未设置或尚未重置的持久性值将显示为 **--**，并且将使用系统的 **DEFAULT** 值（如果有）。

DEFAULT (缺省)

属性的缺省值。如果属性没有缺省值，则显示 **--**。

POSSIBLE (可能)

属性可具有的值的逗号分隔列表。如果这些值在某个数字范围中，则可以显示简写形式 *min - max*。如果可能的值未知或无限制，则显示 **--**。

属性列表取决于链路类型和网络设备驱动程序，其他给定属性的可用值取决于底层网络硬件及其状态。“常规链路属性”部分中记录了一些常规链路属性。但是，以“_”（下条）开头的链路属性特定于给定链路或其底层网络设备，并可进行更改或删除操作。有关详细信息，请参见相应的网络设备驱动程序手册页。

-c \ --parseable

使用稳定的可供计算机解析的格式显示。**-o** 选项需要与此选项一起使用。请参见下面的“可解析的输出格式”。

-P \ --persistent

显示持久性链路属性信息

`-p prop, ... \ --prop=prop, ...`

要显示的属性的逗号分隔列表。请参见子命令说明后面的链路属性部分。

`-Z`

在输出中显示 ZONE 列。

`-z zone[,...]`

请参见上面 `dladm show-link` 下 `-z` 选项的描述。

dladm create-secobj `[-t] [-R root-dir] [-f file] -c class secobj`

在指定 *class* 中创建一个名为 *secobj* 的安全对象，以便在稍后连接到加密网络时用作 WEP 或 WPA 密钥。安全对象的值可以通过交互方式提供，也可以从文件中读取。交互式提示的序列及文件格式取决于安全对象的类。

目前支持 *wep* 和 *wpa* 类。WEP（Wired Equivalent Privacy，有线对等保密）密钥的长度可以为 5 或 13 个字节。它可以作为 ASCII 或十六进制字符串提供，因此，12345 和 0x3132333435 是等效的 5 字节密钥（可以省略 0x 前缀）。含有 WEP 密钥的文件必须包括一个使用 WEP 密钥格式的行。WPA（Wi-Fi Protected Access，Wi-Fi 保护访问）密钥必须作为 ASCII 字符串提供，并且其长度介于 8 至 63 个字节之间。

此子命令只能由“网络链路安全”RBAC 配置文件包含的用户或角色使用。

`-c class \ --class=class`

class 可以为 *wep* 或 *wpa*。请参见前面的讨论。

`-t \ --temporary`

指定创建为临时创建。临时创建将持续到下一次重新引导。

`-R root-dir \ --root-dir=root-dir`

请参见上面的“选项”。

`-f file \ --file=file`

指定应当用于获取安全对象的值的文件。此文件的格式取决于安全对象类。有关使用此选项设置 WEP 密钥的示例，请参见“示例”部分。

dladm delete-secobj `[-t] [-R root-dir] secobj[,...]`

删除一个或多个指定的安全对象。此子命令只能由“网络链路安全”RBAC 配置文件包含的用户或角色使用。

`-t \ --temporary`

指定删除为临时删除。临时删除将持续到下一次重新引导。

`-R root-dir \ --root-dir=root-dir`

请参见上面的“选项”。

dladm show-secobj `[-P] [[-p] -o field[,...]] [secobj,...]`

显示当前或持久性安全对象信息。如果指定了一个或多个安全对象，则会显示每个安全对象的信息，否则显示所有当前或持久性安全对象。

缺省情况下将显示当前安全对象，所有这些安全对象都是以持久方式创建的不会临时删除的对象，或者是临时创建的对象。

出于安全原因，无法显示安全对象的值。

-o *field*[...]、**--output=*field*[...]**

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名称必须为下面列出的字段之一。对于显示的安全对象，可以显示下列字段：

OBJECT

安全对象的名称。

CLASS

安全对象的类。

-p、**--parseable**

使用稳定的可供计算机解析的格式显示。**-o** 选项需要与 **-p** 一起使用。请参见下面的“可解析的输出格式”。

-P、**--persistent**

显示持久性安全对象信息

```
dladm create-vnic [-t] -l link [-R root-dir] [-m value | auto | {factory [-n slot-identifier]} |
{vrrp -A {inet | inet6} -V vrid} | {random [-r prefix]}] [-v vlan-id] [-p prop=value [...]]
vnic-link
```

在指定链路上创建名为 *vnic-link* 的 VNIC。

-t、**--temporary**

指定 VNIC 为临时 VNIC。临时 VNIC 将持续到下一次重新引导。

-R *root-dir*、**--root-dir=*root-dir***

请参见上面的“选项”。

-l *link*、**--link=*link***

link 可以为物理链路或 etherstub。

-m *value* | *keyword*、**--mac-address=*value* | *keyword***

根据指定值或关键字设置 VNIC 的 MAC 地址。如果 *value* 不是关键字，则将其解释为单点传送 MAC 地址，该地址必须对底层 NIC 有效。用户指定的 MAC 地址必须位于由“全局唯一”和“本地管理”类型的 MAC 地址指定的范围内。

可使用下列特殊关键字：

factory [-n *slot-identifier*]、

factory [--slot=*slot-identifier*]

为 VNIC 分配出厂 MAC 地址。当请求出厂 MAC 地址时，可以配合使用 **-m** 和 **-n** 选项，以便指定要使用的 MAC 地址插槽。如果未指定 **-n**，系统将选择一个可用的出厂 MAC 地址。**show-phys** 子命令的 **-m** 选项可用于显示出厂 MAC 地址及其插槽标识符和可用性的列表。

random [-r *prefix*]、

random [--mac-prefix=*prefix*]

为 VNIC 分配随机 MAC 地址。将使用缺省前缀，该前缀包含有效 IEEE OUI，并设置了本地位。该前缀可使用 **-r** 选项进行覆盖。

`vrrp -A {inet | inet6} -V vrid`

根据指定的地址族和 *vrid* 为 VNIC 分配 VRRP 虚拟 MAC 地址。

`auto`

首先尝试并使用出厂 MAC 地址。如果该地址不可用，则分配随机 MAC 地址。如果未指定 `-m` 选项，则缺省操作为 `auto`。

`-v vlan-id`

针对此 VNIC 启用 VLAN 标记。VLAN 标记的 ID 将为 *vlan-id*。

`-p prop=value,... \ --prop prop=value,...`

要设置为指定值的属性的逗号分隔列表。

`dladm delete-vnic [-t] [-R root-dir] vnic-link`

删除指定 VNIC。

`-t \ --temporary`

指定删除为临时删除。临时删除将持续到下一次重新引导。

`-R root-dir \ --root-dir=root-dir`

请参见上面的“选项”。

`dladm show-vnic [-pPZ] [-s [-i interval]] [-o field[,...]] [-l link] [-z zone[,...]] [vnic-link]`

显示所有 VNIC 或某个链路上的所有 VNIC 的 VNIC 配置信息，或者仅显示指定 *vnic-link* 的 VNIC 配置信息。

`-o field[,...] \ --output=field[,...]`

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名称必须是下面列出的某个字段，或者是可以显示所有字段的特殊值 `all`。缺省情况下（无 `-o`），`show-vnic` 显示所有字段。

LINK

VNIC 的名称。

OVER

配置此 VNIC 所基于的物理链路的名称。

SPEED

VNIC 的最大速度（以兆位/秒为单位）。

MACADDRESS

VNIC 的 MAC 地址。

MACADDRTYPE

VNIC 的 MAC 地址类型。`dladm` 可区分下列 MAC 地址类型：

`random`（随机）

为 VNIC 分配的随机地址。

`factory`（出厂）

VNIC 使用的出厂 MAC 地址。

-p、**--parseable**

使用稳定的可供计算机解析的格式显示。**-o** 选项需要与 **-p** 一起使用。请参见下面的“可解析的输出格式”。

-P、**--persistent**

显示持久性 VNIC 配置。

-s、**--statistics**

显示 VNIC 统计信息。在 [dlstat\(1M\)](#) 中，此选项已过时。

-i interval、**--interval=interval**

与 **-s** 选项配合使用，用于指定应显示统计信息的时间间隔（以秒为单位）。在 [dlstat\(1M\)](#) 中，此选项已过时。

-l link、**--link=link**

显示指定链路上的所有 VNIC 的信息。

-Z

在输出中显示 ZONE 列。

-z zone[,...]

请参见上面 `dladm show-link` 下 **-z** 选项的描述。

dladm create-part [-t] [-f] [-R root-dir] -l ib-link [-p prop=value[,...]] -P pkey part-link
在指定链路上创建名为 *part-link* 的 IP-over-IB 链路。仅在 InfiniBand 物理链路上支持此子命令。

-f、**--force**

强制创建分区链路，即使端口上缺少 *pkey*、缺少多播组或此端口已关闭。

-l ib-link、**--link=ib-link**

IP-over-IB 物理链路名称。

-P、**--pkey=pkey**

用于创建分区链路的分区密钥。指定的 *pkey* 将始终视为十六进制密钥，而不管它是否具有 `0x` 前缀。

-p prop=value[,...]

--prop prop=value[,...]

要设置为指定值的属性的逗号分隔列表。下面的“常规链路属性”部分中提供了支持的属性。

-R root-dir、**--root-dir=root-dir**

请参见上面的“选项”。

-t、**--temporary**

指定分区链路创建为临时分区链路创建。临时分区链路将持续到下一次重新引导。

dladm delete-part [-R root-dir] part-link

删除指定的分区链路。

`-R root-dir`、`--root-dir=root-dir`

请参见上面的“选项”。

`-t`、`--temporary`

指定分区链路删除为临时分区链路删除。临时删除将持续到下一次重新引导。

`dladm show-part [-pP] [-l ib-link] [-o field,...] [part-link]`

显示所有分区链路或 *ib-link* 上的所有分区的 IB 分区链路信息，或者仅显示指定 *part-link* 的 IB 分区链路信息。

`-l ib-link`、`--link=ib-link`

显示指定链路上的所有分区的信息。

`-o field,...]`

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名称必须是下面列出的某个字段，或者是可以显示所有字段的特殊值 `all`。缺省情况下（无 `-o`），`show-part` 显示所有字段。

LINK

分区链路的名称。

PKEY

与分区链路关联的 `Pkey`。

OVER

创建此分区链路所基于的物理链路的名称。

STATE

分区链路的当前状态。可能的值包括：`up`、`down` 或 `unknown`。

FLAGS

用于创建分区链路的一组状态标志。可能的值为：

`f`

已强制创建分区（未检查是否可以创建分区）。

`t`

分区链路为临时链路，仅持续到下一次重新引导。

`-P`、`--persistent`

显示持久性 IB 分区链路配置。

`-p`、`--parseable`

使用稳定的可供计算机解析的格式显示。`-o` 选项需要与 `-p` 一起使用。请参见下面的“可解析的输出格式”。

`dladm show-ib [-pP] [-o field,...] [ib-link]`

显示所有或指定 IB 链路上的 IB 物理链路信息。

-o field[,...]

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名称必须是下面列出的某个字段，或者是可以显示所有字段的特殊值 **all**。缺省情况下（无 **-o**），**show-ib** 显示所有字段。

LINK

物理链路的名称。

HCAGUID

HCA 的全局唯一标识符。

PORTGUID

端口的全局唯一标识符。

PORT

端口号。

STATE

物理链路的当前状态。可能的值包括：**up**、**down** 或 **unknown**。

PKEYS

与 **LINK** 字段中指定的 **IP-over-IB** 链路相关联的端口上的可用 **Pkey**。

-P、**--persistent**

显示持久性 IB 物理链路配置。

-p、**--parseable**

使用稳定的可供计算机解析的格式显示。**-o** 选项需要与 **-p** 一起使用。请参见下面的“可解析的输出格式”。

dladm create-etherstub [-t] [-R root-dir] etherstub

创建具有指定名称的 **etherstub**。

-t、**--temporary**

指定 **etherstub** 为临时 **etherstub**。临时 **etherstub** 不会在重新引导后保留。

-R root-dir、**--root-dir=root-dir**

请参见上面的“选项”。

可基于 **etherstub**（而非物理 NIC）创建 **VNIC**。与物理 NIC 相同，此类创建会导致栈在基于同一 **etherstub** 创建的不同 **VNIC** 之间隐式创建虚拟交换机。

dladm delete-etherstub [-t] [-R root-dir] etherstub

删除指定的 **etherstub**。

-t、**--temporary**

指定删除为临时删除。临时删除将持续到下一次重新引导。

-R root-dir、**--root-dir=root-dir**

请参见上面的“选项”。

```
dladm show-etherstub [-Z] [-z zone[,...]] [etherstub]
```

缺省情况下显示配置的所有 etherstub，或者如果指定了 *etherstub*，则显示指定 etherstub。

-Z

在输出中显示 ZONE 列。

-z zone[,...]

请参见上面 `dladm show-link` 下 -z 选项的描述。

```
dladm create-iptun [-t] [-R root-dir] -T type [-a {local|remote}=addr,...] iptun-link
```

创建名为 *iptun-link* 的 IP 隧道链路。通过使用 [ipsecconf\(1M\)](#)，可以借助 IPsec 为此类链路提供额外的保护。

从概念上讲，一个 IP 隧道包括两个部分：两个或多个 IP 节点之间的虚拟链路，以及位于此链路上方的 IP 接口，该接口允许系统传送和接收由底层链路封装的 IP 包。此子命令创建一个虚拟链路。[ipadm\(1M\)](#) 命令用于在此链路上方配置 IP 接口。

-t、--temporary

指定 IP 隧道链路为临时链路。临时隧道将持续到下一次重新引导。

-R root-dir、--root-dir=root-dir

请参见上面的“选项”。

-T type、--tunnel-type=type

指定要创建的隧道的类型。此类型可以是以下类型之一：

ipv4

位于两个 IPv4 节点之间的点对点 IP-over-IP 隧道。此类型的隧道需要 IPv4 源地址和目标地址才能运行。可以在此类隧道上方检测到 IPv4 和 IPv6 接口，以便创建 IPv4-over-IPv4 和 IPv6-over-IPv4 隧道配置。

ipv6

位于两个 IPv6 节点之间的点对点 IP-over-IP 隧道（按照 IETF RFC 2473 中的定义）。此类型的隧道需要 IPv6 源地址和目标地址才能运行。可以在此类隧道上方检测到 IPv4 和 IPv6 接口，以便创建 IPv4-over-IPv6 和 IPv6-over-IPv6 隧道配置。

6to4

6to4 一点对多点隧道（按照 IETF RFC 3056 中的定义）。此类型的隧道需要 IPv4 源地址才能运行。在此类隧道链路上方检测 IPv6 接口，以便配置 6to4 路由器。

-a {local|remote}=addr,...

--address {local|remote}=addr,...

与本地或远程隧道地址相对应的文本 IP 地址或主机名。可单独指定本地或远程地址，也可以同时指定两个地址，并使用逗号进行分隔（例如，`-a local=laddr,remote=raddr`）。

`dladm modify-iptun [-t] [-R root-dir] -a {local|remote}=addr,... iptun-link`
 修改指定 IP 隧道的参数。

`-t`、`--temporary`

指定修改为临时修改。临时修改将持续到下一次重新引导。

`-R root-dir`、`--root-dir=root-dir`

请参见上面的“选项”。

`-a {local|remote}=addr,...`

`--address {local|remote}=addr,...`

为隧道链路指定新的本地或远程地址。有关说明，请参见 `create-iptun`。

`dladm delete-iptun [-t] [-R root-dir] iptun-link`

删除指定的 IP 隧道链路。

`-t`、`--temporary`

指定删除为临时删除。临时删除将持续到下一次重新引导。

`-R root-dir`、`--root-dir=root-dir`

请参见上面的“选项”。

`dladm show-iptun [-PZ] [[-p] -o field[,...]] [-z zone[,...]] [iptun-link]`

显示一个 IP 隧道或所有 IP 隧道的 IP 隧道链路配置。

`-P`、`--persistent`

显示持久性 IP 隧道配置。

`-p`、`--parseable`

使用稳定的可供计算机解析的格式显示。`-o` 选项需要与 `-p` 一起使用。请参见下面的“可解析的输出格式”。

`-o field[,...]`、`--output=field[,...]`

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名称必须是下面列出的某个字段，或者是可以显示所有字段的特殊值 `all`。缺省情况下（无 `-o`），`show-iptun` 显示所有字段。

LINK

IP 隧道链路的名称。

TYPE

`create-iptun` 的 `-T` 选项指定的隧道类型。

FLAGS

一组与 IP 隧道链路相关联的标志。可能的标志包括：

S

IP 隧道链路受 IPsec 策略保护。要显示与隧道链路关联的 IPsec 策略，请输入：

```
# ipsecconf -ln -i tunnel-link
```

有关如何配置 IPsec 策略的更多详细信息，请参见 `ipsecconf(1M)`。

i

已使用 [ipadm\(1M\)](#) 隐式创建 IP 隧道链路，并且当不再引用此 IP 隧道链路时（也即，当删除此隧道上方的最后一个 IP 接口时），将自动删除此链路。有关隐式创建隧道的详细信息，请参见 [ipadm\(1M\)](#)。

LOCAL

本地隧道地址。

REMOTE

远程隧道地址。

-Z

在输出中显示 ZONE 列。

-z *zone*[,...]

请参见上面 `dladm show-link` 下 `-z` 选项的描述。

`dladm show-usage [-a] -f filename [-p plotfile -F format] [-s time] [-e time] [link]`

在 [dlstat\(1M\)](#) `show-link -h` 命令中，此子命令已过时。

可解析的输出格式

许多 `dladm` 子命令都具有一个以计算机可解析的格式显示输出的选项。输出格式是一行或多行以冒号 (:) 分隔的字段。显示的字段特定于使用的子命令，并且在给定子命令的 `-o` 选项的条目下列出。输出仅包括通过 `-o` 选项请求的字段，这些字段按请求的顺序排列。

当请求多个字段时，将在输出之前使用反斜杠 (\) 对任何文本冒号字符进行转义。类似地，也会对文本反斜杠字符进行转义 (\\)。通过使用带有环境变量 `IFS=:` 的 [shell read\(1\)](#) 函数，可以解析此转义格式（请参见下面的“示例”）。请注意，当仅请求一个字段时，不会进行转义。

常规链路属性

支持下列常规链路属性：

autopush

指定在打开链路的 DLPI 设备时推送到与该链路相关联的流的 STREAMS 模块集。该模块集是以空格分隔的模块的列表。

可选的特殊字符序列 `[anchor]` 表示应将 STREAMS 锚放置到以前在列表中指定的模块的流中。指定多个锚或将锚放置在列表开头会产生错误。

最好使用 `autopush` 属性，而不是较通用的 [autopush\(1M\)](#) 命令。

cpus

将给定数据链路的包处理绑定到某个或某组处理器。此值可以是一个或多个处理器 ID 或 ID 范围的逗号分隔列表。如果此列表包含多个处理器，处理将向外传播到所有处理器。系统将维护处理器连接关联以及任何单独的连接包的排序方式。

系统不会专门为链路保留处理器或处理器集。只有与链路处理相关联的内核线程和中断才会绑定到指定处理器或处理器集。如果需要将处理器专用于链路，可以使用 [psrset\(1M\)](#) 创建处理器集，并指定要绑定到此链路的处理器集中的处理器。

如果已在以前的操作中将链路绑定到处理器或处理器集，则将删除此绑定，并改用新的处理器集。

缺省设置为无 CPU 绑定，也就是说，包处理不会绑定到任何特定处理器或处理器集。

在采用 pool 链路属性的链路上，不允许指定 cpus 属性。

cpus-effective

此只读属性显示指定数据链路上进行包处理用的 CPU 的列表。

如果设置了 cpus 属性，cpus-effective 将相同。

如果设置了 pool 属性，将从管理员指定的池中选择 cpus-effective。

如果未设置 pool 和 cpus 属性，系统将为 cpus-effective 选择适当的值。

rxfanout

允许您指定接收端扇出线程的数目。

接收环上接收的通信流量可以跨多个线程扇出并以并行方式进行处理。这在系统中具有大量 CPU 时特别有用。此属性是对特定数据链路的接收端扇出线程数的计数。请注意，此属性允许管理员指定所需的 rxfanout。但是，根据可用 CPU 和硬件 RX 环的数目，系统可能会选择其他（更小或甚至更大）扇出值。

rxfanout-effective

CPU 的数目是接收端扇出的上界，而 rxrings 的数目是下界。因此，实际的接收端扇出计数值可能不同于用户设置的值。

learn_limit

限制要通过网桥链路获取的新 MAC 源或已更改的 MAC 源的数目。当数目超过此值时，则会临时禁用在此链路上进行获取的功能。只有非 VLAN、非 VNIC 类型的链路才具有此属性。

缺省值为 1000。有效值大于或等于 0。

learn_decay

指定受 learn_limit 限制的源更改的衰减速率。系统每 5 秒钟从网桥链路的计数器中减去此数字。只有非 VLAN、非 VNIC 类型的链路才具有此属性。

缺省值为 200。有效值大于或等于 0。

mac-address

设置数据链路的主 MAC 地址。设置此选项后，更改底层数据链路的所有当前和未来 MAC 客户机使用的主 MAC 地址。

maxbw

设置链路的全双工带宽。此带宽指定为一个带有某个容量级后缀（分别表示 Kbps、Mbps 和 Gbps 的 K、M 或 G）的整数。如果未指定单位，输入值将以 Mbps 为单位进行读取。缺省值是无带宽限制。

pool

将给定数据链路的包处理绑定到通过 **poolcfg(1M)** 和 **pooladm(1M)** 进行定义和管理的处理器池。处理绑定与 **cpus** 链路属性的绑定方式相似，但 CPU 列表不是显式的，并且会通过池功能进行维护。

如果已启用池，并且未为链路指定池，**pool_default** 将用于包处理。

对于 **ip-type=exclusive** 的区域，如果通过池区域属性或 **dedicated-cpus** 分配指定池，该池还将用于与此区域关联的所有数据链路。

在采用 **cpus** 链路属性的链路上，不允许指定 **pool** 属性。

pool-effective

如果已启用池功能，此只读属性将显示用于包处理的池。如果管理员未为数据链路分配池，该池将为 **pool_default**。

如果已禁用池功能，则不存在有效池，并且该值将为空。

priority

设置链路的相对优先级。该值可指定为 **high**、**medium** 或 **low** 标记之一。缺省值为 **high**。此优先级不会反映在线路上的任何协议优先级字段中，而是用于系统内的包处理调度。

rxringsavail

一种只读属性，可用于指定接收端上的可用环数。

rxrings

指定 MAC 客户机的接收环端的数目。值 **sw** 表示不应为此 MAC 客户机分配任何 RX 环且将以软件为基础。值 **hw** 表示此 MAC 客户机可获取一个 RX 环（如果可用），否则将以软件为基础。非零值表示为此 MAC 客户机保留此数目的环（如果可用），否则将失败。如果未指定此属性，MAC 客户机可获取一个 RX 环（如果可用），否则将以软件为基础。

rxhwclntavail

一种只读属性，可用于指定可创建的基于 RX 硬件的额外 MAC 客户机数目。

txringsavail

一种只读属性，可用于指定传送端上的可用环数。

txrings

指定 MAC 客户机的传送环的数目。值 **sw** 表示不应为此 MAC 客户机分配任何 TX 环。值 **hw** 表示此 MAC 客户机可获取一个 TX 环（如果可用），否则将以软件为基础。非零值表示为此 MAC 客户机保留此数目的环（如果可用），否则将失败。如果未指定此属性，MAC 客户机可获取一个 TX 环（如果可用），否则将以软件为基础。

txhwclntavail

一种只读属性，可用于指定可创建的基于 TX 硬件的额外 MAC 客户机数目。

stp

在网桥链路上启用或禁用跨树协议。将此值设置为 **0** 可禁用跨树协议，并将链路置于转发模式，同时启用 BPDU 保护。此模式适用于仅连接到终端节点的点对点链路。只有非 VLAN、非 VNIC 类型的链路才具有此属性。启用 STP 所需的缺省值为 **1**。

forward

启用或禁用 VLAN 的转发。将此值设置为 **0** 可禁用 VLAN 链路的网桥转发。禁用网桥转发会从网桥的“允许集”中删除此 VLAN。为配置的 VLAN 启用网桥转发时，所需的缺省值为 **1**。

default_tag

设置在此链路中收发的不带标记的包所采用的缺省 VLAN ID。只有非 VLAN、非 VNIC 类型的链路才具有此属性。将此值设置为 **0** 可禁用进出端口的不带标记的包的网桥转发。缺省值为 VLAN ID **1**。有效值介于 0 至 4094 之间。缺省 VLAN ID 也称为端口 VLAN 标识符 (PVID)。

不能使用与底层链路的缺省 VLAN 值匹配的 VLAN ID 创建带有标记的 VLAN 或带有 VLAN 标记的 VNIC 链路。链路上所有不带标记的包都已经与缺省 VLAN (PVID) 关联。要成功创建 VLAN ID 等于缺省 VLAN 值的带有标记的 VLAN 或带有 VLAN 标记的 VNIC 链路，必须首先将底层链路的 `default_tag` 属性更改为其他 VLAN 值。

当 `default_tag=0` 时，链路上所有不带标记的包都不再与任何 VLAN 关联。因此，您可以使用介于 1 至 4094 之间的任何 VLAN ID 创建 VLAN 链路。请注意，系统可能会删除端点中使用 PVID 错误标记的所有接收的包。如果给定链路上的所有端点的 PVID 并不完全相符，则会出现此情况。链路上的所有端点必须使用相同的 PVID，并且不能使用 PVID 来标记通信。

stp_priority

设置 STP 和 RSTP 端口优先级值，此值用于确定网桥上的首选根端口。数值越低，优先级越高。缺省值为 **128**。有效值介于 0 至 255 之间。

stp_cost

设置使用链路的 STP 和 RSTP 成本。缺省值为 **auto**，该值根据链路速度设置成本，并使用 **100** 表示 10Mbps，**19** 表示 100Mbps，**4** 表示 1Gbps，**2** 表示 10Gbps。有效值介于 1 至 65535 之间。

stp_edge

启用或禁用网桥边缘端口检测。如果设置为 **0 (false)**，则即使未显示任何类型的网桥 PDU，系统也会假定此端口将连接到其他网桥。缺省值为 **1**，表示系统将自动检测边缘端口。

stp_p2p

设置网桥点对点运行模式。可能的值包括：**true**、**false** 和 **auto**。如果设置为 **auto**，则会自动搜索点对点连接。如果设置为 **true**，则会强制端口模式使用点对点连接。如果设置为 **false**，则会强制端口模式使用普通多点模式。缺省值为 **auto**。

stp_mcheck

触发系统以便在此链路上运行 RSTP 强制 BPDU 迁移检查过程。将此属性值设置为 1 即可触发此过程。此属性自动重置为原来的 0。除非满足下列条件，否则无法设置此值：

- 链路处于桥接状态
- 网桥受跨树协议保护
- 网桥 `force-protocol` 值至少为 2 (RSTP)

缺省值为 0（第 1 部分）。

protection

启用一种或多种链路保护。有效值包括：

mac-nospoof

MAC 地址防欺骗。外发包的源 MAC 地址必须与链路的已配置 MAC 地址匹配。系统将删除不匹配的包。如果链路属于某个区域，启用 `mac-nospoof` 将防止区域所有者修改此链路的 MAC 地址。

ip-nospoof

IP 地址防欺骗。此保护类型与链路属性 `allowed-ips` 配合运行。

`allowed-ips` 是包含 IP（IPv4 或 IPv6）地址的列表。缺省情况下，此列表为空。此列表中隐式存在的地址包括：符合 RFC 2464 的链路本地 IPv6 地址（从链路的 MAC 地址派生）；从 DHCP 回复获取的 IPv4/IPv6 地址；未指定的 IPv4/IPv6 地址（均为零）。

如果外发 IP 包的源地址位于 `allowed-ips` 中，则可以传递此包。

如果外发 ARP 包的发送方协议地址位于 `allowed-ips` 中，则可以传递此包。

将 `allowed-ips` 设置为包含一个或多个 IP 地址的集合以保护数据链路时，将无法对不在该集中的 IP 地址进行配置，此时系统会向用户返回 `EPERM` 错误。此外，接口不能用于转发 IP 包，尝试对此接口设置 `ipadm(1M)` 转发属性将遇到 `EPERM` 错误。

dhcp-nospoof

DHCP 客户机 ID（适用于 DHCPv6 的 DUID）和硬件地址防欺骗。此保护类型与链路属性 `allowed-dhcp-cids` 配合运行。

`allowed-dhcp-cids` 列表中项的格式应与 `/etc/default/dhcpagent` 文件中 `CLIENT ID` 字段的格式相同。唯一区别在于，当指定 DUID 时，应使用 `.`（句点）替换 `,`（逗号）。有关详细信息，请参见 `dhcpagent(1M)`。

仅当满足下列条件时，才能传递外发 DHCP (v4/v6) 包。

- 未配置 `allowed-dhcp-cids`，且包类型为：
 - DHCPv4，客户机 ID 字段必须与配置的 MAC 地址匹配。
 - DHCPv6，DUID 类型必须为 1 或 3，并且 DUID 的链路层地址部分必须与配置的 MAC 地址匹配。

- 已配置 `allowed-dhcp-cids`，且包类型为：
 - DHCPv4，客户机 ID 字段必须与此列表中的某个 ID 或配置的 MAC 地址匹配。
 - DHCPv6，DUID 字段必须与此列表中的某个 ID 匹配，或者 DUID 类型必须为 1 或 3，并且 DUID 的链路层地址部分必须与配置的 MAC 地址匹配。

`restricted`

此保护将外发包类型限定为 IPv4、IPv6 和 ARP。

`zone`

指定链路所属的区域。通过 `dladm` 只能临时修改此属性，因此必须指定 `-t` 选项。要修改区域赋值使之在重新引导后保留，请使用 [zonecfg\(1M\)](#)。可能的值包括当前在系统上运行的任何专用 IP 区域。缺省情况下，按照 [zonecfg\(1M\)](#) 执行区域绑定。

Wifi 链路属性

支持下列 WiFi 链路属性。请注意，将给定属性设置为给定值的功能取决于驱动程序和硬件。

`channel`

指定要使用的通道。当处于 IBSS 模式时，只能通过特定 WiFi 链路修改此属性。缺省值和允许值范围因管理域而异。

`powermode`

指定 WiFi 链路的电源管理模式。可能的值包括 `off`（禁用电源管理）、`max`（最省电模式）和 `fast`（性能敏感型电源管理）。缺省值为 `off`。

`radio`

指定 WiFi 链路的无线电模式。可能的值包括 `on` 或 `off`。缺省值为 `on`。

`speed`

指定 WiFi 链路的固定速度（以兆位/秒为单位）。可能的值集取决于驱动程序和硬件（但可通过 `show-linkprop` 显示）；常用速度包括 1、2、11 和 54。缺省情况下，没有固定速度。

以太网链路属性

仅在只读模式下支持下面的 MII 属性（如 [ieee802.3\(5\)](#) 中所述）。

- `duplex`
- `state`
- `adv_autoneg_cap`
- `adv_10gfdx_cap`
- `adv_1000fdx_cap`
- `adv_1000hdx_cap`
- `adv_100fdx_cap`
- `adv_100hdx_cap`
- `adv_10fdx_cap`
- `adv_10hdx_cap`

各 `adv_` 属性（例如，`adv_10fdx_cap`）还具有一个对应的读/写 `en_` 属性（例如，`en_10fdx_cap`），用于控制在自动协商时使用的参数。在没有电源管理时，可通过 `adv_` 速度/双工参数提供经过协商的、可在当前硬件中使用的值。但是，如果启用电源管理，当前在硬件中公开的速度/双工功能可能是初始链路参数协商中使用的位集的子集。因此，可将 MII `adv_` 参数标记为只读，并将另一组 `en_` 参数用于配置初始协商时的速度和双工属性。

请注意，`adv_autoneg_cap` 没有对应的 `en_autoneg_cap`：`adv_autoneg_cap` 本身是禁用/启用自动协商的 0/1 开关，因此不受电源管理的影响。

此外，系统还会报告下列以太网属性：

`flowctrl`

建立将通过设备宣传的流量控制模式。有效输入为下列项之一：

`auto`

设备上的流量控制模式是动态确定的。要查看设备上设置的实际流量控制模式，请检查 `flowctrl-effective` 链路属性。

`no`

不启用流量控制。

`rx`

接收并处理传入暂停帧。

`tx`

当发生拥塞时，将暂停帧传送到同级设备，但忽略接收的暂停帧。

`pfc`

传送包含应暂停的通信流量的优先级值的暂停帧。接收暂停帧，并对已在帧中指定其优先级值的通信流量执行操作。

`bi`

双向流量控制。

请注意，此值的实际设置受设备和链路合作伙伴允许的功能的限制。

`gvrp-timeout`

指定 VID 声明广播之间的等待时间（以毫秒为单位）。

`flowctrl-effective`

设备上配置的实际流量控制模式。当 `flowctrl` 属性设置为 `auto` 时，这指示已生效的流量控制模式。该属性是只读属性。

`mtu`

设备支持的最大客户机 SDU（Send Data Unit，发送数据单位）。有效范围为 68-65536。

ntcs

设备上支持的通信流量类的数目。支持 DCB（Data Center Bridging，数据中心桥接）扩展的设备可以支持多个通信流量类。此属性可用于确定设备是否支持 DCB 扩展。该属性是只读属性。

pfcmap

此属性用于指示启用 PFC（Priority-based flow control，基于优先级的流量控制）的 802.1p 优先级值。这是一个 8 位掩码，其中的每个位都表示是否为相应的优先级启用 PFC。对于启用 PFC 的优先级，设备将在发生拥塞时为该优先级传送暂停帧。这仅在 `ntcs` 大于零且 `flowctrl-effective` 为 `pfc` 时才相关。

speed

（只读）设备的运行速度 (Mbps)。

tagmode

此链路属性用于控制将在链路上传送的包中插入 802.1Q VLAN 标记的条件。可以为此属性分配两个模式值：

normal 在下列条件下，在外发包中插入 VLAN 标记：

- 此包属于某个 VLAN。
- 用户已请求优先级标记。

vlanonly 仅当外发包属于某个 VLAN 时才插入 VLAN 标记。如果在此模式下插入标记，并且用户还请求了非零优先级，则遵循此优先级并将其包含在 VLAN 标记中。

缺省值为 `vlanonly`。

vlan-announce

此属性控制自动 VLAN ID 声明。启用此属性后，它将广播设备上配置的任何 VNIC 或 VLAN 的 VID。它支持物理链路和聚合。可能的值为：

off

不会发送 VID 声明。

gvrp

使用 GVRP 协议发送的声明（按照 802.1D 中的定义）。请参见 `gvrp-timeout` 以配置广播频率。

InfiniBand 链路属性 仅在 IB 分区对象链路上支持下列属性。

linkmode

在 IB 分区数据链路上设置链路传输服务类型。缺省值为 `cm`。有效值包括：

cm

连接模式。此模式使用缺省 MTU 65520，并且支持的最大 MTU 为 65535 字节。如果远程节点无法使用连接模式，则会自动改用不可靠的数据报模式。

ud

不可靠的数据报模式。此模式使用缺省 MTU 2044，并且支持的最大 MTU 为 4092 字节。

IP 隧道链路属性

支持下列 IP 隧道链路属性。

hoplimit

指定用于封装隧道链路的外部 IP 头的 IPv4 TTL 或 IPv6 跳数限制。所有隧道类型都存在此属性。缺省值为 64（第 1 部分）。

encaplimit

指定 IPv6 隧道的 IPv6 封装限制（按照 RFC 2473 中的定义）。此值是给定隧道包的隧道嵌套限制。缺省值为 4（第 1 部分）。值为 0 表示将禁用封装限制。

示例

示例 1 配置聚合

要使用密钥 1 在设备 bge0（链路名称 net0）和 bge1（链路名称 net1）的聚合上方配置数据链路，请输入以下命令：

```
# dladm create-aggr -d net0 -d net1 1
```

示例 2 连接到 WiFi 链路

要通过一个 WiFi 链路连接到系统上不安全的最佳可用网络（按照为 connect-wifi 指定的优先级规则），请输入以下命令：

```
# dladm connect-wifi
```

示例 3 创建 WiFi 密钥

要以交互方式创建 WEP 密钥 mykey，请输入以下命令：

```
# dladm create-secobj -c wep mykey
```

此外，也可以使用文件内容以非交互方式创建 WEP 密钥 mykey：

```
# umask 077
# cat >/tmp/mykey.$$ <<EOF
12345
EOF
# dladm create-secobj -c wep -f /tmp/mykey.$$ mykey
# rm /tmp/mykey.$$
```

示例 4 连接到指定的加密 WiFi 链路

要使用密钥 mykey 在链路 ath0 上连接到 ESSID wlan，请输入以下命令：

```
# dladm connect-wifi -k mykey -e wlan ath0
```

示例 5 更改链路属性

要在链路 pcwl0 上将 powermode 设置为值 fast，请输入以下命令：

```
# dladm set-linkprop -p powermode=fast pcwl0
```

示例6 连接到受 WPA 保护的 WiFi 链路

创建 WPA 密钥 `psk` 并输入以下命令：

```
# dladm create-secobj -c wpa psk
```

然后，若要在链路 `ath0` 上使用密钥 `psk` 连接到 ESSID `wlan`，请输入以下命令：

```
# dladm connect-wifi -k psk -e wlan ath0
```

示例7 重命名链路

要将 `bge0` 链路重命名为 `mgmt0`，请输入以下命令：

```
# dladm rename-link bge0 mgmt0
```

示例8 更换网卡

假定由于硬件故障，需要将 `bge0` 设备（其链路命名为 `mgmt0`，如上一示例所示）替换为 `ce0` 设备。以物理方式删除 `bge0` NIC，然后将其替换为新的 `ce0` NIC。要将新添加的 `ce0` 设备与 `mgmt0` 配置（以前与 `bge0` 关联）关联，请输入以下命令：

```
# dladm rename-link ce0 mgmt0
```

示例9 删除网卡

假定在上一示例中，目的不是将 `bge0` NIC 更换为其他 NIC，而是删除但不更换此硬件。此种情况下，不应将 `mgmt0` 数据链路配置指定为与不同物理设备相关联（如上一示例所示），而应将其删除。请输入以下命令，以便删除与 `mgmt0` 数据链路相关联的数据链路配置，该数据链路的物理硬件（在本示例中为 `bge0`）已被删除：

```
# dladm delete-phys mgmt0
```

示例10 使用可解析的输出捕获一个字段

下面的赋值将链路 `net0` 的 MTU 保存为名为 `mtu` 的变量。

```
# mtu=$(dladm show-link -p -o mtu net0)
```

示例11 使用可解析的输出在链路上迭代

下面的脚本显示系统中每个链路的状态。

```
# dladm show-link -p -o link,state | while IFS=: read link state; do
    print "Link $link is in state $state"
done
```

示例12 配置 VNIC

通过一个物理链路 `bge0` 创建两个名为 `hello0` 和 `test1` 的 VNIC：

```
# dladm create-vnic -l bge0 hello0
# dladm create-vnic -l bge0 test1
```

示例 13 配置 VNIC 并分配带宽和优先级

通过一个物理链路 `bge0` 创建两个名为 `hello0` 和 `test1` 的 VNIC，并将 `hello0` 作为具有较高优先级的 VNIC，它具有出厂时分配的 MAC 地址，其最大带宽为 50 Mbps。将 `test1` 作为具有较低优先级的 VNIC，它采用随机 MAC 地址，其最大带宽为 100 Mbps。

```
# dladm create-vnic -l bge0 -m factory -p maxbw=50,priority=high hello0
# dladm create-vnic -l bge0 -m random -p maxbw=100M,priority=low test1
```

示例 14 使用出厂 MAC 地址配置 VNIC

首先，列出可用的出厂 MAC 地址，并选择其中之一：??

```
# dladm show-phys -m bge0
```

LINK	SLOT	ADDRESS	INUSE	CLIENT
net0	primary	0:e0:81:27:d4:47	yes	net0
net0	1	8:0:20:fe:4e:a5	no	
net0	2	8:0:20:fe:4e:a6	no	
net0	3	8:0:20:fe:4e:a7	no	

创建一个名为 `hello0` 的 VNIC，并使用插槽 1 的地址：

```
# dladm create-vnic -l bge0 -m factory -n 1 hello0
# dladm show-phys -m bge0
```

LINK	SLOT	ADDRESS	INUSE	CLIENT
net0	primary	0:e0:81:27:d4:47	yes	net0
net0	1	8:0:20:fe:4e:a5	yes	hello0
net0	2	8:0:20:fe:4e:a6	no	
net0	3	8:0:20:fe:4e:a7	no	

示例 15 使用用户指定的 MAC 地址创建 VNIC，并将其绑定到处理器集

使用用户指定的 MAC 地址和处理器绑定 0, 2, 4-6 创建一个名为 `hello0` 的 VNIC。

```
# dladm create-vnic -l bge0 -m 8:0:20:fe:4e:b8 -p cpus=0,2,4-6 hello0
```

示例 16 创建无物理 NIC 的虚拟网络

首先，创建一个名称为 `stub1` 的 etherstub：

```
# dladm create-etherstub stub1
```

在此 etherstub 上创建两个名为 `hello0` 和 `test1` 的 VNIC。此操作会隐式创建一个连接 `hello0` 和 `test1` 的虚拟交换机。

```
# dladm create-vnic -l stub1 hello0
# dladm create-vnic -l stub1 test1
```

示例 17 显示网桥信息

下面的命令使用了不带选项和带有不同选项的 `show-bridge` 子命令。

示例 17 显示网桥信息 (续)

```
# dladm show-bridge
BRIDGE      PROTECT ADDRESS          PRIORITY DESROOT
foo          stp      32768/8:0:20:bf:f 32768     8192/0:d0:0:76:14:38
bar          stp      32768/8:0:20:e5:8 32768     8192/0:d0:0:76:14:38

# dladm show-bridge -l foo
LINK        STATE      UPTIME    DESROOT
hme0        forwarding 117       8192/0:d0:0:76:14:38
qfe1        forwarding 117       8192/0:d0:0:76:14:38

# dladm show-bridge -s foo
BRIDGE      DROPS      FORWARDS
foo          0           302

# dladm show-bridge -ls foo
LINK        DROPS      RECV      XMIT
hme0        0          360832    31797
qfe1        0          322311    356852

# dladm show-bridge -f foo
DEST        AGE      FLAGS    OUTPUT
8:0:20:bc:a7:dc 10.860  --      hme0
8:0:20:bf:f9:69  --      L        hme0
8:0:20:c0:20:26 17.420  --      hme0
8:0:20:e5:86:11  --      L        qfe1
```

示例 18 创建 IPv4 隧道

可通过下面的命令序列在 66.1.2.3 和 192.4.5.6 之间创建并显示一个名为 mytunnel0 的持久性 IPv4 隧道链路：

```
# dladm create-iptun -T ipv4 -a local=66.1.2.3,remote=192.4.5.6 mytunnel0
# dladm show-iptun mytunnel0
LINK        TYPE  FLAGS  SOURCE          DESTINATION
mytunnel0    ipv4  --     66.1.2.3        192.4.5.6
```

然后，可以在此隧道链路上创建点对点 IP 接口：

```
# ipadm create-ip mytunnel0
# ipadm create-addr -T static -a local=10.1.0.1,remote=10.1.0.2 \
mytunnel0/addr
# ipadm show-addr mytunnel0/addr
ADDROBJ     TYPE    STATE    ADDR
mytunnel0/addr static  ok       10.1.0.1->10.1.0.2
```

示例 19 创建 6to4 隧道

下面的命令创建了创建一个 6to4 隧道链路。6to4 路由器的 IPv4 地址为 75.10.11.12。

示例 19 创建 6to4 隧道 (续)

```
# dladm create-iptun -T 6to4 -a local=75.10.11.12 sitetunnel0
# dladm show-iptun sitetunnel0
LINK          TYPE  FLAGS  SOURCE          DESTINATION
sitetunnel0   6to4  --     75.10.11.12     --
```

下面的命令在此隧道上创建了一个 IPv6 接口：

```
# ipadm create-ip sitetunnel0
# ipadm show-addr sitetunnel0/_a
ADDROBJ      TYPE    STATE      ADDR
sitetunnel0/_a  static  ok         2002:4b0a:b0c::1/16
```

请注意，系统在 6to4 接口上自动配置 IPv6 地址。有关如何在 6to4 隧道链路上配置 IPv6 地址的说明，请参见 [ipadm\(1M\)](#)。

示例 20 使用链路保护

启用链路保护：

```
# dladm set-linkprop \
-p protection=mac-nospoof,restricted,ip-nospoof,dhcp-nospoof vnic0
```

禁用链路保护：

```
# dladm reset-linkprop -p protection vnic0
```

修改 allowed-ips 列表：

```
# dladm set-linkprop -p allowed-ips=10.0.0.1,10.0.0.2 vnic0
```

修改 allowed-dhcp-cids 列表：

```
# dladm set-linkprop -p allowed-dhcp-cids=hello vnic0
```

显示生成的配置：

```
# dladm show-linkprop -p protection,allowed-ips vnic0
```

LINK	PROPERTY	PERM	VALUE	DEFAULT	POSSIBLE
vnic0	protection	rw	mac-nospoof, restricted, ip-nospoof, dhcp-nospoof	--	mac-nospoof, restricted, ip-nospoof, dhcp-nospoof
vnic0	allowed-ips	rw	10.0.0.1, 10.0.0.2	--	--
vnic0	allowed-dhcp-cids	rw	hello	--	--

示例 21 创建 IB 分区

下面的命令使用分区密钥 0xffff 在物理链路 ibp0 上创建了一个 ffff.ibp0 分区。

```
# dladm create-part -P ffff -l ibp0 ffff.ibp0
```

示例 22 显示 IB 分区信息

下面的命令显示了 IB 分区信息。

```
# dladm show-part
LINK          PKEY OVER          STATE    FLAGS
ffff.ibp0     FFFF ibp0         up       ----
```

示例 23 显示 IB 数据链路信息

下面的命令显示了 IB 数据链路信息。

```
# dladm show-ib
LINK          HCAGUID          PORTGUID          PORT STATE  PKEYS
ibp0          3BA000100CD7C     3BA000100CD7D     1   down   FFFF
ibp1          3BA000100CD7C     3BA000100CD7E     2   down   FFFF
ibp3          5AD0000033634     5AD0000033636     2   up     FFFF,8001
ibp2          5AD0000033634     5AD0000033635     1   up     FFFF,8001
```

示例 24 删除分区

下面的命令删除了分区 ffff.ibp0。

```
# dladm delete-part ffff.ibp0
```

示例 25 使用 show-link 显示分区信息

下面的命令使用 show-link 子命令显示分区信息。

```
# dladm show-link
LINK          CLASS    MTU    STATE    OVER
e1000g0       phys     1500   up       --
e1000g1       phys     1500   unknown  --
ibp0          phys     65520  down     --
ibp3          phys     65520  up       --
ibp2          phys     65520  up       --
ibp1          phys     65520  down     --
pffff.ibp0    part     2044   down     ibp0
p8001.ibp2    part     65520  unknown  ibp2
```

示例 26 从全局区域显示所有区域中的链路

下面所示的 show-link 命令将从全局区域显示所有区域中的数据链路。不在全局区域中的链路将以区域名前缀后跟斜杠 (/) 分隔符的方式显示。

在本示例中，net0 是在全局区域中创建的 VNIC，zone1/net0 是自动为 zone1 创建的 VNIC，zone2/net0 是自动为 zone2 创建的 VNIC。

示例 26 从全局区域显示所有区域中的链路 (续)

```
# dladm show-link
LINK          CLASS    MTU    STATE    OVER
e1000g0       phys     1500   up       --
e1000g1       phys     8170   unknown  --
e1000g2       phys     1500   unknown  --
e1000g3       phys     1500   unknown  --
net0          vnic     1500   up       e1000g0
zone1/net0    vnic     1500   up       e1000g0
zone2/net0    vnic     1500   up       e1000g0
```

示例 27 显示全局区域中的链路

以下 show-link 命令仅显示全局区域中的数据链路。

```
# dladm show-link -z global
LINK          CLASS    MTU    STATE    OVER
e1000g0       phys     1500   up       --
e1000g1       phys     8170   unknown  --
e1000g2       phys     1500   unknown  --
e1000g3       phys     1500   unknown  --
net0          vnic     1500   up       e1000g0
```

示例 28 显示指定区域的链路

以下 show-link 命令显示特定非全局区域中的数据链路。

```
# dladm show-link -z zone1
LINK          CLASS    MTU    STATE    OVER
zone1/net0    vnic     1500   up       e1000g0
```

示例 29 从全局区域显示指定区域的链路

以下 show-link 命令从全局区域显示特定非全局区域中的数据链路。

```
# dladm show-link -z zone1
LINK          CLASS    MTU    STATE    OVER
zone1/net0    vnic     1500   up       e1000g0
```

示例 30 显示非全局区域中的链路

下面所示的 show-link 是从 zone1 进行调用的，仅显示该区域的数据链路。

请注意，在 show-link 输出中，不会显示 zone1/ 前缀。不显示前缀的原因是命令是从区域内部调用的。

```
# zlogin zone1
# dladm show-link -z zone1
LINK          CLASS    MTU    STATE    OVER
net0          vnic     1500   up       ?
```


示例 31 使用 -Z 选项显示当前区域

下面的命令假定以下情况：

- 链路 net1 当前已分配给 zoneA。条目 net1 和 zoneA/net1 表示同一链路。这两个条目的 ZONE 列相同且为链路当前分配给区域的名称。
- 链路 net2 未分配给任何非全局区域。
- 链路 zoneB/net2 是自动为 zoneB 创建的 VNIC。
- 链路 zoneC/net2 是自动为 zoneC 创建的 VNIC。
- 链路 zoneD/net2 是在 zoneD 内创建的 IP 隧道。与 net1 不同的是，net2 的每个条目均表示一个不同的链路。这些条目的 ZONE 列是不同的。

```
# dladm show-link -Z
LINK      ZONE      CLASS      MTU      STATE      OVER
e1000g0   global    phys       1500     up         - -
e1000g1   global    phys       1500     up         - -
net1      zoneA     vnic       1500     up         e1000g0
zoneA/net1 zoneA     vnic       1500     up         e1000g0
net2      global    vnic       1500     up         e1000g1
zoneB/net2 zoneB     vnic       1500     up         e1000g1
zoneC/net2 zoneC     vnic       1500     up         e1000g1
zoneD/net2 zoneD     iptun      65515    up         - -
```

属性

有关下列属性的描述，请参见 [attributes\(5\)](#)：

/usr/sbin

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）

/sbin

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）

请注意，对于 /usr/sbin 和 /sbin，已确定 show-aggr、show-link 和 show-vnic 子命令的 -s 和 -i 选项为过时选项。

请注意，对于 /usr/sbin 和 /sbin，show-linkprop 的 *effective 属性均具有 Volatile（可变）接口稳定性。

请注意，随上面的 `dladm` 子命令一起介绍的与网桥相关的子命令要求安装 `pkg://solaris/network/bridging` 软件包。

另请参见

[acctadm\(1M\)](#)、[autopush\(1M\)](#)、[dhcpage\(1M\)](#)、[dlstat\(1M\)](#)、[ifconfig\(1M\)](#)、[ipadm\(1M\)](#)、[ipseccor](#)

附注

在聚合子命令中引用聚合的首选方法是使用其链路名称。此外，还支持使用整数 *key* 引用聚合以便实现向后兼容，但这不是必需的。当创建聚合时，如果指定了 *key*（而非链路名称），`dladm` 将以 *aggrkey* 的形式自动生成此聚合的链路名称。

引用名 dlnmgtm - 数据链路管理守护进程

用法概要 /usr/sbin/dlnmgtm
svc:/network/datalink-management:default

描述 dlnmgtm 是处理网络数据链路接口的管理事件的系统守护进程。通过服务管理工具 (service management facility, SMF) 服务实例对其进行控制：
svc:/network/datalink-management:default
不能直接调用该守护进程。它不构成管理接口或编程接口的一部分。管理数据链路的管理接口是 [dlnadm\(1M\)](#)。

选项 该守护进程没有任何选项。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Private（专用）

另请参见 [dlnadm\(1M\)](#)、[attributes\(5\)](#)

引用名

dlstat – 报告数据链路统计信息

用法概要

```
dlstat [-r | -t] [-Z] [-i interval] [-z zone[,...]] [link]

dlstat [-a | -A] [-Z] [-i interval] [-p] [-o field[,...]]
      [-u R|K|M|G|T|P] [-z zone[,...]] [link]

dlstat show-phys [-r | -t] [-Z] [-i interval] [-a]
      [-p] [-o field[,...]] [-u R|K|M|G|T|P] [-z zone[,...]] [link]

dlstat show-link [-r | -t] [-Z] [-i interval] [-a]
      [-p] [-o field[,...]] [-u R|K|M|G|T|P] [-z zone[,...]] [link]

dlstat show-link -h [-Z] [-a] -f filename [-d] [-F format] [-s time]
      [-e time] [-z zone[,...]] [link]

dlstat show-aggr [-r | -t] [-Z] [-i interval] [-p] [-o field[,...]]
      [-u R|K|M|G|T|P] [-z zone[,...]] [link]
```

描述

dlstat 命令报告有关数据链路的运行时统计信息。[dladm\(1M\)](#) show-phys 向 dlstat show-phys 提供链路名称信息。[dladm\(1M\)](#) show-link 向 dlstat show-link 提供链路名称信息。[dladm\(1M\)](#) show-aggr 向 dlstat show-aggr 提供链路聚合信息。

在上面的 SYNOPSIS（用法概要）中显示了 dlstat 的六种命令格式。前两种格式没有子命令，而其余格式有子命令。下面的“子命令”对所有这六种格式进行了介绍。

选项

dlstat 命令具有以下选项和操作数，这些选项和操作数在下述“子命令”下显示的很多命令格式中是通用的（除非另行显式标记）。

-a

转储所有总量统计信息字段。

-i interval

指定刷新统计信息的间隔（以秒为单位）。缺省间隔为 1 秒。

-o field[,...]

显示要输出的输出字段的逗号分隔列表，不区分大小写。字段名称必须是下面列出的字段之一，或特殊值 all 以显示所有支持的字段。

支持的 RX 字段的列表：

- link (链路)
- index
- rbytes
- ipkts
- intrs (仅在不带有子命令或带有 show-link 的情况下适用)
- polls (仅在不带有子命令或带有 show-link 的情况下适用)
- idrops (仅在不带有子命令或带有 show-link 的情况下适用)

TX 端字段的列表：

- link (链路)

- index
- obytes
- opkts
- odrops（仅在不带有子命令或带有 show-link 的情况下适用）

对于 show-phys 子命令，index 列标识物理设备中的各个 RX 和 TX 硬环回。对于 show-link 子命令，index 列标识数据链路中的 RX 和 TX 硬件通道。有关硬环回和硬件通道的区别的说明，请参见《Oracle Solaris 管理：网络接口和网络虚拟化》。

-p
以稳定的机器可解析格式显示输出。

-r
仅显示接收端统计信息。包括接收的字节和包、硬件和软件删除等。有关完整列表，请参见“示例”。

-t
仅显示传输端统计信息。包括发送的字节和包、删除等。有关完整列表，请参见“示例”。

-u R|K|M|G|T|P
如果使用了该选项，则可以选择要显示的所有统计信息的单位，例如 R：原始计数，K：千位，M：兆位，T：兆兆位，P：千兆兆位。如果没有使用该选项，则将根据情况使用不同单位显示统计信息，采用 xy.zU 格式，其中 x、y 和 z 是数字，U 是相应的单位。

-Z
在输出中显示 ZONE 列。

-z zone[,...]
仅显示指定区域中链路的统计信息。缺省情况下，当 dlstat 从全局区域运行时，将显示所有区域中链路的统计信息。

从非全局区域运行时，dlstat 仅显示该区域中链路的统计信息。非全局区域无法看到其他区域中的链路。

link

如果指定了该选项，则仅为命名链路、物理设备 (show-phys) 或聚合 (show-aggr) 显示统计信息。否则，为所有链路、设备或聚合显示统计信息。

子命令

dlstat 支持以下命令格式。

dlstat [-r | -t] [-Z] [-i interval] [-z zone[,...]] [link]

重复检查所有链路并报告统计信息。按链路利用率的降序顺序排列输出。如果没有指定链路，则系统为所有链路显示统计信息。针对每个链路显示流量统计信息，而不是针对每个物理设备。例如，对于物理链路上配置的 VNIC，基础物理链路的统计信息中不会反映流过该 VNIC 的流量。但是，链路统计信息将包含与在该链路顶部配置的用户定义的流相匹配的流量。

此命令格式具有一个上面的“选项”中没有介绍的选项：

-A

转储此数据链路的所有统计信息字段。此命令的输出统计信息包括其他所有 `dlstat` 命令报告的所有统计信息。

`dlstat [-a | -A] [-Z] [-i interval] [-p] [-o field[,...]] [-u R|K|M|G|T|P] [-z zone[,...]] [link]`
允许指定要显示的统计信息。

此命令格式的选项在上面的“选项”中进行了介绍。

`dlstat show-phys [-r | -t] [-Z] [-i interval] [-a] [-p] [-o field[,...]] [-u R|K|M|G|T|P] [-z zone[,...]] [link]`

显示物理设备的统计信息。

此子命令的选项在上面的“选项”中进行了介绍。

`dlstat show-link [-r | -t] [-Z] [-i interval] [-a] [-p] [-o field[,...]] [-u R|K|M|G|T|P] [-z zone[,...]] [link]`

显示链路的统计信息。

`dlstat show-link -h [-Z] [-a] -f filename [-d] [-F format] [-s time] [-e time] [-z zone[,...]] [link]`

显示存储的扩展记帐文件中的网络使用历史。要使用此语法，必须已经使用 [acctadm\(1M\)](#) 预先配置并启用网络记帐。缺省输出是启用扩展记帐的整个周期中现有链路的网络使用情况的摘要。

link 参数如上面的“选项”中所述。

-a

显示启用扩展记帐的指定周期内的所有历史网络使用情况。这包括有关已经删除的链路的使用情况信息。

-f *filename*

指定从中读取网络使用历史的扩展记帐记录的文件。

-d

显示具有日志信息的日期。日期格式为 *mm/dd/yyyy*。

-F *format*

指定网络使用历史信息的输出格式。gnuplot 是唯一支持的格式。

-s *time*

-e *time*

指定数据显示的开始时间和停止时间。时间格式为 *MM/DD/YYYY*，*hh:mm:ss*。hh 采用 24 小时制表示法。

`dlstat show-aggr [-r | -t] [-Z] [-i interval] [-p] [-o field[,...]] [-u R|K|M|G|T|P] [-z zone[,...]] [link]`

显示聚合的每端口统计信息。

此子命令的选项在上面的“选项”中进行了介绍。

示例**示例1 显示统计信息**

要显示所有链路的统计信息，请输入以下命令。统计信息显示为小数点前 3 位数字、小数点后 2 位数字且具有相应的单位。

```
# dlstat
      LINK    IPKTS  RBYTES   OPKTS  OBYTES
e1000g0 101.88K  32.86M  40.16K  4.37M
nxge1   4.50M   6.78G   1.38M  90.90M
vnic1      8    336        0      0
net0    73.96K   6.81M        0      0
zone1/net0 144.47K  13.32M    247  16.29K
zone2/net0 132.89K  12.25M    236  15.82K
```

示例2 显示RX端统计信息

以下命令显示每秒钟接收端的统计信息。

```
# dlstat -r -i 1
      LINK    IPKTS  RBYTES   INTRS   POLLS   IDROPS
e1000g0 101.91K  32.86M  87.56K  14.35K      0
nxge1   9.61M  14.47G   5.79M   3.82M      0
vnic1      8    336        0      0      0
e1000g0      0      0        0      0      0
nxge1   82.13K 123.69M  50.00K  32.13K      0
vnic1      0      0        0      0      0
.          .      .      .      .      .
.          .      .      .      .      .
```

示例3 显示每个硬件环的统计信息

以下命令基于每个接收端硬件环显示统计信息。

```
# dlstat show-phys -r nxge1
      LINK TYPE  INDEX   IPKTS  RBYTES
nxge1  rx      0        0      0
nxge1  rx      1        0      0
nxge1  rx      2    1.73M  2.61G
nxge1  rx      3        0      0
nxge1  rx      4    8.44M 12.71G
nxge1  rx      5    5.68M  8.56G
nxge1  rx      6    4.90M  7.38G
nxge1  rx      7        0      0
```

示例4 显示每个通道的统计信息

以下命令基于每个接收端通道显示统计信息。首先，具有专用硬件通道的接口：

```
# dlstat show-link -r nxge1
      LINK TYPE  ID INDEX   IPKTS  RBYTES   INTRS   POLLS   IDROPS
nxge1  rx  local  --      0      0      0      0      0
nxge1  rx   hw    1      0      0      0      0      0
```

示例 4 显示每个通道的统计信息 (续)

nxge1	rx	hw	2	1.73M	2.61G	1.33M	400.22K	0
nxge1	rx	hw	3	0	0	0	0	0
nxge1	rx	hw	4	8.44M	12.71G	4.35M	4.09M	0
nxge1	rx	hw	5	5.68M	8.56G	3.72M	1.97M	0
nxge1	rx	hw	6	4.90M	7.38G	3.11M	1.80M	0
nxge1	rx	hw	7	0	0	0	0	0

接下来，没有专用硬件通道（即，仅软件通道）的接口：

```
# dlstat show-link -r ixgbe0
LINK TYPE ID INDEX IPKTS RBYTES INTRIS POLLS IDROPS
ixgbe0 rx local -- 0 0 0 0 0
ixgbe0 rx sw -- 794.28K 1.19G 794.28K 0 0
```

示例 5 显示传输端统计信息

以下命令按照五秒钟间隔显示传输端统计信息。

```
# dlstat -t -i 5
LINK OPKTS OBYTES ODROPS
e1000g0 40.24K 4.37M 0
nxge1 9.76M 644.14M 0
vnic1 0 0 0
e1000g0 0 0 0
nxge1 26.82K 1.77M 0
vnic1 0 0 0
. . .
. . .
. . .
```

示例 6 显示传输端环统计信息

以下命令显示传输端硬件环统计信息。

```
# dlstat show-phys -t nxge1
LINK TYPE INDEX OPKTS OBYTES
nxge1 tx 0 44 3.96K
nxge1 tx 1 0 0
nxge1 tx 2 1.48M 121.68M
nxge1 tx 3 2.45M 201.11M
nxge1 tx 4 1.47M 120.82M
nxge1 tx 5 0 0
nxge1 tx 6 1.97M 161.57M
nxge1 tx 7 4.59M 376.21M
nxge1 tx 8 2.43M 199.24M
nxge1 tx 9 0 0
nxge1 tx 10 3.23M 264.69M
nxge1 tx 11 1.88M 153.96M
```


示例7 显示传输端通道统计信息

以下命令显示传输端通道统计信息。

```
# d1stat show-link -t nxge1
LINK TYPE      ID INDEX  OPKTS  OBYTES  ODROPS
nxge1  tx      hw      0      32    1.44K    0
nxge1  tx      hw      1       0       0       0
nxge1  tx      hw      2    1.48M   97.95M    0
nxge1  tx      hw      3    2.45M  161.87M    0
nxge1  tx      hw      4    1.47M   97.25M    0
nxge1  tx      hw      5       3     276       0
nxge1  tx      hw      6    1.97M  130.05M    0
nxge1  tx      hw      7    4.59M  302.80M    0
nxge1  tx      hw      8    2.42M  302.80M    0
nxge1  tx      hw      9       0       0       0
nxge1  tx      hw     10    3.23M  213.05M    0
nxge1  tx      hw     11    1.88M  123.93M    0
```

示例8 选择特定的统计信息集

以下命令显示您如何选择特定的统计信息集。

```
# d1stat show-link -r -o LINK,TYPE,ID,INDEX,INTRS,POLLS nxge1
LINK TYPE      ID INDEX  INTRS  POLLS
nxge1  rx    local  --      0      0
nxge1  rx    other  --      0      0
nxge1  rx     hw     1       0      0
nxge1  rx     hw     2    2.47M  753.90K
nxge1  rx     hw     3       0      0
nxge1  rx     hw     4    8.24M   7.72M
nxge1  rx     hw     5    6.96M   3.68M
nxge1  rx     hw     6    5.82M   3.36M
nxge1  rx     hw     7       0      0
```

示例9 显示历史网络使用情况

通过将扩展记帐工具 [acctadm\(1M\)](#) 与以下命令一起使用，可以存储网络使用历史统计信息：

```
# acctadm -e basic -f /var/log/net.log net
acctadm net
Network accounting: active
    Network accounting file: /var/log/net.log
    Tracked Network resources: basic
Untracked Network resources: src_ip,dst_ip,src_port,dst_port,protocol,
                             dsfield
```

然后，可以使用以下命令以摘要格式检索保存的历史数据：

```
# d1stat show-link -h -f /var/log/net.log
LINK      DURATION  IPACKETS  RBYTES      OPACKETS  OBYTES      BANDWIDTH
```

示例9 显示历史网络使用情况 (续)

e1000g0 80 1031 546908 0 0 2.44 Kbps

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

/usr/sbin

属性类型	属性值
可用性	system/core-os
接口稳定性	请参见下文

屏幕输出为 Uncommitted（未确定）。调用为 Committed（已确定）。

另请参见 [acctadm\(1M\)](#)、[dladm\(1M\)](#)、[ifconfig\(1M\)](#)、[kstat\(1M\)](#)、[netstat\(1M\)](#)、[attributes\(5\)](#)

引用名 dmesg – collect system diagnostic messages to form error log

用法概要 /usr/bin/dmesg
/usr/sbin/dmesg

描述 dmesg is made obsolete by [syslogd\(1M\)](#) for maintenance of the system error log.

dmesg looks in a system buffer for recently printed diagnostic messages and prints them on the standard output.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [syslogd\(1M\)](#), [attributes\(5\)](#)

引用名	dminfo – report information about a device entry in a device maps file														
用法概要	<pre>dminfo [-v] [-a] [-f <i>pathname</i>] dminfo [-v] [-a] [-f <i>pathname</i>] -n <i>dev</i> -name... dminfo [-v] [-a] [-f <i>pathname</i>] -d <i>dev</i> -path... dminfo [-v] [-a] [-f <i>pathname</i>] -t <i>dev</i> -type... dminfo [-v] [-f <i>pathname</i>] -u <i>dm</i> -entry</pre>														
描述	dminfo reports and updates information about the device_maps(4) file.														
选项	The following options are supported <table><tr><td>-a</td><td>Succeed if any of the requested entries are found. If used with -v, all entries that match the requested case(s) are printed.</td></tr><tr><td>-d <i>dev-path</i></td><td>Search by <i>dev-path</i>. Search device_maps(4) for a device special pathname in the <i>device_list</i> field matching the <i>dev-path</i> argument. This option cannot be used with -n, -t or -u.</td></tr><tr><td>-f <i>pathname</i></td><td>Use a device_maps file with <i>pathname</i> instead of /etc/security/device_maps.</td></tr><tr><td>-n <i>dev-name</i></td><td>Search by <i>dev-name</i>. Search device_maps(4) for a <i>device_name</i> field matching <i>dev-name</i>. This option cannot be used with -d, -t or -u.</td></tr><tr><td>-t <i>dev-type</i></td><td>Search by <i>dev-type</i>. Search device_maps(4) for a <i>device_type</i> field matching the given <i>dev-type</i>. This option cannot be used with -d, -n or -u.</td></tr><tr><td>-u <i>dm-entry</i></td><td>Update the device_maps(4) file. This option is provided to add entries to the device_maps(4) file. The <i>dm-entry</i> must be a complete device_maps(4) file entry. The <i>dm-entry</i> has fields, as in the device_maps file. It uses the colon (:) as a field separator, and white space as the <i>device_list</i> subfield separators. The <i>dm-entry</i> is not made if any fields are missing, or if the <i>dm-entry</i> would be a duplicate. The default device maps file can be updated only by the super user.</td></tr><tr><td>-v</td><td>Verbose. Print the requested entry or entries, one line per entry, on the standard output. If no entries are specified, all are printed.</td></tr></table>	-a	Succeed if any of the requested entries are found. If used with -v, all entries that match the requested case(s) are printed.	-d <i>dev-path</i>	Search by <i>dev-path</i> . Search device_maps(4) for a device special pathname in the <i>device_list</i> field matching the <i>dev-path</i> argument. This option cannot be used with -n, -t or -u.	-f <i>pathname</i>	Use a device_maps file with <i>pathname</i> instead of /etc/security/device_maps.	-n <i>dev-name</i>	Search by <i>dev-name</i> . Search device_maps(4) for a <i>device_name</i> field matching <i>dev-name</i> . This option cannot be used with -d, -t or -u.	-t <i>dev-type</i>	Search by <i>dev-type</i> . Search device_maps(4) for a <i>device_type</i> field matching the given <i>dev-type</i> . This option cannot be used with -d, -n or -u.	-u <i>dm-entry</i>	Update the device_maps(4) file. This option is provided to add entries to the device_maps(4) file. The <i>dm-entry</i> must be a complete device_maps(4) file entry. The <i>dm-entry</i> has fields, as in the device_maps file. It uses the colon (:) as a field separator, and white space as the <i>device_list</i> subfield separators. The <i>dm-entry</i> is not made if any fields are missing, or if the <i>dm-entry</i> would be a duplicate. The default device maps file can be updated only by the super user.	-v	Verbose. Print the requested entry or entries, one line per entry, on the standard output. If no entries are specified, all are printed.
-a	Succeed if any of the requested entries are found. If used with -v, all entries that match the requested case(s) are printed.														
-d <i>dev-path</i>	Search by <i>dev-path</i> . Search device_maps(4) for a device special pathname in the <i>device_list</i> field matching the <i>dev-path</i> argument. This option cannot be used with -n, -t or -u.														
-f <i>pathname</i>	Use a device_maps file with <i>pathname</i> instead of /etc/security/device_maps.														
-n <i>dev-name</i>	Search by <i>dev-name</i> . Search device_maps(4) for a <i>device_name</i> field matching <i>dev-name</i> . This option cannot be used with -d, -t or -u.														
-t <i>dev-type</i>	Search by <i>dev-type</i> . Search device_maps(4) for a <i>device_type</i> field matching the given <i>dev-type</i> . This option cannot be used with -d, -n or -u.														
-u <i>dm-entry</i>	Update the device_maps(4) file. This option is provided to add entries to the device_maps(4) file. The <i>dm-entry</i> must be a complete device_maps(4) file entry. The <i>dm-entry</i> has fields, as in the device_maps file. It uses the colon (:) as a field separator, and white space as the <i>device_list</i> subfield separators. The <i>dm-entry</i> is not made if any fields are missing, or if the <i>dm-entry</i> would be a duplicate. The default device maps file can be updated only by the super user.														
-v	Verbose. Print the requested entry or entries, one line per entry, on the standard output. If no entries are specified, all are printed.														
退出状态	<table><tr><td>0</td><td>Successful completion.</td></tr><tr><td>1</td><td>Request failed.</td></tr><tr><td>2</td><td>Incorrect syntax.</td></tr></table>	0	Successful completion.	1	Request failed.	2	Incorrect syntax.								
0	Successful completion.														
1	Request failed.														
2	Incorrect syntax.														
文件	/etc/security/device_maps														

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [device_maps\(4\)](#), [attributes\(5\)](#)

引用名	dns-sd – 多播 DNS (Multicast DNS, mDNS) 和 DNS 服务发现 (DNS Service Discovery, DNS-SD) 测试工具
用法概要	<pre>dns-sd -R name type domain port [key=value ...] dns-sd -B type domain dns-sd -L name type domain dns-sd -Q FQDN rrtype rrclass dns-sd -C FQDN rrtype rrclass dns-sd -P name type domain port host IP [key=value ...] dns-sd -E -F -A -U -N -T -M -I</pre>
描述	dns-sd 命令是一种网络诊断工具，类似于 ping(1M) 或 traceroute(1M)。但是，与这些工具不同，其大多数功能不是在 dns-sd 可执行文件本身中实现，而是通过任何应用程序都可使用的库代码实现。dns-sd 使用的库 API 在 <code>/usr/include/dns_sd.h</code> 中进行了介绍。 dns-sd 命令主要用于交互使用。由于其命令行参数和输出格式可能会发生更改，因此从 shell 脚本调用该命令可能是无法预测的。此外，DNS 服务发现的异步性质使其无法轻易地用于面向脚本的编程。这种形式的异步交互最适合以下应用程序：多线程应用程序或使用主事件处理循环在击键、网络数据和其他异步事件发生时接收通知的应用程序。
选项	所支持的选项如下： -R name type domain port [key=value ...] 在具有给定 <i>name</i> 和 <i>type</i> 的指定域中注册（通告）服务，以在当前计算机的指定 <i>port</i> 上侦听。 <i>name</i> 可以是包含合法 unicode 字符（包括没有任何限制的点、空格、斜杠、冒号等）的任意 unicode 文本，其最大长度是 63 个 UTF-8 字节。 <i>type</i> 的格式必须为 <code>"_app-proto._tcp"</code> 或 <code>"_app-proto_udp"</code>，其中 <code>"app-proto"</code> 是 http://www.dns-sd.org 中注册的应用程序协议名称，在服务类型 (RFC 2782) 链接下。 <i>domain</i> 是要注册服务的域。在当前实现中，仅支持本地多播域 <code>"local"</code>。将来可以在具有工作的 DNS 更新服务器 [RFC2136] 的任意域中注册。域 <code>"."</code> 表示“选择合适缺省值”，当前表示 <code>"local"</code> 域。 <i>port</i> 是从 0 到 65535 的编号，并且是服务侦听的 TCP 或 UDP 端口号。通过在端口 0 上注册服务，应用程序可以显式通告服务的不可用性。 存储在已通告服务的 DNS TXT 记录中的 <i>key/value</i> 对可能会选择性地介绍服务的其他属性。允许的键和值与 http://www.dns-sd.org 中注册的服务一起列在服务类型 (RFC 2782) 链接下。

-B type domain

浏览 *domain* 中 *type* 服务的实例。

有关有效的类型，请参见 <http://www.dns-sd.org> 中服务类型 (RFC 2782) 链接下的内容。忽略域名或使用 "." 意味着“选择合适缺省值”。

-L name type domain

查找并显示联系和使用命名服务所需的信息。此信息包括可以使用该服务的计算机的主机名、服务侦听的端口号以及（如果有）说明服务属性的 TXT 记录属性。

在典型应用程序中，很少发生浏览，但每次使用服务时都发生查找（或“解析”）。例如，用户不会经常浏览网络以选择缺省打印机，但是选择了缺省打印机后，用户每次按 Cmd-P 进行打印时，该命名服务都将解析为其当前 IP 地址和端口号。

-Q FQDN rrtype rrclass

任何资源记录类型和类的通用查询。

-C FQDN rrtype rrclass

任何资源记录类型和类的通用查询。此选项还重新确认每个查询结果。重新确认记录指示 **mdnsd(1M)** 验证记录的有效性。如果记录无效，**mdnsd(1M)** 会在守护进程高速缓存中刷新该记录，还会在网络上其他 **mdnsd(1M)** 高速缓存中刷新该记录。

-P name type domain port host IP [key=value ...]

在具有给定 *name* 和 *type* 的指定域中注册（通告）服务，该服务侦听指定端口且在其他主机上可访问该服务。此选项应该用于由代理通告可在其他主机上访问的服务。必须指定访问服务的主机名和 IPv4 地址。

-E

发现建议的注册域。此选项返回建议的注册服务的域。通过查询 **resolv.conf(4)** 中的名称服务器返回建议的注册域。

-F

发现建议的浏览域。此选项返回建议的浏览服务的域。通过查询 **resolv.conf(4)** 中的名称服务器返回建议的浏览域。

-A

测试使用多播 DNS 注册服务，并测试使用多播 DNS 的 DNS 记录的添加、更新和删除操作。

-U

测试使用多播 DNS 注册服务，并测试为使用多播 DNS 注册的服务更新 DNS TXT 记录。

-N

测试为使用多播 DNS 注册的服务添加大型 NULL 记录。

-T

测试为使用多播 DNS 注册的服务添加大型 TXT 记录。

- M
测试创建具有多个 TXT 记录的注册。
- I
测试注册并立即更新 TXT 记录。

示例

示例1 通告打印服务
以下命令通告此计算机的端口 515 上存在 LPR 打印服务，因此 DNS-SD 兼容打印客户机可以使用该服务：

```
dns-sd -R "My Test" _printer._tcp. . 515 pdl=application/postscript
```

要使该注册可用，端口 515 上应该提供 LPR 服务。通告不存在的服务起不了作用。

示例2 通告 Web 页
以下命令通告此计算机端口 80 上的 HTTP 服务器提供支持的 Web 页，因此该页将在 Safari 中的 Bonjour 列表上显示或在其他 DNS-SD 兼容 Web 客户机上显示：

```
dns-sd -R "My Test" _http._tcp . 80 path=/path-to-page.html
```

示例3 在本地网络上查找通告的 Web 页
以下命令在本地网络（Safari 显示的相同列表）上查找通告的 Web 页：

```
dns-sd -B _http._tcp
```

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/network/dns/mdns
接口稳定性	Volatile（可变）

另请参见

[mdnsd\(1M\)](#)、[ping\(1M\)](#)、[ping\(1M\)](#)、[resolv.conf\(4\)](#)、[attributes\(5\)](#)

引用名	dnsssec-dsfromkey – DNSSEC DS RR 生成工具
用法概要	<pre>dnsssec-dsfromkey [-v level] [-1] [-2] [-a alg] keyfile dnsssec-dsfromkey -s [-v level] [-1] [-2] [-a alg] [-c class] [-d dir] keyfile</pre>
描述	dnsssec-dsfromkey
选项	所支持的选项如下： -1 将 SHA-1 用作摘要算法。缺省情况下，同时使用 SHA-1 和 SHA-256。 -2 将 SHA-256 用作摘要算法。 -a <i>algorithm</i> 选择摘要算法。<i>algorithm</i> 的值必须是 SHA-1 (SHA1) 或 SHA-256 (SHA256) 之一。这些值不区分大小写。 -v <i>level</i> 设置调试级别。 -s 密钥集模式：代替密钥文件名称，该参数是密钥集文件的 DNS 域名。-c 和 -d 选项仅在该模式下才有意义。 -c <i>class</i> 指定 DNS 类（缺省类是 IN）；仅在密钥集模式下才有用。 -d <i>directory</i> 在作为目录的 <i>directory</i> 中查找密钥集文件；不处于密钥集模式下时忽略。
示例	要从 Kexample.com.+003+26160 密钥文件名称构建 SHA-256 DS RR，请使用以下命令： <pre># dnsssec-dsfromkey -2 Kexample.com.+003+26160</pre> 此命令将生成与下面类似的输出： <pre>example.com. IN DS 26160 5 2 3A1EADA7A74B8D0BA86726B0C227AA85AB8BBD2B2004F41A868A54F0 C5EA0B94</pre>
文件	该密钥文件可以通过密钥标识 <i>Knnnnn.+aaa+iinii</i> 或完整文件名 <i>Knnnnn.+aaa+iinii.key</i> 指定，与由 dnsssec-keygen(1M) 生成的一样。 密钥集文件名基于目录、字符串 <i>keyset-</i> 和 <i>dnsname</i> 构建而成。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	service/network/dns/bind
接口稳定性	Volatile（可变）

另请参见 [dnssec-keygen\(1M\)](#)、[dnssec-signzone\(1M\)](#)、[attributes\(5\)](#)

RFC 3658、RFC 4509

请参见 BIND 9 Administrator's Reference Manual。从本手册页发布之日起，将在 <https://www.isc.org/software/bind/documentation> 上提供该文档。

注意 密钥文件错误可以生成“未找到文件”消息，即使该文件存在也是如此。

引用名	dnssec-keyfromlabel – DNSSEC 密钥生成工具
用法概要	<pre>dnssec-keyfromlabel -a <i>algorithm</i> -l <i>label</i> [-c <i>class</i>] [-f <i>flag</i>] [-k] [-n <i>nametype</i>] [-p <i>protocol</i>] [-t <i>type</i>] [-v <i>level</i>] <i>name</i></pre>
描述	dnssec-keyfromlabel 从加密硬件设备使用指定的标签检索密钥，并为 DNSSEC（安全 DNS）生成密钥文件，如 RFC 2535 和 RFC 4034 中所定义的那样。
选项	所支持的选项如下： -a <i>algorithm</i> 选择加密算法。<i>algorithm</i> 的值必须是 RSAMD5 (RSA) 或 RSASHA1、DSA、NSEC3RSASHA1、NSEC3DSA 或 DH (Diffie-Hellman) 之一。这些值不区分大小写。 请注意，对于 DNSSEC，RSASHA1 是强制实现的算法；DSA 是建议的算法。另请注意，DH 自动设置 -k 标志。 -L <i>label</i> 指定加密硬件 (PKCS#11) 设备中的密钥标签。 -n <i>nametype</i> 指定密钥的所有者类型。<i>nametype</i> 的值必须是 ZONE（对于 DNSSEC 区域密钥 (KEY/DNSKEY)）、HOST 或 ENTITY（对于与主机相关的密钥 (KEY)）、USER（对于与用户相关的密钥 (KEY)）或 OTHER (DNSKEY)。这些值不区分大小写。 -c <i>class</i> 表明包含密钥的 DNS 记录应该具有指定类。如果没有指定，将使用类 IN。 -f <i>flag</i> 在 KEY/DNSKEY 记录的标志字段中设置指定的标志。唯一识别的标志是 KSK（Key Signing Key，密钥签名密钥）DNSKEY。 -h 显示 dnssec-keyfromlabel 的选项和参数的简短摘要。 -k 生成 KEY 记录，而不是 DNSKEY 记录。 -p <i>protocol</i> 为生成的密钥设置协议值。协议是 0 到 255 之间的数字。缺省值是 3 (DNSSEC)。在 RFC 2535 及其后续 RFC 中列出了此参数的其他可能值。 -T <i>type</i> 表明密钥的用途。<i>type</i> 必须是 AUTHCONF、NOAUTHCONF、NOAUTH 或 NOCONF 之一。缺省值是 AUTHCONF。AUTH 指的是验证数据的能力，CONF 指的是加密数据的能力。 -v <i>level</i> 设置调试级别。

生成的密钥文件 当 `dnssec-keyfromlabel` 成功完成时，会向标准输出中显示 `Knnnn.+aaa+iiii` 格式的字符串。这是其生成的密钥文件的标识字符串，其意义如下。

- `nnnn` 是密钥名称。
- `aaa` 是算法的数字表示。
- `iiii` 是密钥标识符（或足迹）。

`dnssec-keyfromlabel` 创建两个文件，并根据显示的字符串命名这两个文件。`Knnnn.+aaa+iiii.key` 包含公钥，`Knnnn.+aaa+iiii.private` 包含私钥。

第一个文件包含 DNS KEY 记录，该记录可以直接插入到区域文件，也可以使用 `$INCLUDE` 语句插入。

第二个文件包含算法特定的字段。出于安全原因，此文件不具有一般读取权限。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/network/dns/bind
接口稳定性	Volatile（可变）

另请参见 [dnssec-keygen\(1M\)](#), [dnssec-signzone\(1M\)](#), [attributes\(5\)](#)

RFC 2539, RFC 2845, RFC 4033

请参见 BIND 9 Administrator's Reference Manual。从本手册页发布之日起，将在 <https://www.isc.org/software/bind/documentation> 上提供该文档。

引用名	dnssec-keygen – DNSSEC 密钥生成工具
用法概要	<pre>dnssec-keygen -a algorithm -b keysize -n nametype [-ehk] [-c class] [-f flag] [-g generator] [-p protocol] [-r randomdev] [-s strength] [-t type] [-v level] name</pre>
描述	dnssec-keygen 实用程序为 DNSSEC（安全 DNS）生成密钥，如 RFC 2535 和 RFC 4034 中定义的那样。它还可以生成与 TSIG（Transaction Signature，事务签名）一起使用的密钥，如 RFC 2845 中定义的那样。
选项	所支持的选项如下： -a algorithm 选择加密算法。algorithm 的值必须是 RSAMD5 (RSA) 或 RSASHA1 之一、DSA、NSEC3RSASHA1、NSEC3DSA、DH (Diffie-Hellman) 或 HMAC-MD5。这些值不区分大小写。 对于 DNSSEC，RSASHA1 是强制实现的算法；DSA 是建议的算法。对于 TSIG，HMAC-MD5 是强制算法。 注 -HMAC-MD5 和 DH 自动设置 -k 标志。 -b keysize 指定密钥中的位数。密钥大小的选择取决于使用的算法。RSAMD5 和 RSASHA1 密钥必须在 512 和 2048 位之间。Diffie-Hellman 密钥必须在 128 和 4096 位之间。DSA 密钥必须在 512 和 1024 位之间，并且必须是 64 的整数倍。HMAC-MD5 密钥必须在 1 位和 512 位之间。 -c class 表明包含密钥的 DNS 记录应该具有指定类。如果没有指定，将使用类 IN。 -e 如果生成 RSAMD5 或 RSASHA1 密钥，则使用大指数。 -f flag 在 KEY/DNSKEY 记录的标志字段中设置指定的标志。唯一识别的标志是 KSK（Key Signing Key，密钥签名密钥）DNSKEY。 -g generator 如果生成 Diffie Hellman 密钥，则使用此 generator。允许的值是 2 和 5。如果没有指定 generator，则将使用 RFC 2539 中的已知素数（如果可能）；否则，缺省值是 2。 -h 列出 dnssec-keygen 的选项和参数的简短摘要。 -k 生成 KEY 记录，而不是 DNSKEY 记录。

-n *nametype*

指定密钥的所有者类型。*nametype* 的值必须是 **ZONE**（对于 DNSSEC 区域密钥 (KEY/DNSKEY)）、**HOST** 或 **ENTITY**（对于与主机相关的密钥 (KEY)）、**USER**（对于与用户相关的密钥 (KEY)）或 **OTHER** (DNSKEY)。这些值不区分大小写。缺省值是 **ZONE**（用于生成 DNSKEY）。

-p *protocol*

为生成的密钥设置协议值。*protocol* 参数是 0 到 255 之间的数字。缺省值是 3 (DNSSEC)。在 RFC 2535 及其后续版本中列出了此参数的其他可能值。

-r *randomdev*

指定随机源。如果操作系统不提供 `/dev/random` 或等效设备，则缺省的随机源是键盘输入。*randomdev* 指定字符设备的名称或包含要使用的随机数据的文件（而非缺省文件）。特殊值 "keyboard" 表示应该使用键盘输入。

-s *strength*

指定密钥的强度值。*strength* 参数是 0 到 15 之间的数字，且当前尚未在 DNSSEC 中定义其用途。

-T *type*

表明密钥的用途。*type* 必须是 **AUTHCONF**、**NOAUTHCONF**、**NOAUTH** 或 **NOCONF** 之一。缺省值是 **AUTHCONF**。**AUTH** 指的是验证数据的能力，**CONF** 指的是加密数据的能力。

-v *level*

设置调试级别。

生成的密钥

当 `dnssec-keygen` 成功完成时，会向标准输出中输出 `Knnnn.+aaa+iiii` 格式的字符串。这是其生成的密钥的标识字符串。

- *nnnn* 是密钥名称。
- *aaa* 是算法的数字表示。
- *iiii* 是密钥标识符（或足迹）。

`dnssec-keygen` 实用程序创建两个文件，并根据列出的字符串命名这两个文件。

- `Knnnn.+aaa+iiii.key` 包含公钥。
- `Knnnn.+aaa+iiii.private` 包含私钥。

`.key` 文件包含 DNS KEY 记录，该记录可以直接插入到区域文件，也可以使用 `$INCLUDE` 语句插入。

`.private` 文件包含算法特定的字段。出于安全原因，此文件不具有一般读取权限。

对于对称加密算法（如 HMAC-MD5），将生成 `.key` 和 `.private` 文件，即使公钥和私钥等效也如此。

示例

示例1 生成 768 位 DSA 密钥

要为域 example.com 生成 768 位 DSA 密钥，则将发出以下命令：

```
dnssec-keygen -a DSA -b 768 -n ZONE example.com
```

该命令将列出以下格式的字符串：

```
Kexample.com.+003+26160
```

将创建以下文件：

```
Kexample.com.+003+26160.key
Kexample.com.+003+26160.private
```

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/network/dns/bind
接口稳定性	Volatile（可变）

另请参见

[dnssec-signzone\(1M\)](#)、[attributes\(5\)](#)

RFC 2539, RFC 2845, RFC 4033

请参见 BIND 9 Administrator's Reference Manual。从本手册页发布之日起，将在 <https://www.isc.org/software/bind/documentation> 上提供该文档。

引用名	dnsssec-makekeyset – DNSSEC zone signing tool
用法概要	dnsssec-makekeyset [-ahp] [-s <i>start-time</i>] [-e <i>end-time</i>] [-r <i>randomdev</i>] [-t <i>tll</i>] [-v <i>level</i>] <i>key</i> ...
描述	The dnsssec-makekeyset utility generates a key set from one or more keys created by dnsssec-keygen(1M) . It creates a file containing a KEY record for each key, and self-signs the key set with each zone key. The output file is of the form <i>keyset-nnnn.</i> , where <i>nnnn</i> is the zone name.
选项	-a Verify all generated signatures. -e <i>end-time</i> Specify the date and time when the generated SIG records expire. As with <i>start-time</i>, an absolute time is indicated in YYYYMMDDHHMMSS notation. A time relative to the start time is indicated with +N, which is N seconds from the start time. A time relative to the current time is indicated with now+N. If no <i>end-time</i> is specified, 30 days from the start time is used as a default. -h Print a short summary of the options and arguments to dnsssec-makekeyset(). -p Use pseudo-random data when signing the zone. This is faster, but less secure, than using real random data. This option may be useful when signing large zones or when the entropy source is limited. -r <i>randomdev</i> Specify the source of randomness. If the operating system does not provide a /dev/random or equivalent device, the default source of randomness is keyboard input. The <i>randomdev</i> argument specifies the name of a character device or file containing random data to be used instead of the default. The special value keyboard indicates that keyboard input should be used. -s <i>start-time</i> Specify the date and time when the generated SIG records become valid. This can be either an absolute or relative time. An absolute start time is indicated by a number in YYYYMMDDHHMMSS notation; 20000530144500 denotes 14:45:00 UTC on May 30th, 2000. A relative start time is indicated by +N, which is N seconds from the current time. If no <i>start-time</i> is specified, the current time is used. -t <i>tll</i> Specify the TTL (time to live) of the KEY and SIG records. The default is 3600 seconds. -v <i>level</i> Set the debugging level.
操作数	The following operands are supported: <i>key</i> The list of keys to be included in the keyset file. These keys are expressed in the form <i>Knnnn.+aaa+iiii</i> as generated by dnsssec-keygen.

示例

示例 1 Generates a keyset containing the DSA key for example.com.

The following command generates a keyset containing the DSA key for example.com generated in the [dnssec-keygen\(1M\)](#) manual page.

```
dnssec-makekeyset -t 86400 -s 20000701120000 -e +2592000 \
Kexample.com.+003+26160
```

In this example, dnssec-makekeyset () creates the file keyset-example.com. This file contains the specified key and a self-generated signature.

The DNS administrator for example.com could send keyset-example.com. to the DNS administrator for .com for signing, if the .com zone is DNSSEC-aware and the administrators of the two zones have some mechanism for authenticating each other and exchanging the keys and signatures securely.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWbind9
Interface Stability	Volatile

另请参见

[dnssec-keygen\(1M\)](#), [dnssec-signkey\(1M\)](#), [attributes\(5\)](#)

RFC 2535

BIND 9 Administrator Reference Manual

附注

Source for BIND9 is available in the SUNWbind9S package.

引用名	dnsssec-signkey – DNSSEC key set signing tool
用法概要	dnsssec-signkey [-ahp] [-c <i>class</i>] [-e <i>end-time</i>] [-r <i>randomdev</i>] [-s <i>start-time</i>] [-v <i>level</i>] <i>keyset</i> <i>key</i> ...
描述	The dnsssec-signkey utility signs a keyset. Typically the keyset will be for a child zone and will have been generated by dnsssec-makekeyset(1M) . The child zone's keyset is signed with the zone keys for its parent zone. The output file is of the form <i>signedkey-nnnn.</i> , where <i>nnnn</i> is the zone name.
选项	The following options are supported: -a Verify all generated signatures. -c <i>class</i> Specify the DNS class of the key sets. -e <i>end-time</i> Specify the date and time when the generated SIG records expire. As with <i>start-time</i>, an absolute time is indicated in YYYYMMDDHHMMSS notation. A time relative to the start time is indicated with +<i>N</i>, which is <i>N</i> seconds from the start time. A time relative to the current time is indicated with now+<i>N</i>. If no <i>end-time</i> is specified, 30 days from the start time is used as a default. -h Prints a short summary of the options and arguments to dnsssec-signkey(). -p Use pseudo-random data when signing the zone. This is faster, but less secure, than using real random data. This option may be useful when signing large zones or when the entropy source is limited. -r <i>randomdev</i> Specify the source of randomness. If the operating system does not provide a /dev/random or equivalent device, the default source of randomness is keyboard input. <i>randomdev</i> specifies the name of a character device or file containing random data to be used instead of the default. The special value keyboard indicates that keyboard input should be used. -s <i>start-time</i> Specify the date and time when the generated SIG records become valid. This can be either an absolute or relative time. An absolute start time is indicated by a number in YYYYMMDDHHMMSS notation; 20000530144500 denotes 14:45:00 UTC on May 30th, 2000. A relative start time is indicated by +<i>N</i>, which is <i>N</i> seconds from the current time. If no <i>start-time</i> is specified, the current time is used. -v <i>level</i> Set the debugging level.
操作数	The following operands are supported: <i>key</i> The keys used to sign the child's keyset. <i>keyset</i> The file containing the child's keyset.

示例

示例 1 Sign the keyset file for example.com.

The DNS administrator for a DNSSEC-aware .com zone would use the following command to sign the *keyset* file for example.com created by `dnssec-makekeyset` with a key generated by `dnssec-keygen`:

```
dnssec-signkey keyset-example.com. Kcom.+003+51944
```

In this example, `dnssec-signkey` creates the file `signedkey-example.com`, which contains the example.com keys and the signatures by the .com keys.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWbind9
Interface Stability	Volatile

另请参见

[dnssec-keygen\(1M\)](#), [dnssec-makekeyset\(1M\)](#), [dnssec-signzone\(1M\)](#), [attributes\(5\)](#)

附注

Source for BIND9 is available in the SUNWbind9S package.

引用名	dnssec-signzone – DNSSEC 区域签名工具
用法概要	<pre> dnssec-signzone [-Aaghtz] [-c class] [-d directory] [-e end-time] [-f output-file] [-H iterations] [-I input-format] [-i interval] [-k key] [-l domain] [-N soa-serial-format] [-n ncpus] [-O output-format] [-o origin] [-r randomdev] [-s start-time] [-v level] [-3 salt] zonefile [key]... </pre>
描述	<code>dnssec-signzone</code> 实用程序可以对区域签名。它生成 NSEC 和 RRSIG 记录并生成已签名版本的区域。已签名区域中委托的安全状态（即子区域是否安全）取决于每个子区域是否存在 keyset 文件。
选项	所支持的选项如下： -A 生成 NSEC3 链时，在所有 NSEC3 记录上设置 OPTOUT 标志且不会为不安全的委托生成 NSEC3 记录。 -a 验证所有生成的签名。 -c class 指定区域的 DNS 类。 -d directory 在 <i>directory</i> 中查找 keyset 文件。 -e end-time 指定生成的 RRSIG 记录过期的日期和时间。与 <i>start-time</i> 一样，绝对时间采用 YYYYMMDDHHMMSS 表示法表示。相对于开始时间的时间用 +N 表示，即开始时间后的 N 秒。相对于当前时间的时间用 now+N 表示。如果没有指定 <i>end-time</i>，则缺省值是开始时间后的第 30 天。 -f output-file 包含已签名区域的输出文件的名称。缺省情况是将 .signed 附加到输入文件名称后。 -g 为 keyset 文件中的子区域生成 DS 记录。将删除现有 DS 记录。 -H iterations 生成 NSEC3 链时，使用 <i>iterations</i> 指定的迭代次数。缺省为 100。 -h 列出 <code>dnssec-signzone()</code> 的选项和参数的简短摘要。 -I input-format 输入区域文件的格式。可能的格式是 <code>text</code>（缺省）和 <code>raw</code>。此选项主要用于动态的已签名区域，以便可以直接对包含更新的非文本格式的转储区域文件签名。使用此选项对非动态区域毫无作用。

-i interval

将循环间隔指定为当前时间的偏移（以秒为单位）。之前已签名的区域作为输入进行传递时，可以重新对记录签名。如果 RRSIG 记录在循环间隔后过期，则保留该记录。否则，该记录被视为很快将过期并将被替换。

缺省循环间隔是签名结束时间和开始时间之差的四分之一。如果没有指定 *end-time* 和 *start-time*，则 `dnssec-signzone` 将生成有效期为 30 天、循环间隔为 7.5 天的签名。由于任何现有 RRSIG 记录在 7.5 天之内过期，因此会将其替换。

-j jitter

使用固定签名生存期对区域签名时，签名时发出的所有 RRSIG 记录均同时到期。如果以增量方式对区域签名（即之前签名的区域作为输入传递给签名者），需要大约同时重新生成所有过期的签名。*jitter* 选项指定用于随机指定签名过期时间的 *jitter* 窗口，从而随时间以一定增量分散重新生成签名的时间。

在某种程度上，签名生命期 *jitter* 也通过分散高速缓存失效时间使验证器和服务器受益。也就是说，如果所有高速缓存的大量 RRSIG 不同时到期，则相对于所有验证器需要几乎同时重取来说，前者出现拥塞的可能性更小。

-k key

将指定的 *key* 视作密钥签名密钥，忽略任何密钥标志。可以多次指定此选项。

-l domain

除密钥 (DNSKEY) 和 DS 集以外，还生成 DLV 集。域附加到记录名称。

-N soa-serial-format

已签名区域的 SOA 序列号格式。可能的格式是 *keep*（缺省）、*increment* 和 *unixtime*，如下所述。

keep

不修改 SOA 序列号。

increment

使用 RFC 1982 运算递增 SOA 序列号。

unixtime

将 SOA 序列号设置为自 Unix 时间戳起过去的秒数。

-n nthreads

指定要使用的线程数。缺省情况下，为每个检测出的 CPU 启动一个线程。

-O output_format

包含已签名区域的输出文件的格式。可能的格式是 *text*（缺省）和 *raw*。

-o origin

指定区域起始点。如果没有指定，则将区域文件的名称假定起始点。

-p

签名区域时使用伪随机数据。这比使用实际随机数据更快，但更不安全。签名大型区域或熵源有限时，此选项可能会有用。

-r randomdev

指定随机源。如果操作系统不提供 `/dev/random` 或等效设备，则缺省的随机源是键盘输入。`randomdev` 指定字符设备的名称或包含要使用的随机数据的文件（而非缺省的 `/dev/random`）。特殊值 `keyboard` 表示应该使用键盘输入。

-s start-time

指定生成的 RRSIG 记录开始有效的日期和时间。此时间可以是绝对时间，也可以是相对时间。绝对开始时间由采用 YYYYMMDDHHMMSS 表示法的数字表示；20000530144500 表示 2000 年 5 月 30 日 14:45:00 (UTC)。相对开始时间由 +N 表示，即当前时间后的 N 秒。如果没有指定 `start-time`，则使用当前时间减去一小时（以允许时钟相位差）后的时间。

-t

完成时列出统计信息。

-v level

设置调试级别。

-z

确定签名内容时，忽略密钥上的 KSK 标志。

-3 salt

使用指定的十六进制编码 `salt` 生成 NSEC3 链。破折号（-）可以用于表示生成 NSEC3 链时没有使用 `salt`。

操作数

支持下列操作数：

zonefile

包含要签名的区域的文件。

key

指定应在对区域签名时使用的密钥。如果没有指定密钥，则检查该区域以获得区域顶点的 DNSKEY 记录。如果找到这些记录且当前目录中有匹配的私钥，则这些记录将用于进行域名。

示例

示例 1 使用 DSA 密钥对区域签名

以下命令使用 [dnssec-keygen\(1M\)](#) 手册页中的示例中生成的 DSA 密钥对 `example.com` 区域签名（该示例为 `Kexample.com.+003+17247`）。区域的密钥必须在主文件（`db.example.com`）中。该调用在当前目录中查找密钥集文件，以便从这些文件生成 DS 记录（-g）。

```
% dnssec-signzone -g -o example.com db.example.com \
Kexample.com.+003+17247
db.example.com.signed
%
```

在上述示例中，`dnssec-signzone` 创建文件 `db.example.com.signed`。应该在 `named.conf` 文件的区域语句中引用此文件。

示例2 重新签名之前已签名的区域

以下命令使用缺省参数对以前已签名的区域重新签名。假定私钥位于当前目录中。

```
% cp db.example.com.signed db.example.com
% dnssec-signzone -o example.com db.example.com \
db.example.com.signed
%
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/network/dns/bind
接口稳定性	Volatile（可变）

另请参见 [dnssec-keygen\(1M\)](#)、[attributes\(5\)](#)

RFC4033

请参见 BIND 9 Administrator's Reference Manual。从本手册页发布之日起，将在 <https://www.isc.org/software/bind/documentation> 上提供该文档。

引用名 domainname – set or display name of the current domain

用法概要 domainname [*name-of-domain*]

描述 Without an argument, domainname displays the name of the current domain name used in RPC exchanges, usually referred to as the NIS domain name. This name typically encompasses a group of hosts or passwd entries under the same administration. The domainname command is used by various components of Solaris to resolve names for entries such as are found in passwd, hosts and aliases. By default, naming services such as NIS use domainname to resolve names.

With solaris.smf.manage.domainname [rbac\(5\)](#) privileges, you can set the name of the domain by specifying the domain name as the value of the config/domainname property of the nis/domain service with a sequence of commands such as the following:

```
% svccfg -s nis/domain setprop config/domainname = \  
hostname: nisdomain.example.com  
% svccfg -s nis/domain:default refresh  
% svcadm enable nis/domain
```

In addition to setting the value of config/domainname, the preceding sequence enables the nis/domain service. This service is not enabled by default. The nis/domain service sets up all the shared NIS configuration, of which domainname is the one mandatory configuration parameter.

The domain name for various naming services can also be set by other means. For example, ypinit can be used to specify a different domain for all NIS calls. The domain name of the machine is usually set during boot time through the domainname command by the svc:/system/identity:domain service. If the new domain name is not saved in the /etc/defaultdomain file, the machine reverts to the old domain after it reboots.

The [sendmail\(1M\)](#) daemon, as shipped with Solaris, and the sendmail implementation provided by sendmail.org (formerly referred to as “Berkeley 8.x sendmail”) both attempt to determine a local host's fully qualified host name at startup and both pursue follow-up actions if the initial search fails. It is in these follow-up actions that the two implementations differ.

Both implementations use a standard Solaris or Unix system call to determine its fully qualified host name at startup, following the name service priorities specified in [nsswitch.conf\(4\)](#). To this point, the Solaris and sendmail.org versions behave identically.

If the request for a fully qualified host name fails, the sendmail.org sendmail sleeps for 60 seconds, tries again, and, upon continuing failure, resorts to a short name. The Solaris version of sendmail makes the same initial request, but then, following initial failure, calls domainname. If successful, the sleep is avoided.

On a Solaris machine, if you run the sendmail.org version of sendmail, you get the startup behavior (omitting the domainname call) described above. If you run the Solaris sendmail, the domainname call is made if needed.

If the Solaris `sendmail` cannot determine the fully qualified host name, use [check-hostname\(1M\)](#) as a troubleshooting aid. This script can offer guidance as to appropriate corrective action.

文件

- `/etc/defaultdomain`
- `/etc/nsswitch.conf`

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	system/core-os

另请参见

[svcs\(1\)](#), [check-hostname\(1M\)](#), [hostconfig\(1M\)](#), [named\(1M\)](#), [sendmail\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [ypinit\(1M\)](#), [aliases\(4\)](#), [defaultdomain\(4\)](#), [hosts\(4\)](#), [nsswitch.conf\(4\)](#), [passwd\(4\)](#), [attributes\(5\)](#), [rbac\(5\)](#), [smf\(5\)](#)

附注

The `domainname` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/system/identity:domain
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名 drd – Logical Domain Dynamic Reconfiguration daemon

用法概要 /usr/lib/ldoms/drd

描述 The drd daemon is part of the framework that enables the addition and removal of resources from a Logical Domain. This framework is collectively called Dynamic Reconfiguration (DR).

drd is responsible for various aspects of DR on a Logical Domain and must be enabled to ensure proper DR functionality. It is started at boot time and has no configuration options.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/ldoms
Interface Stability	Uncommitted

另请参见 [svcs\(1\)](#), [svcadm\(1M\)](#), [syslog\(3C\)](#), [syslog.conf\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

错误 drd uses [syslog\(3C\)](#) to report status and error messages. All of the messages are logged with the LOG_DAEMON facility. Error messages are logged with the LOG_ERR and LOG_NOTICE priorities, and informational messages are logged with the LOG_INFO priority. The default entries in the /etc/syslog.conf file log all the drd error messages to the /var/adm/messages log.

附注 The drd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/platform/sun4v/drd:default

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	drvconfig – apply permission and ownership changes to devices
用法概要	<pre>drvconfig [-bn] [-a <i>alias_name</i>] [-c <i>class_name</i>] [-i <i>drivername</i>] [-m <i>major_num</i>] [-r <i>root_dir</i>]</pre>
描述	devfsadm(1M) is now the preferred command and should be used instead of <code>drvconfig</code>. The default operation of <code>drvconfig</code> is to apply permission and ownership changes to devices. Normally, this command is run automatically after a new driver has been installed (with add_drv(1M)) and the system has been rebooted.
选项	The following options are supported: -a<i>alias_name</i> Add the name <i>alias_name</i> to the list of aliases that this driver is known by. This option, if used, must be used with the -m <i>major_num</i>, the -b and the -i <i>drivername</i> options. -b Add a new major number to name binding into the kernel's internal name_to_major tables. This option is not normally used directly, but is used by other utilities such as add_drv(1M). Use of the -b option requires that -i and -m be used also. No /devices entries are created. -c<i>class_name</i> The driver being added to the system exports the class <i>class_name</i>. This option is not normally used directly, but is used by other utilities. It is only effective when used with the -b option. -i<i>drivername</i> Only configure the devices for the named driver. The following options are used by the implementation of add_drv(1M) and rem_drv(1M), and may not be supported in future versions of Solaris: -m<i>major_num</i> Specify the major number <i>major_num</i> for this driver to add to the kernel's name_to_major binding tables. -n Do not try to load and attach any drivers, or if the -i option is given, do not try to attach the driver named <i>drivername</i>. -r<i>root_dir</i> Perform operations under <i>root_dir</i>, rather than directly under root.
退出状态	0 Successful completion. non-zero An error occurred.
文件	/devices Device nodes directory /etc/minor_perm Minor mode permissions /etc/name_to_major Major number binding /etc/driver_classes Driver class binding file

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[sh\(1\)](#), [add_drv\(1M\)](#), [modinfo\(1M\)](#), [modload\(1M\)](#), [modunload\(1M\)](#), [rem_drv\(1M\)](#), [update_drv\(1M\)](#), [path_to_inst\(4\)](#), [attributes\(5\)](#), [devfs\(7FS\)](#)

引用名	dsbitmap – 计算 Availability Suite 位图卷的大小														
用法概要	<pre>dsbitmap -h dsbitmap -p data_volume [bitmap_volume] dsbitmap -r data_volume [bitmap_volume]</pre>														
描述	dsbitmap 命令计算需要与指定数据卷一起使用的 Availability Suite 位图卷的大小。														
选项	所支持的选项如下： -h 列出 dsbitmap 命令的用法消息 -p <i>data_volume</i> [<i>bitmap_volume</i>] 对于给定的 <i>data_volume</i>，dsbitmap 计算并显示相关 Availability Suite Point in Time 位图卷的所需大小。显示所有可能的 Availability Suite Point in Time 集配置的位图卷大小。这些配置包括独立投影、完整大小依赖投影和紧凑型依赖投影。如果提供可选的 <i>bitmap_volume</i> 参数，dsbitmap 将确定该卷是否足够大，以用作 <i>data_volume</i> 的位图卷。 -r <i>data_volume</i> [<i>bitmap_volume</i>] 对于给定的 <i>data_volume</i>，dsbitmap 计算并显示相关 Availability Suite Remote Mirror 位图卷的所需大小。显示所有可能的 Availability Suite Remote Mirror 集配置的位图卷大小。这些配置包括与内存队列、磁盘队列和 32 位引用计数的同步复制和异步复制。如果提供可选的 <i>bitmap_volume</i> 参数，dsbitmap 将确定该卷是否足够大，以用作 <i>data_volume</i> 的位图卷。														
用法	通常，系统管理员在配置 Sun StorageTek Availability Suite 软件的初始阶段使用 dsbitmap，以便确定所需的位图卷大小，然后检查已创建的位图卷是否合适。														
退出状态	将返回下列退出值： <table><tr><td>0</td><td>成功完成。如果指定了位图卷的名称，则该卷足够大以供所有潜在用途。</td></tr><tr><td>1</td><td>出现错误。</td></tr><tr><td>2</td><td>命令行上提供了一个无效选项。</td></tr><tr><td>3</td><td>如果指定的位图卷要作为 Availability Suite Remote Mirror 位图来用于磁盘队列的异步集，则不够大，但是如果要用于所有其他 Remote Mirror 集配置，则足够大。</td></tr><tr><td>4</td><td>指定的位图卷不够大，不能作为 Availability Suite Remote Mirror 位图用于任何 Remote Mirror 集配置。</td></tr><tr><td>5</td><td>指定的位图卷不够大，不能作为 Availability Suite Remote Mirror 位图用于任何 Remote Mirror 集配置。</td></tr><tr><td>6</td><td>如果指定的位图卷要作为 Availability Suite Point in Time 位图来用于紧凑型依赖投影，则不够大，但是如果要用于所有其他 Point in Time 集配</td></tr></table>	0	成功完成。如果指定了位图卷的名称，则该卷足够大以供所有潜在用途。	1	出现错误。	2	命令行上提供了一个无效选项。	3	如果指定的位图卷要作为 Availability Suite Remote Mirror 位图来用于磁盘队列的异步集，则不够大，但是如果要用于所有其他 Remote Mirror 集配置，则足够大。	4	指定的位图卷不够大，不能作为 Availability Suite Remote Mirror 位图用于任何 Remote Mirror 集配置。	5	指定的位图卷不够大，不能作为 Availability Suite Remote Mirror 位图用于任何 Remote Mirror 集配置。	6	如果指定的位图卷要作为 Availability Suite Point in Time 位图来用于紧凑型依赖投影，则不够大，但是如果要用于所有其他 Point in Time 集配
0	成功完成。如果指定了位图卷的名称，则该卷足够大以供所有潜在用途。														
1	出现错误。														
2	命令行上提供了一个无效选项。														
3	如果指定的位图卷要作为 Availability Suite Remote Mirror 位图来用于磁盘队列的异步集，则不够大，但是如果要用于所有其他 Remote Mirror 集配置，则足够大。														
4	指定的位图卷不够大，不能作为 Availability Suite Remote Mirror 位图用于任何 Remote Mirror 集配置。														
5	指定的位图卷不够大，不能作为 Availability Suite Remote Mirror 位图用于任何 Remote Mirror 集配置。														
6	如果指定的位图卷要作为 Availability Suite Point in Time 位图来用于紧凑型依赖投影，则不够大，但是如果要用于所有其他 Point in Time 集配														

置，则足够大。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	storage/avs/avs-cache-management
接口稳定性	Committed（已确定）

另请参见 [iiadm\(1M\)](#)、[sndradm\(1M\)](#)、[attributes\(5\)](#)

- C group -i

初始化仅与指定组相关的配置数据库条目。此选项删除之前或当前的配置信息时，将在继续操作之前提示您确认该删除操作。
- C group -p parser_config_file

与 -i 选项一起使用时，dscfg 将解析器配置文件装入持久性配置，它将重新格式化仅与指定组相关的配置数据库条目。
- C group -a config_file

将指定的 config_file 恢复到仅与指定组相关的配置数据库条目。此选项不会对该文件执行任何错误检查。使用此选项将使配置文件无效。
- C group -l

列出与指定资源组相关的当前配置数据库的内容。
- C group -s config_file_location

指定 Sun Cluster 配置数据库所驻留的位置，该位置应该是 DID 设备。
- D device_group

检查指定的 device_group 是否被 Sun Cluster 已知以及在该节点上是否可用。它显示表明设备组状态的信息并返回以下值，

(使用回显 \$? 检索返回值)：

1

设备组位于 Sun Cluster 中，且在该节点上处于活动状态；

0

设备组位于 Sun Cluster 中，但在其他节点上处于活动状态；

-1

设备组不位于 Sun Cluster 中。

用法 dscfg 命令通常从其他脚本（如软件包安装脚本）和 dscfgadm(1M) 运行。它不适用于常规使用。对于 dscfgadm 不包含的 dscfg 的用法，请参见《Availability Suite Troubleshooting Guide》。

文件 /etc/dscfg_format
解析器配置文件

/etc/dscfg_local
没有群集标记的所有条目的配置位置

/etc/dscfg_cluster
指定 Sun Cluster 配置数据库的位置的引用文件。

属性 有关以下属性的说明，请参见 attributes(5)：

属性类型	属性值
可用性	storage/avs/avs-cache-management

属性类型	属性值
接口稳定性	Committed（已确定）

另请参见

[dscfgadm\(1M\)](#)、[iiadm\(1M\)](#)、[scmadm\(1M\)](#)、[sndradm\(1M\)](#)、[svadm\(1M\)](#)、[ds.log\(4\)](#)、[attributes](#)

引用名	dscfgadm – Availability Suite 配置位置和服务的管理命令
用法概要	<pre>dscfgadm -i dscfgadm -e [-r] [-p] dscfgadm -d [-r] dscfgadm -s</pre>
描述	dscfgadm 命令通过提供用于设置配置位置并启用和禁用 Availability Suite 服务的工具，来控制 Availability Suite 配置位置和服务。
选项	如果没有为 dscfgadm 命令指定任何参数，实用程序将交互式逐步完成设置配置位置并启动服务的过程。初始化配置位置以进行使用之前，将验证该位置以确保其满足大小和文件类型等标准。 dscfgadm 命令支持以下选项。 -d [-r] 禁用 Availability Suite SMF 服务。包括停止守护进程和暂停当前在 Availability Suite 软件下启用的任何 Point-in-Time Copy 或 Remote Mirror 集。不与其他选项一起执行时，-d 选项将禁用所有 Availability Suite 服务。系统引导过程中一直保持该设置。 -r 传递给 -d 或 -e 选项时，仅对 Remote Mirror 服务起作用。 -e [-r] [-p] 启用 Availability Suite SMF 服务。包括启动守护进程和恢复之前已在 Availability Suite 软件下配置的任何 Point-in-Time Copy 或 Remote Mirror 集。不与其他选项一起执行时，-e 选项将启用所有 Availability Suite 服务。系统引导过程中一直保持该设置。 -r 传递给 -d 或 -e 选项时，仅对 Remote Mirror 服务起作用。 -p 传递给 -e 选项时，仅启用 Point-in-Time Copy 服务。 -i 显示有关 Availability Suite SMF 服务的信息。 -s 允许您设置包含特定于 Sun Cluster 配置的信息的配置数据库的位置。 -x 执行程序时，显示详细调试信息。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	storage/avs/avs-cache-management
接口稳定性	Committed（已确定）

另请参见 [dscfg\(1M\)](#)、[iiadm\(1M\)](#)、[scmadm\(1M\)](#)、[sndradm\(1M\)](#)、[svadm\(1M\)](#)、[attributes\(5\)](#)

引用名	dscfglockd – Availability Suite 服务锁定守护进程
用法概要	<code>/usr/lib/dscfglockd [-e <i>program</i> -f <i>file</i>]</code>
描述	dscfglockd 守护进程在群集内的节点之间协调 StorageTek 配置数据库锁定请求。
选项	dscfglockd 支持以下选项： -e <i>program</i> 执行 <i>program</i> 中的脚本和参数。所需的可执行文件和任何参数都必须通过适当引用为单个参数来提供给此选项。该参数由 sh(1) 进行处理。 -f <i>file</i> 从 <i>file</i> 中读取对等主机名称列表。 如果没有指定参数，dscfglockd 将充当本地锁定守护进程，但是将与可能联系本地主机的其他守护进程（作为其自己的配置过程的一部分）协调锁定请求。
退出状态	0 守护进程启动成功。 >0 守护进程启动失败。
文件	<code>/lib/svc/method/svc-scm</code> 用于启动和停止 dscfglockd 的 Shell 脚本。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	storage/avs/avs-cache-management
接口稳定性	Committed（已确定）

另请参见 [sh\(1\)](#)、[attributes\(5\)](#)

引用名	dsstat – 报告 Sun StorageTek Availability Suite I/O 统计信息
用法概要	<pre>dsstat -m mode [-r report-options] [-d display-options] [-s volume-sets] [-f -F] [-z] [interval [count]]</pre>
描述	dsstat 命令收集并报告 Sun StorageTek Availability Suite 产品的 I/O 统计信息。
选项	dsstat 支持以下选项： -m mode 指定运行模式。有效的模式为： cache 高速缓存统计信息 ii Point-in-Time Copy 统计信息 sndr Remote Mirror 统计信息 要同时显示多个服务的统计信息，请使用 -m 开关指定所需的每个模式。例如： <pre>% dsstat -m ii -m sndr % dsstat -m ii -m sndr -m cache</pre> 要确定哪些统计信息报告自哪个服务，请使用如下所述的 <i>role</i> 字段。当作为多个服务之一请求高速缓存统计信息时，<i>rkps</i> 和 <i>wkps</i> 统计信息将进一步分为 <i>crkps</i>、<i>drkps</i>、<i>cwkps</i> 和 <i>dwkps</i>。如果未指定 -m 开关，则 dsstat 将显示上述所有有效模式的缺省统计信息。请参见下面的“字段说明”了解对这些字段的描述。 -r report-options 指定要显示的卷组件。每一项由单个字符表示，可以选择多个项。<i>report-options</i> 根据上面指定的模式而不同。该选项不用于 cache 模式。 ii 模式下的有效 <i>report-options</i> 为： m 主卷统计信息。 ■ 对于“ii 模式下的 <i>report-options</i>”：m、s、b、o ■ 对于“sndr 模式下的 <i>report-options</i>”：b、n ■ 对于“cache 模式下的 <i>display-options</i>”：r、w、d、c、s、f ■ 对于“ii 模式下的 <i>display-options</i>”：r、w、t、s、p、f s 影子卷统计信息。 b 位图卷统计信息。 o 溢出卷统计信息（如有连接）。 在 sndr 模式下，有效的 <i>report-options</i> 为： b 位图卷统计信息。

n 网络卷统计信息。

-d *display-options*

指定要显示的统计信息。统计信息类型由单个字符表示；可以指定多个类型。

在 *cache* 模式下，有效的 *display-options* 为：

r 详细的读取统计信息。

w 详细的写入统计信息。

s 摘要统计信息。

f 高速缓存行为标志。

以下 *display-options* 只适用于 *cache* 模式，并且需要与模式选项 (-m) 结合使用

d 降级数据统计信息。

c 写入取消统计信息。

cache 的缺省 *display-options* 是 *sf*。

在 *ii* 模式下，有效的 *display-options* 为：

r 详细的读取统计信息。

w 详细的写入统计信息。

t 计时统计信息。

s 摘要统计信息。

p 要求同步的卷的百分比。

f 卷类型/状态标志。

ii 的缺省 *display-options* 是 *spf*。在 *sndr* 模式下，有效的 *display-options* 为：

r 详细的读取统计信息。

w 详细的写入统计信息。

t 计时统计信息。

s 摘要统计信息。

p 要求同步的卷的百分比。

- f 卷类型/状态标志。
- q 异步队列统计信息。

以下 *display-option* 只适用于 *sndr* 模式，并且需要与模式选项 (-m) 结合使用。

- q 异步队列统计信息。

sndr 的缺省 *display-options* 是 *spf*。

指定摘要 *display-option* 将导致 *rw t display-options* 被忽略。

- s *volume-sets* 过滤输出以只包含指定的 *volume-sets*，其中 *volume-sets* 是以逗号分隔的卷名列表。显示 Remote Mirror 统计信息时，指定的名称将与 Remote Mirror 主卷进行比较，以确定其是否匹配。此外，Remote Mirror 卷应指定为 <host>:<volume>，其中 <host> 是主要主机或辅助主机，<volume> 是该主机用以识别卷集的卷名。指定多个 Remote Mirror 卷集时，请注意指定的各个卷集将使用上述规则分别进行评估。显示 Point-in-Time Copy 统计信息时，指定的卷名将与 Point-in-Time Copy 影子卷进行比较，以确定其是否匹配。
- f 每个报告周期输出字段标题。
- F 报告开始时，输出一次字段标题。
- z 取消显示具有零值或没有活动的报告行。

操作数

dsstat 命令行支持以下操作数：

interval
为各个报告指定 *interval*（以秒为单位）。如果未指定 *interval*，单个报告将以一秒的 *interval* 输出。

由于给定集中的条目数不断变化以及集数量不断变化，标题信息可能会显示在要报告的集的中部。为了避免此情况，可使用 -f 或 -F 选项在所需时间显示标题信息。

count
指定要生成的报告数。如果未指定 *count*，输出将继续直至中断。

字段说明

除非另有说明，否则所有字段都是基于指定 *interval* 期间收集数据的每秒平均数。

- name 报告的实体的名称
- t 卷类型
- s 卷状态
- r 高速缓存读取行为
- w 高速缓存写入行为

pct	要求同步的卷的百分比
role	报告的项的角色
tps	读取 + 写入的总数
rtps	读取数
wtps	写入数
kps	读取 + 写入的总千字节数
rkps	读取的千字节数
wkps	写入的千字节数
crkps	从高速缓存读取的千字节数
drkps	从磁盘读取的千字节数
cwkps	写入高速缓存的千字节数
dwkps	写入磁盘的千字节数
ckps	高速缓存传入传出的千字节数
dkps	磁盘传入传出的千字节数
svt	每次运行的服务时间
hit	<i>interval</i> 期间的读取命中次数
ds/s	从高速缓存降级的千字节数
cn/s	写入取消数
q	使用的异步队列类型
qi	当前已排队的项数
qk	当前已排队的千字节数
qhwi	已排队项的高界限
qhwk	已排队千字节数的高界限

name 字段只显示名称的最后 16 个字符。

cache 的有效高速缓存行为是：

C	高速缓存读取/写入
D	磁盘读取/写入

ii 的有效卷类型是：

I	独立影子卷

D 依赖影子卷

ii 的有效卷状态是：

C 复制正在进行

- 无复制正在进行

sndr 的有效卷类型是：

P 这是此卷集的主要主机

S 这是此卷集的辅助主机

sndr 的有效卷状态是：

L 正在记录对此卷的更改

Q 卷处于队列模式下

R 正在将更改复制到辅助主机

SY 正在进行同步。正在发送数据。

RS 正在进行反向同步。正在接收数据。

SN 需要同步

RN 需要反向同步

VF 卷失败

BF 位图失败

QF 队列失败

sndr 的有效队列类型是：

D 已启用基于磁盘的队列

M 已启用基于内存的队列

sndr 的卷角色是：

net 网络卷统计信息

bmp 位图卷统计信息

ii 的卷角色是：

mst 主卷统计信息

shd 影子卷统计信息

bmp 位图卷统计信息

ovr 溢出卷统计信息

示例

示例1 报告高速缓存统计信息

以下示例显示报告高速缓存统计信息，其中详细划分了对高速缓存/磁盘的读取和写入。报告生成的时间间隔为五秒。报告限制到集 /dev/rdisk/clt1d0s0。

```
# dsstat -m cache -d rw -s /dev/rdisk/clt1d0s0 5
- read -           - write -
name              ckps   dkps   hit   ckps   dkps   hit
dev/rdisk/clt1d0s0 0      0     0.00   0      0     0.00
dev/rdisk/clt1d0s0 3    2396  0.13   983    763  100.00
dev/rdisk/clt1d0s0 2399   799  75.00  2815   2686 100.00
dev/rdisk/clt1d0s0 3200   800  80.00  2755   2908 100.00
dev/rdisk/clt1d0s0 3999   799  83.33  2809   2868 100.00
dev/rdisk/clt1d0s0 4800   800  85.71  2867   2931 100.00
```

示例2 报告主卷、影子卷和位图卷的统计信息

使用缺省输出报告 Point-in-Time Copy 主卷、影子卷和位图卷的统计信息。报告生成的时间间隔为两秒。

```
# dsstat -m ii -r msb 2
name              t  s    pct role   kps   tps   svt
dev/rdisk/c0t1d0s5 I  C  96.15 mst  19921  38   22
dev/rdisk/c0t1d0s6          shd  9960  19   20
dev/rdisk/c0t1d0s7          bmp   39   77   2
dev/rdisk/c0t1d0s5 I  C  94.24 mst  19623  38   22
dev/rdisk/c0t1d0s6          shd  9939  19   20
dev/rdisk/c0t1d0s7          bmp   39   77   2
dev/rdisk/c0t1d0s5 I  C  92.34 mst  19969  39   22
dev/rdisk/c0t1d0s6          shd  9984  19   20
dev/rdisk/c0t1d0s7          bmp   39   78   2
```

示例3 报告 Remote Mirror 网络统计信息

使用详细的读取、写入统计信息报告 Remote Mirror 的网络统计信息。报告包括卷类型/状态标志和百分比。报告生成的时间间隔为两秒。将报告限制到集 /dev/rdisk/c0t1d0s0。

```
# dsstat -m sndr -r n -d rwpf -s /dev/rdisk/c0t1d0s0 2
name              t  s    pct role   rkps  rtps  wkps  wtps
dev/rdisk/c0t1d0s0 P  L 100.00 sec    0    0    0    0
dev/rdisk/c0t1d0s0 P  SY 99.90 sec    0    0   288    9
dev/rdisk/c0t1d0s0 P  SY 97.90 sec    0    0  5296  165
dev/rdisk/c0t1d0s0 P  SY 95.81 sec    0    0  5184  161
dev/rdisk/c0t1d0s0 P  SY 93.81 sec    0    0  5280  164
dev/rdisk/c0t1d0s0 P  SY 91.71 sec    0    0  5198  162
```

- 退出状态
- 将返回下列退出值：
- 0

成功完成。
- 1

成功完成，无统计信息需要报告。
- 2

遇到了无效参数。
- 3

无内存可用于创建 `kstat` 统计信息。
- 4

发生未知错误。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	storage/avs/avs-cache-management
接口稳定性	Committed（已确定）

另请参见

[dscfg\(1M\)](#)、[svadm\(1M\)](#)、[ds.log\(4\)](#)、[rdc.cf\(4\)](#)、[attributes\(5\)](#)

引用名 dsvclockd – DHCP service lock daemon

用法概要 /usr/lib/inet/dsvclockd [-d 1 | 2] [-f] [-v]

描述 The dsvclockd daemon is a lock manager that works in conjunction with the Dynamic Host Configuration Protocol (DHCP) Data Service Library (libdhcpsvc). It provides shared or exclusive access to the [dhcp_network\(4\)](#) and [dhcptab\(4\)](#) tables. This service is used by the SUNWbinfiles and SUNWfiles DHCP data store modules. See [dhcp_modules\(5\)](#).

dsvclockd is started on demand by libdhcpsvc. The dsvclockd daemon should be started manually only if command line options need to be specified.

- 选项** The following options are supported:
- d 1 | 2 Set debug level. Two levels of debugging are currently available, 1 and 2. Level 2 is more verbose.
 - f Run in the foreground instead of as a daemon process. When this option is used, messages are sent to standard error instead of to [syslog\(3C\)](#).
 - v Provide verbose output.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/dhcp
Interface Stability	Uncommitted

另请参见 [syslog\(3C\)](#), [dhcp_network\(4\)](#), [dhcptab\(4\)](#), [dhcp_modules\(5\)](#), [attributes\(5\)](#)

引用名

dtrace – DTrace 动态跟踪编译器和跟踪实用程序

用法概要

```
dtrace [-32 | -64] [-aACeFGHhLqSvVwZ] [-b bufsz] [-c cmd]
      [-D name [=value]] [-I path] [-L path] [-o output]
      [-s script] [-U name] [-x arg [=val]]
      [-X a | c | s | t] [-p pid]
      [-P provider [[predicate] action]]
      [-m [provider:] module [[predicate] action]]
      [-f [[provider:] module:] function [[predicate] action]]
      [-n [[provider:] module:] function:] name [[predicate] action]]
      [-i probe-id [[predicate] action]]
```

描述

DTrace 是一个用于 Solaris 操作系统的全面的动态跟踪框架。DTrace 提供一个强大的基础结构，使管理员、开发者和服务人员能简明地回答有关操作系统和用户程序行为的任意问题。

《[Solaris Dynamic Tracing Guide](#)》介绍如何使用 DTrace 观察、调试和调优系统行为。请参见此书以详细了解 DTrace 功能，包括随附的 DTrace 监测工具、检测过程提供器以及 D 编程语言。

dtrace 命令提供了一个可访问 DTrace 工具提供的基本服务的通用界面，包括：

- 用于列出 DTrace 当前发布的探测器和提供器集的选项
- 允许使用任何探测器说明符（提供器、模块、函数和名称）直接启用探测器的选项
- 用于运行 D 编译器并编译一个或多个直接在命令行上编写的 D 程序文件或程序的选项
- 用于生成匿名跟踪程序的选项
- 用于生成程序稳定性报告的选项
- 用于修改 DTrace 跟踪和缓冲行为并启用其他 D 编译器功能的选项

可以使用 dtrace 创建 D 脚本，方式是在 `#!` 声明中使用它来创建解释程序文件。还可在不使用 `-e` 选项实际启用跟踪的情况下，使用 dtrace 尝试编译 D 程序并确定其属性。请参见“选项”部分。请参见《[Solaris Dynamic Tracing Guide](#)》，获得如何使用 dtrace 实用程序执行这些任务的详细示例。

选项

`-P`、`-m`、`-f`、`-n` 以及 `-i` 选项接受的参数可包括括在斜杠 `//` 中的可选 D 语言 *predicate* 与括在大括号 `{}` 中的 D 语言 *action* 语句列表。在命令行上指定的 D 程序代码必须适当用引号括起，以避免 Shell 进行元字符解释。

所支持的选项如下：

`-32` | `-64`

D 编译器使用操作系统内核的本机数据模型生成程序。可使用 `isainfo -b` 命令确定当前操作系统的数据模型。如果指定了 `-32` 选项，则 dtrace 将强制 D 编译器使用 32 位数据模型编译 D 程序。如果指定了 `-64` 选项，则 dtrace 将强制 D 编译器使用 64 位数据模型编译 D 程序。通常这些选项不是必需的，因为 dtrace 会选择本机数据模型作为缺省数据模型。数据模型会影响整数类型的大小和其他语言属性。对

任一数据模型编译的 D 程序可在 32 位和 64 位内核上执行。-32 和 -64 选项还可确定 -G 选项生成的 ELF 文件格式（ELF32 或 ELF64）。

-a

声明匿名跟踪状态并显示跟踪数据。可将 -a 选项与 -e 选项组合在一起，以强制 dtrace 在使用匿名跟踪状态后立即退出，而不是继续等待新数据。有关匿名跟踪的更多信息，请参见《Solaris Dynamic Tracing Guide》。

-A

生成 `driver.conf(4)` 指令以进行匿名跟踪。该选项将构造一组 `dtrace(7D)` 配置文件指令，为匿名跟踪启用指定探测器，然后退出。缺省情况下，dtrace 尝试将指令存储到文件 `/kernel/drv/dtrace.conf` 中。如果使用 -o 选项指定备用输出文件，则可修改此行为。

-b *bufsz*

设置主体跟踪缓冲区大小 (*bufsz*)。跟踪缓冲区大小可以使用任何大小后缀 *k*、*m*、*g* 或 *t*。如果不能分配缓冲区空间，则 dtrace 将根据 *bufsize* 属性的设置减小缓冲区大小或退出。

-c *cmd*

运行指定的命令 *cmd*，并在完成时退出。如果命令行上有多个 -c 选项，dtrace 将在所有命令退出后退出，并报告每个子进程终止时的退出状态。第一个命令的进程 ID 可供在命令行上（或者使用 -s 选项通过 *\$target* 宏变量）指定的任何 D 程序使用。有关宏变量的更多信息，请参见《Solaris Dynamic Tracing Guide》。

-C

在编译 D 程序之前对它们运行 C 预处理程序 `cpp(1)`。可使用 -D、-U、-I 和 -H 选项将这些选项传递到 C 预处理程序。如果使用 -x 选项，则可选择 C 标准的符合程度。有关调用 C 预处理程序时由 D 编译器定义的一组令牌的说明，请参见 -X。

-D *name* [=*value*]

调用 `cpp(1)`（使用 -c 选项启用）时定义 *name*。如果指定等号 (=) 和其他 *value*，则 *name* 将赋予相应的值。此选项将 -D 选项传递到每个 `cpp` 调用。

-e

在编译所有请求并使用匿名跟踪状态（-a 选项）之后、启用任何探测器之前退出。可以结合使用该选项与 -a 选项以列出匿名跟踪数据，然后退出。还可将该选项与 D 编译器选项结合使用。此结合使用将验证无需真正执行程序并启用相应检测过程验就对程序进行编译。

-f [[*provider*:]*module*:]*function* [[*predicate*]*action*]]

指定要跟踪或列出的函数名称（-l 选项）。相应参数可以包括下列任何探测器说明形式：*provider:module:function*、*module:function* 或 *function*。未指定的探测器说明字段将留为空白，并且与任何探测器相匹配，无论这些字段中的值是什么。如果说明中仅指定了 *function* 限定符，则会匹配所有包含相应 *function* 的探测器。可将可选 D 探测器子句用作 -f 参数的后缀。在命令行上，可以一次指定多个 -f 选项。

-F

通过标识函数的入口和返回位置来合并跟踪输出。对函数入口的探测报告将进行缩进，并在其输出内容前附带 -> 前缀。对函数返回的探测报告将不进行缩进，并在其输出内容前附带 <- 前缀。系统调用进入探测器报告将进行缩进，并在其输出内容前附带 => 前缀。系统调用返回探测器报告将不进行缩进，并在其输出内容前附带 <= 前缀。

-G

生成包含嵌入式 DTrace 程序的 ELF 文件。在程序中指定的 DTrace 探测器将保存在可与其他程序链接的可重定位 ELF 对象中。如果 -o 选项存在，则会使用指定为此操作数参数的路径名来保存 ELF 文件。如果 -o 选项不存在，且 DTrace 程序随附有名称为 *filename.d* 的文件，则 ELF 文件将使用名称 *filename.o* 进行保存。否则，ELF 文件将使用名称 *d.out* 进行保存。

-H

在调用 **cpp(1)**（使用 -c 选项启用）时列出所包含文件的路径名。此选项会将 -H 选项传递到每个 **cpp** 调用，从而使调用将路径名的列表显示到 **stderr**，每个路径名显示为一行。

-h

以指定的提供器定义生成包含与探测器对应的宏的头文件。该选项应用于生成由其他源文件包含以备稍后与 -G 选项结合使用的头文件。如果 -o 选项存在，则会使用指定为此选项参数的路径名来保存头文件。如果 -o 选项不存在，且 DTrace 程序随附有名称为 *filename.d* 的文件，则头文件将使用名称 *filename.h* 进行保存。

-i *probe-id*[[*predicate*] *action*]

指定要跟踪或列出的探测器标识符 (*probe-id*)（-l 选项）。可以使用十进制整数指定探测器 ID，如 **dtrace -l** 所示。可将可选 D 探测器子句用作 -i 参数的后缀。一次可以指定多个 -i 选项。

-I *path*

在调用 **cpp(1)**（使用 -c 选项启用）时，将指定的目录 *path* 添加到 **#include** 文件的搜索路径。此选项将 -I 选项传递到每个 **cpp** 调用。指定的 *path* 将插入到缺省目录列表之前的搜索路径中。

-L *path*

将指定的目录 *path* 添加到 DTrace 库的搜索路径。DTrace 库用于包含可在编写 D 程序时使用的常见定义。指定的 *path* 将添加到缺省库搜索路径后面。

-l

列出探测器而不是启用它们。如果指定了 -l 选项，**dtrace** 将生成与使用 -P、-m、-f、-n、-i 和 -s 选项指定的说明相匹配的探测器报告。如果未指定这些选项中的任何一个，该选项将列出所有探测器。

-m[[*provider:*] *module*: [[*predicate*] *action*]]

指定要跟踪或列出的模块名称（-l 选项）。相应参数可以包括下列任何探测器说明形式：*provider:module* 或 *module*。未指定的探测器说明字段将留为空白，并且与任何探测器相匹配，无论这些字段中的值是什么。如果说明中仅指定了 *module* 限

定符，则会匹配所有包含相应 *module* 的探测器。可将可选 D 探测器子句用作 *-m* 参数的后缀。在命令行上，一次可以指定多个 *-m* 选项。

-n *[[[provider:] module:] function:] name [[predicate] action]*

指定要跟踪或列出的探测器名称（*-l* 选项）。相应参数可以包括下列任何探测器说明形式：*provider:module:function:name*、*module:function:name*、*function:name* 或 *name*。未指定的探测器说明字段将留为空白，并且与任何探测器相匹配，无论这些字段中的值是什么。如果说明中仅指定了 *name* 限定符，则会匹配所有包含相应 *name* 的探测器。可将可选 D 探测器子句用作 *-n* 参数的后缀。在命令行上，一次可以指定多个 *-n* 选项。

-o output

针对 *-A*、*-G*、*-h* 以及 *-l* 选项或跟踪的数据本身指定 *output* 文件。如果 *-A* 选项存在而 *-o* 不存在，则缺省输出文件为 */kernel/drv/dtrace.conf*。如果 *-G* 选项存在并且 *-s* 选项的参数为 *filename.d* 形式，而 *-o* 不存在，则缺省输出文件为 *filename.o*。否则，缺省输出文件为 *d.out*。

请注意，随着使用 *-o* 选项连续调用 *dtrace*，*dtrace* 不会覆盖输出文件，而是附加到输出文件。

-p pid

获取指定的进程 ID *pid*，高速缓存其符号表，然后在完成时退出。如果命令行上有多个 *-p* 选项，*dtrace* 将在所有命令退出后退出，并报告每个进程终止时的退出状态。第一个进程 ID 可供在命令行上（或使用 *-s* 选项通过 *\$target* 宏变量）指定的任何 D 程序使用。有关宏变量的更多信息，请参见《[Solaris Dynamic Tracing Guide](#)》。

-P provider [[predicate] action]

指定要跟踪或列出的提供器名称（*-l* 选项）。其余的探测器说明字段（模块、函数和名称）将留为空白，并且与任何探测器相匹配，无论这些字段中的值是什么。可将可选 D 探测器子句用作 *-P* 参数的后缀。在命令行上，可以一次指定多个 *-P* 选项。

-q

设置静默模式。*dtrace* 将取消显示消息（如与指定选项和 D 程序相匹配的探测器数目），并且将不列出标题、CPU ID、探测器 ID，也不会输出中插入新行。只有 D 程序语句（如 *trace()* 和 *printf()*）跟踪并格式化的数据才会显示在 *stdout* 中。

-s

编译指定的 D 程序源文件。如果 *-e* 选项存在，则将编译该程序，但不会启用检测过程。如果 *-l* 选项存在，则将编译该程序，并且会列出与其相匹配的一组探测器，但不会启用检测过程。如果 *-e*、*-l*、*-G* 或 *-A* 均不存在，则启用 D 程序指定的检测过程并开始跟踪。

-S

显示 D 编译器中间代码。对于为每个 D 程序生成的中间代码，D 编译器会将它的报告生成到 *stderr*。

-U *name*

在调用 **cpp(1)**（使用 **-c** 选项启用）时取消定义指定的 *name*。此选项将 **-U** 选项传递到每个 **cpp** 调用。

-v

设置详细模式。如果指定了 **-v** 选项，**dtrace** 将生成程序稳定性报告，说明所指定 D 程序的最低接口稳定性和相关性级别。DTrace 稳定性级别在《[Solaris Dynamic Tracing Guide](#)》中进行了详细阐述。

-V

报告 **dtrace** 支持的最高 D 编程接口版本。版本信息将列显到 **stdout**，然后 **dtrace** 命令将退出。有关 DTrace 版本控制功能的更多信息，请参见《[Solaris Dynamic Tracing Guide](#)》。

-w

允许使用 **-s**、**-P**、**-m**、**-f**、**-n** 或 **-i** 选项在指定的 D 程序中执行破坏性操作。如果未指定 **-w** 选项，**dtrace** 将不允许编译或启用包含破坏性操作的 D 程序。

-x *arg* [=val]

启用或修改 DTrace 运行时选项或 D 编译器选项。选项列表可在《[Solaris Dynamic Tracing Guide](#)》中找到。将通过指定名称来启用布尔选项。将通过使用等号 (=) 分隔选项名称和值来设置包含值的选项。

-X a | c | s | t

指定在调用 **cpp(1)**（使用 **-c** 选项启用）时应选择的 ISO C 标准的符合程度。根据参数字母的值，**-x** 选项参数会影响 **__STDC__** 宏的值和存在性。

-x 选项支持以下参数：

- a 缺省。ISO C 及 K&R 兼容性扩展，包含 ISO C 所需的语义更改。如果未指定 **-x**，则此模式为缺省模式。在将 **cpp** 与 **-xa** 选项一起调用时，预定义宏 **__STDC__** 具有值 0。
- c 一致性。严格符合 ISO C（不包含 K&R C 兼容性扩展）。将 **cpp** 与 **-xc** 选项一起调用时，预定义宏 **__STDC__** 具有值 1。
- s 仅 K&R C。将 **cpp** 与 **-xs** 选项一起调用时，不会定义宏 **__STDC__**。
- t 转换。ISO C 及 KRC 兼容性扩展，不包含 ISO C 所需的语义更改。将 **cpp** 与 **-xt** 选项一起调用时，预定义宏 **__STDC__** 具有值 0。

因为 **-x** 选项仅影响 D 编译器调用 C 预处理程序的方式，所以从 D 的角度看来，**-xa** 和 **-xt** 选项是等同的，提供这两个选项都是为了简化 C 生成环境中设置的重复使用。

无论 **-x** 模式如何，始终会指定下列附加 C 预处理程序定义，并且它们在所有模式中有效：

- **__sun**
- **__unix**

- `__SVR4`
- `__sparc`（仅在 SPARC 系统上）
- `__sparcv9`（仅当在 SPARC 系统上编译 64 位程序时）
- `__i386`（仅当在 x86 系统上编译 32 位程序时）
- `__amd64`（仅当在 x86 系统上编译 64 位程序时）
- `__'uname -s' 'uname -r'`（例如，`__SunOS_5_10`）
- `__SUNW_D=1`
- `__SUNW_D_VERSION=0xMMmmmmuuu`

其中，*MM* 是以十六进制表示的主要发行版值，*mmm* 是以十六进制表示的次要发行版值，而 *uuu* 是以十六进制表示的微发行版值。有关 DTrace 版本控制的更多信息，请参见《[Solaris Dynamic Tracing Guide](#)》。

-Z
允许与任何探测器都不匹配的探测器说明。在未指定 **-z** 选项时，如果在 D 程序文件中（使用 **-s** 选项）或在命令行上（使用 **-P**、**-m**、**-f**、**-n** 或 **-i** 选项）指定的任何探测器说明中包含不与任何已知探测器匹配的说明，则 **dtrace** 会报告错误并退出。

操作数
可以在 **dtrace** 命令上指定零或多个附加参数，以定义一组宏变量（**\$1**、**\$2** 等等）。附加参数可以用于使用 **-s** 选项指定的 D 程序中或命令行上。宏变量的使用在《[Solaris Dynamic Tracing Guide](#)》中有进一步的描述。

退出状态
将返回下列退出值：

- 0 成功完成。

对于 D 程序请求，退出状态 **0** 指示已成功编译程序，并且已成功启用探测器，或者已成功检索匿名状态。即使指定的跟踪请求遇到错误或删除，**dtrace** 也会返回 **0**。
- 1 出现错误。

对于 D 程序请求，退出状态 **1** 指示程序编译失败或者无法满足指定的请求。
- 2 指定的命令行选项或参数无效。

属性
有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/dtrace
接口稳定性	请参见下文。

命令行语法是 Committed（已确定）。人可阅读的输出是 Uncommitted（未确定）。

另请参见

[cpp\(1\)](#)、[isainfo\(1\)](#)、[ssh\(1\)](#)、[libdtrace\(3LIB\)](#)、[driver.conf\(4\)](#)、[attributes\(5\)](#)、[dtrace\(7D\)](#)
《《Solaris Dynamic Tracing Guide》》

用法

使用 `-p` 标志后，`dtrace` 将在检查目标进程并报告结果时停止这些进程。进程被停止后将无法执行任何操作。这意味着，例如，如果 X 服务器由在 X 服务器控制下的窗口中运行的 `dtrace` 进行检查，则整个窗口系统可能会出现死锁现象，因为 `proc` 工具将尝试将其结果显示到无法刷新的窗口。在此情况下，使用 [ssh\(1\)](#) 从另一个系统登录，中止违例 `proc` 工具以清除死锁。

引用名

dumpadm – configure operating system crash dump

用法概要

```
/usr/sbin/dumpadm [-nuy] [-c content-type] [-d dump-device]
                  [-m mink | minm | min%] [-s savecore-dir]
                  [-r root-dir] [-z on | off]
```

描述

The dumpadm program is an administrative command that manages the configuration of the operating system crash dump facility. A crash dump is a disk copy of the physical memory of the computer at the time of a fatal system error. When a fatal operating system error occurs, a message describing the error is printed to the console. The operating system then generates a crash dump by writing the contents of physical memory to a predetermined dump device, which is typically a local disk partition. The dump device can be configured by way of dumpadm. Once the crash dump has been written to the dump device, the system will reboot.

Fatal operating system errors can be caused by bugs in the operating system, its associated device drivers and loadable modules, or by faulty hardware. Whatever the cause, the crash dump itself provides invaluable information to your support engineer to aid in diagnosing the problem. As such, it is vital that the crash dump be retrieved and given to your support provider. Following an operating system crash, the [savecore\(1M\)](#) utility is executed automatically during boot to retrieve the crash dump from the dump device and write it to your file system in compressed form, to a file name `vmdump.X`, where *X* is an integer identifying the dump. Afterwards, [savecore\(1M\)](#) can be invoked on the same or another system to expand the compressed crash dump to a pair of files named `unix.X` and `vmcore.X`. The directory in which the crash dump is saved on reboot can be configured using dumpadm.

For systems with a UFS root file system, the default dump device is configured to be an appropriate swap partition. Swap partitions are disk partitions reserved as virtual memory backing store for the operating system. Thus, no permanent information resides in swap to be overwritten by the dump. See [swap\(1M\)](#). For systems with a ZFS root file system, dedicated ZFS volumes are used for swap and dump areas. For further information about setting up a dump area with ZFS, see the ZFS Administration Guide. To view the current dump configuration, use the dumpadm command with no arguments:

```
example# dumpadm
```

```
    Dump content: kernel pages
    Dump device: /dev/dsk/c0t0d0s1 (swap)
Savecore directory: /var/crash
    Savecore enabled: yes
    Save compressed: yes
```

When no options are specified, dumpadm displays the current crash dump configuration. The example above shows the set of default values: the dump content is set to kernel memory pages only, the dump device is a swap disk partition, the directory for savecore files is set to `/var/crash/`, savecore is set to run automatically on reboot and save the crash dump in a compressed format.

When one or more options are specified, dumpadm verifies that your changes are valid, and if so, reconfigures the crash dump parameters and displays the resulting configuration. You must be root to view or change dump parameters.

Upon system installation, dumpadm establishes a dump device of sufficient size, based on kernel memory and other internal information, to accommodate a dump file. If you subsequently attempt to create a dump device that is too small to store the dump file, dumpadm issues an error message and the operation fails.

选项

The following options are supported:

-c content-type

Modify the dump configuration so that the crash dump consists of the specified dump content. The content should be one of the following:

kernel

Kernel memory pages only.

all

All memory pages.

curproc

Kernel memory pages, and the memory pages of the process whose thread was currently executing on the CPU on which the crash dump was initiated. If the thread executing on that CPU is a kernel thread not associated with any user process, only kernel pages will be dumped.

-d dump-device

Modify the dump configuration to use the specified dump device. The dump device may be one of the following:

dump-device

A specific dump device specified as an absolute pathname, such as `/dev/dsk/cNtNdNsN` when the system is running a UFS root file system. Or, specify a ZFS volume, such as `/dev/zvol/dsk/rpool/dump`, when the system is running a ZFS root file system.

swap

If the special token `swap` is specified as the dump device, dumpadm examines the active swap entries and selects the most appropriate entry to configure as the dump device. See [swap\(1M\)](#). Refer to the NOTES below for details of the algorithm used to select an appropriate swap entry. When the system is first installed with a UFS root file system, dumpadm uses the value for `swap` to determine the initial dump device setting. A given ZFS volume cannot be configured for both the swap area and the dump device.

-m mink | minm | min%

Create a `minfree` file in the current savecore directory indicating that savecore should maintain at least the specified amount of free space in the file system where the savecore directory is located. The `min` argument can be one of the following:

k

A positive integer suffixed with the unit k specifying kilobytes.

m

A positive integer suffixed with the unit m specifying megabytes.

%

A % symbol, indicating that the minfree value should be computed as the specified percentage of the total current size of the file system containing the savecore directory.

The savecore command will consult the minfree file, if present, prior to writing the dump files. If the size of these files would decrease the amount of free disk space below the minfree threshold, no dump files are written and an error message is logged. The administrator should immediately clean up the savecore directory to provide adequate free space, and re-execute the savecore command manually. The administrator can also specify an alternate directory on the savecore command-line.

-n

Modify the dump configuration to not run savecore automatically on reboot. This is not the recommended system configuration; if the dump device is a swap partition, the dump data will be overwritten as the system begins to swap. If savecore is not executed shortly after boot, crash dump retrieval may not be possible.

-r *root-dir*

Specify an alternate root directory relative to which dumpadm should create files. If no -r argument is specified, the default root directory / is used.

-s *savecore-dir*

Modify the dump configuration to use the specified directory to save files written by savecore. The directory should be an absolute path and exist on the system. If upon reboot the directory does not exist, it will be created prior to the execution of savecore. See the NOTES section below for a discussion of security issues relating to access to the savecore directory. The default savecore directory is /var/crash/.

-u

Forcibly update the kernel dump configuration based on the contents of /etc/dumpadm.conf. Normally this option is used only on reboot when starting svc:/system/dumpadm:default, when the dumpadm settings from the previous boot must be restored. Your dump configuration is saved in the configuration file for this purpose. If the configuration file is missing or contains invalid values for any dump properties, the default values are substituted. Following the update, the configuration file is resynchronized with the kernel dump configuration.

-y

Modify the dump configuration to automatically run savecore on reboot. This is the default for this dump setting. See NOTES.

`-z on | off`
Modify the dump configuration to control the operation of `savecore` on reboot. The options are `on`, to enable saving core files in a compressed format, and `off`, to automatically uncompress the crash dump file. The default is `on`, because crash dump files can be very large and require less file system space if saved in a compressed format.

示例

示例 1 Reconfiguring The Dump Device To A Dedicated Dump Device:
The following command reconfigures the dump device to a dedicated dump device:
example# `dumpadm -d /dev/dsk/c0t2d0s2`

```
Dump content: kernel pages
Dump device: /dev/dsk/c0t2d0s2 (dedicated)
Savecore directory: /var/crash
Savecore enabled: yes
Save compressed: yes
```

退出状态

The following exit values are returned:

- 0
Dump configuration is valid and the specified modifications, if any, were made successfully.
- 1
A fatal error occurred in either obtaining or modifying the dump configuration.
- 2
Invalid command line options were specified.

文件

`/dev/dump`
Dump device.

`/etc/dumpadm.conf`
Contains configuration parameters for `dumpadm`. Modifiable only through that command.

`savecore-directory/minfree`
Contains minimum amount of free space for `savecore-directory`. See [savecore\(1M\)](#).

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[svcs\(1\)](#), [uname\(1\)](#), [savecore\(1M\)](#), [svcadm\(1M\)](#), [swap\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注	The system crash dump service is managed by the service management facility, smf(5), under the service identifier: <pre>svc:/system/dumpadm:default</pre> Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M). The service's status can be queried using the svcs(1) command.
Dump Device Selection	When the special swap token is specified as the argument to <code>dumpadm -d</code> the utility will attempt to configure the most appropriate swap device as the dump device. <code>dumpadm</code> configures the largest swap block device as the dump device; if no block devices are available for swap, the largest swap entry is configured as the dump device. If no swap entries are present, or none can be configured as the dump device, a warning message will be displayed. While local and remote swap files can be configured as the dump device, this is not recommended.
Dump Device/Swap Device Interaction (UFS File Systems Only)	In the event that the dump device is also a swap device, and the swap device is deleted by the administrator using the <code>swap -d</code> command, the <code>swap</code> command will automatically invoke <code>dumpadm -d swap</code> in order to attempt to configure another appropriate swap device as the dump device. If no swap devices remain or none can be configured as the dump device, the crash dump will be disabled and a warning message will be displayed. Similarly, if the crash dump is disabled and the administrator adds a new swap device using the <code>swap -a</code> command, <code>dumpadm -d swap</code> will be invoked to re-enable the crash dump using the new swap device. Once <code>dumpadm -d swap</code> has been issued, the new dump device is stored in the configuration file for subsequent reboots. If a larger or more appropriate swap device is added by the administrator, the dump device is not changed; the administrator must re-execute <code>dumpadm -d swap</code> to reselect the most appropriate device from the new list of swap devices.
Minimum Free Space	If the <code>dumpadm -m</code> option is used to create a <code>minfree</code> file based on a percentage of the total size of the file system containing the <code>savecore</code> directory, this value is not automatically recomputed if the file system subsequently changes size. In this case, the administrator must re-execute <code>dumpadm -m</code> to recompute the <code>minfree</code> value. If no such file exists in the <code>savecore</code> directory, <code>savecore</code> will default to a free space threshold of one megabyte. If no free space threshold is desired, a <code>minfree</code> file containing size 0 can be created.
Security Issues	If, upon reboot, the specified <code>savecore</code> directory is not present, it will be created prior to the execution of <code>savecore</code> with permissions 0700 (read, write, execute by owner only) and owner <code>root</code>. It is recommended that alternate <code>savecore</code> directories also be created with similar permissions, as the operating system crash dump files themselves may contain secure information.
Default for savecore	System installation software might reserve a dedicated dump device (for example, a disk slice or a ZFS volume). In such a case, the <code>dumpadm</code> default can be set to <code>-n</code>, meaning that <code>savecore</code> does not run automatically when the system reboots. A crash image will be preserved on the dump device. Run <code>/usr/sbin/savecore</code> manually as <code>root</code> to retrieve the crash image and copy it to a file under <code>/var/crash</code>. The crash image will remain on the dump device until overwritten by a later one.

引用名	editmap – query and edit single records in database maps for sendmail
用法概要	<pre>editmap -C file [-N] [-f] [-q -u -x] maptype mapname key ["value"]...</pre>
描述	The editmap command queries or edits one record in a database maps used by the keyed map lookups in sendmail(1M). Arguments are passed on the command line and output (for queries) is directed to standard output. Depending on how it is compiled, editmap handles up to three different database formats, selected using the <i>maptype</i> parameter. See OPERANDS. If the TrustedUser option is set in the sendmail configuration file and editmap is invoked as root, the generated files are owned by the specified TrustedUser.
选项	The following options are supported: -C <i>file</i> Use the specified sendmail configuration file (<i>file</i>) to look up the TrustedUser option. -f Disable the folding of all upper case letters in the key to lower case. Normally, all upper case letters in the key are folded to upper case. This is intended to mesh with the -f flag in the K line in sendmail.cf. The value is never case folded. -N Include the null byte that terminates strings in the map (for alias maps). -q Query the map for the specified key. If found, print value to standard output and exit with 0. If not found then print an error message to stdout and exit with EX_UNAVAILABLE. -u Update the record for key with value or inserts a new record if one doesn't exist. Exits with 0 on success or EX_IOERR on failure. -x Delete the specific key from the map. Exit with 0 on success or EX_IOERR on failure.
操作数	The following operands are supported: <i>key</i> The left hand side of a record. Each record is of the form: <i>key value</i> <i>key</i> and <i>value</i> are separated by white space. <i>mapname</i> File name of the database map being created. <i>maptype</i> Specifies the database format. The following <i>maptype</i> parameters are available: - dbm Specifies DBM format maps. - bt ree Specifies B-Tree format maps.

key and *value* are separated by white space.

See `attributes(5)` for descriptions of the following attributes:

```
makemap(1M), sendmail(1M), attributes(5)
```

引用名	edquota – edit user quotas for ufs file system
用法概要	<pre>edquota [-p <i>proto_user</i>] <i>username</i>...</pre> <pre>edquota -t</pre>
描述	edquota is a quota editor. One or more users may be specified on the command line. For each user a temporary file is created with an ASCII representation of the current disk quotas for that user for each mounted ufs file system that has a quotas file, and an editor is then invoked on the file. The quotas may then be modified, new quotas added, etc. Upon leaving the editor, edquota reads the temporary file and modifies the binary quota files to reflect the changes made. The editor invoked is vi(1) unless the EDITOR environment variable specifies otherwise. Only the super-user may edit quotas. In order for quotas to be established on a file system, the root directory of the file system must contain a file, owned by root, called quotas. (See quotaon(1M).) <i>proto_user</i> and <i>username</i> can be numeric, corresponding to the UID of a user. Unassigned UIDs may be specified; unassigned names may not. In this way, default quotas can be established for users who are later assigned a UID. If no options are specified, the temporary file created will have one or more lines of the format, where a block is considered to be a 1024 byte (1K) block: <pre>fs mount_point blocks (soft =number, \ hard =number) inodes (soft =number, \ hard =number)</pre> The <i>number</i> fields may be modified to reflect desired values. The following options are supported: -p Duplicate the quotas of the <i>proto_user</i> specified for each <i>username</i> specified. This is the normal mechanism used to initialize quotas for groups of users. -t Edit the soft time limits for each file system. If the time limits are zero, the default time limits in <code>/usr/include/sys/fs/ufs_quota.h</code> are used. The temporary file created will have one or more lines of the form <pre>fs mount_point blocks time limit = number tmunit, files time limit = number tmunit</pre> <i>tmunit</i> may be one of "month", "week", "day", "hour", "min" or "sec"; characters appended to these keywords are ignored, so you may write "months" or "minutes" if you prefer. The <i>number</i> and <i>tmunit</i> fields may be modified to set desired values. Time limits are printed in the greatest possible time unit such that the value is greater than or equal to one. If "default" is printed after the <i>tmunit</i>, this indicates that the value shown is zero (the default).
选项	

- 用法

See [largefile\(5\)](#) for the description of the behavior of edquota when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).
- 文件

quotas

quota file at the file system root

/etc/mnttab

table of mounted file systems
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

- 另请参见

[vi\(1\)](#), [quota\(1M\)](#), [quotacheck\(1M\)](#), [quotaon\(1M\)](#), [repquota\(1M\)](#), [attributes\(5\)](#), [largefile\(5\)](#), [quotactl\(7I\)](#)
- 附注

All UIDs can be assigned quotas.

引用名	eeprom – EEPROM 显示和装载实用程序
用法概要	<code>/usr/sbin/eeprom [-] [-f device] [parameter[=value]]</code>
描述	<code>eeprom</code> 显示或更改 EEPROM 中的参数值。它按照给定的顺序处理参数。处理带有 <i>value</i> 的 <i>parameter</i> 时，<code>eeprom</code> 对 EEPROM 进行指定的更改；否则，显示 <i>parameter</i> 的值。如果未给定任何参数说明符，<code>eeprom</code> 会显示所有 EEPROM 参数的值。‘-’（连字符）标志指定将从标准输入（每行一个 <i>parameter</i> 或 <i>parameter=value</i>）读取参数和值。 仅超级用户可以更改 EEPROM 内容。 <code>eeprom</code> 验证 EEPROM 校验和，如果校验和不正确，则进行报警。 <i>platform-name</i> 是平台实现的名称，可使用 <code>uname(1)</code> 的 <code>-i</code> 选项找到。
SPARC	基于 SPARC 的系统通过 <code>security-mode</code> 、 <code>security-password</code> 和 <code>security-#badlogins</code> 属性，来使用 <code>eeprom</code> 实现固件口令保护。
x86	使用驻留在特定平台的引导区域中的文件模拟 EEPROM 存储。<code>/boot/solaris/bootenv.rc</code> 文件模拟 EEPROM 存储。 由于基于 x86 的系统通常在系统 BIOS 中实现口令保护，因此 <code>eeprom</code> 程序中不提供口令保护支持。虽然可以在基于 x86 的系统上设置 <code>security-mode</code>、<code>security-password</code> 和 <code>security-#badlogins</code> 属性，但是这些属性在基于 x86 的系统上没有任何特殊含义或行为。
选项	<code>-f device</code> 使用 <i>device</i> 作为 EEPROM 设备。
操作数	
仅限于 x86	<i>acpi-user-options</i> 一种配置变量，用于控制电源管理规范“高级配置和电源接口 (Advanced Configuration and Power Interface, ACPI)”的使用。此变量可接受的值取决于您所使用的 Solaris 操作系统的发行版。 对于所有发行版的 Solaris 10 和 Solaris 11，值 <code>0x0</code> 表示如果 ACPI 在系统上可用，将尝试使用它。值 <code>0x2</code> 表示禁用 ACPI。 对于 Solaris 10 1/06 发行版，值 <code>0x8</code> 表示如果 ACPI 在系统中可用，将以与之前发行版的 Solaris 10 兼容的模式尝试使用它。Solaris 10 1/06 的缺省值为 <code>0x8</code>。 对于 1/06 发行版之后的 Solaris 10 以及 Solaris 11，缺省值为 <code>0x0</code>。 大多数用户可以放心接受缺省值（在可用时启用 ACPI）。如果怀疑在 Solaris 1/06 之后的 Solaris 发行版中存在与使用 ACPI 有关的问题，建议首先尝试使用值 <code>0x8</code>；如果未获得满意的结果，再使用值 <code>0x02</code>。

console

指定控制台设备。可能的值包括 `ttya`、`ttyb`、`text`、`graphics` 和 `force-text`。在 `text` 模式下，控制台输出进入帧缓冲区，而输入来自键盘。`text` 模式的变体 `graphics` 会显示具有动画效果的图像，直到按下某个键或控制台登录（`sulogin` 或 `kmdb`）需要进行控制台交互。`text` 的另外一个变体 `force-text` 将通过将 VGA 适配器设置为 VGA 文本模式来避免将其用作位映射设备。如果此属性不存在，控制台设备回退到 `input-device` 和 `output-device` 指定的设备。如果 `console` 属性以及 `input-device` 和 `output-device` 属性对都不存在，控制台缺省为帧缓冲区和键盘。

screen-#columns screen-#rows

当 `screen` 在位映射设备上设置为 `graphics` 或 `text` 时，`screen-#columns` 和 `screen-#rows` 允许指定所需数量的文本列和行。它们分别缺省为 80 和 24。

Nvram 配置参数

并非所有 OpenBoot 系统均支持所有参数。缺省值取决于系统以及 PROM 修订。请参见在 `ok` (OpenBoot) 提示符下输入 `printenv` 命令时，该命令的 "Default Value"（缺省值）列中的输出，以确定系统的缺省值。

auto-boot?

如果为 `true`，在打开电源或重置后自动引导。缺省值为 `true`。在 x86 上，此参数由 `grub` 菜单文件控制。请参见 [installgrub\(1M\)](#)。

ansi-terminal?

用于控制终端仿真程序行为的配置变量。值 `false` 使终端仿真程序停止解释 ANSI 转义序列，而是将其回显至输出设备。缺省值为 `true`。

boot-args

保存传递至引导子系统的参数字符串。例如，可以使用 `boot-args=' - install dhcp'` 请求自定义 `jumpstart` 安装。请参见 [boot\(1M\)](#)、[kmdb\(1\)](#) 和 [kernel\(1M\)](#)。

boot-command

`auto-boot?` 为 `true` 时执行的命令。缺省值为 `boot`。

boot-device

从其进行引导的设备。*boot-device* 可能包含 0 个或多个以空格分隔的设备说明符。每个设备说明符可以是 `prom` 设备别名或 `prom` 设备路径。从第一个设备说明符开始，引导 `prom` 将尝试打开列表中随后的每一个设备说明符。成功打开的第一个设备说明符将用作从中进行引导的设备。缺省值为 `disk net`。

boot-file

要进行引导的文件（空字符串会使辅助引导程序选择缺省值）。缺省值为空字符串。

boot-from

引导设备和文件（仅限 OpenBoot PROM 版本 1.x）。缺省值为 `vmunix`。

boot-from-diag

诊断引导设备和文件（仅限 OpenBoot PROM 版本 1.x）。缺省值为 `le( )unix`。

boot-ncpus

用于控制系统引导应使用的处理器数的配置变量。缺省情况下，系统引导时使用所支持的最大处理器数。

comX-noprobe

其中，X 是串行端口号，用于阻止串行端口 X 上的设备探测器。

diag-device

诊断引导源设备。缺省值为 `net`。

diag-file

在诊断模式下从其进行引导的文件。缺省值为空字符串。

diag-level

诊断级别。值包括 `off`、`min`、`max` 和 `menus`。可能还具有特定平台的值。如果设置为 `off`，则不调用 POST。如果调用 POST，该值可用作 POST 的参数并由其进行解释。缺省值为 `platform-dependent`。

diag-switch?

如果为 `true`，则在诊断模式下运行。在大多数桌面系统上，缺省值为 `false`；在大多数服务器上，缺省值为 `true`。

error-reset-recovery

在错误重置陷阱后恢复。缺省值为特定平台的设置。

在支持此变量的平台上，它会代替

`watchdog-reboot?`、`watchdog-sync?`、`redmode-reboot?`、`redmode-sync?`、`sir-sync?` 和 `xir-sync?` 参数。

选项包括：

none

列出说明重置陷阱的消息，并转到 OpenBoot PROM 用户界面中的 *aka* OK 提示符。

sync

在重置陷阱后，调用 OpenBoot PROM 的 `sync` 一词。在外部启动重置 (`externally initiated reset, XIR`) 陷阱后，一些平台可能会将其视为 `none`。

boot (引导归档文件)

在重置陷阱后重新引导。在 XIR 陷阱后，一些平台可能会将其视为 `none`。

fcode-debug?

如果为 `true`，则包含插件设备 FCode 的名称参数。缺省值为 `false`。

hardware-revision

系统版本信息。

input-device

打开电源时所用的输入设备（通常为 `keyboard`、`ttya` 或 `ttyb`）。缺省值为 `keyboard`。

keyboard-click?

如果为 **true**，则启用键盘单击。缺省值为 **false**。

keyboard-layout

为非自识别键盘（7c 类型）指定布局名称的字符串。调用 **kbd -s**，以获取可接受的布局名称列表。请参见 **kbd(1)**。

keymap

自定义键盘的键盘映射。

last-hardware-update

系统更新信息。

load-base

客户机程序的缺省装载地址。缺省值为 **16384**。

local-mac-address?

如果为 **true**，网络驱动程序会使用其自身的 MAC 地址，而非系统地址。缺省值为 **false**。

mfg-mode

POST 的生产模式参数。可能值包括 **off** 或 **chamber**。值作为参数传递至 POST。缺省值为 **off**。

mfg-switch?

如果为 **true**，将重复系统自检直到被 STOP-A 中断。缺省值为 **false**。

nvrnrc

NVRAMRC 的内容。缺省值为空。

network-boot-arguments

PROM 将用于进行网络引导的参数。缺省值为空字符串。**network-boot-arguments** 可用于指定要使用的引导协议 (RARP/DHCP) 以及在进程上要使用的系统知识范围。

网络引导支持的参数的语法为：

[protocol,] [key=value,]*

所有参数均是可选的，显示顺序随意。除非参数位于列表末尾，否则需要使用逗号。如果进行了指定，参数将优先于任何缺省值，或者，如果使用 DHCP 进行引导，则优先于 DHCP 服务器为那些参数提供的配置信息。

上述的 *protocol* 指定要使用的地址搜索协议。

下列配置参数被指定为 *key=value* 属性对。

tftp-server

TFTP 服务器的 IP 地址

file

针对 WAN Boot 要使用 TFTP 或 URL 下载的文件

host-ip
客户机的 IP 地址（采用点分十进制记法）

router-ip
缺省路由器的 IP 地址（采用点分十进制记法）

subnet-mask
子网掩码（采用点分十进制记法）

client-id
DHCP 客户机标识符

hostname
DHCP 事务中要使用的主机名

http-proxy
HTTP 代理服务器规范 (IPADDR[:PORT])

tftp-retries
TFTP 重试的最大次数

dhcp-retries
DHCP 重试的最大次数

如果未指定任何参数（即 **network-boot-arguments** 为空字符串），PROM 将使用特定平台的缺省地址搜索协议。

如果在指定其他系统参数时缺少协议参数，则意味着需要手动配置。

手动配置要求客户机带有引导所需的所有信息。如果使用手动配置，必须在 **network-boot-arguments** 中提供 PROM 装载第二阶段引导程序所需的信息，而第二阶段引导程序所需的信息可以作为 **boot** 程序的参数进行指定，也可以通过 **boot** 程序的交互式命令解释程序进行指定。

使用手动配置时 PROM 所需的信息包括引导客户机的 IP 地址、引导文件的名称，以及提供引导文件映像的服务器的地址。根据网络配置，可能还需要指定要使用的缺省路由器的子网掩码和地址。

oem-banner
自定义 OEM 标题（通过将 **oem-banner?** 设置为 **true** 启用）。缺省值为空字符串。

oem-banner?
如果为 **true**，则使用自定义 OEM 标题。缺省值为 **false**。

oem-logo
字节数组自定义 OEM 徽标（通过将 **oem-logo?** 设置为 **true** 启用）。以十六进制形式显示。

oem-logo?
如果为 **true**，则使用自定义 OEM 徽标（否则，使用 Oracle 徽标）。缺省值为 **false**。

pci-mem64?

如果为 `true`，OpenBoot PROM 会将 64 位 PCI 内存地址分配至能够支持 64 位地址的 PCI 设备。

此变量仅在 SPARC 平台上可用，且是可选的。一些版本的 SunOS 不支持 PCI MEM64 地址，且如果 OpenBoot PROM 分配 PCI MEM64 地址，它们会以意外的方式出现故障。

缺省值取决于系统。如果变量存在，缺省值适用于随特定平台提供的 SunOS 的最低版本。

output-device

打开电源时所用的输出设备（通常为 `screen`、`ttya` 或 `ttyb`）。缺省值为 `screen`。

redmode-reboot?

指定 `true`，以在 `redmode` 重置陷阱后重新引导。缺省值为 `true`。（仅限 Sun Enterprise 10000。）

redmode-sync?

指定 `true`，以在 `redmode` 重置陷阱后调用 OpenBoot PROM 的 `sync` 一词。缺省值为 `false`。（仅限 Sun Enterprise 10000。）

rootpath

指定操作系统的根设备。

sbus-probe-list

指定以何种顺序对哪些 SBus 插槽进行探测。缺省值为 `0123`。

screen-#columns

屏幕列数（字符数/行）。缺省值为 `80`。

screen-#rows

屏幕行数。缺省值为 `34`。

scsi-initiator-id

主机适配器的 SCSI 总线地址，范围为 0-7。缺省值为 7。

sd-targets

映射 SCSI 磁盘单元（仅限 OpenBoot PROM 版本 1.x）。缺省值为 `31204567`，这表示单元 0 映射至目标 3，单元 1 映射至目标 1，依此类推。

security-#badlogins

错误安全口令的尝试次数。在基于 x86 的系统上，此属性没有任何特殊含义或行为。

security-mode

固件安全级别（选项：`none`、`command` 和 `full`）。如果设置为 `command` 或 `full`，系统将提示您输入 PROM 安全口令。缺省值为 `none`。在基于 x86 的系统上，此属性没有任何特殊含义或行为。

security-password

固件安全口令（从不显示）。仅当 **security-mode** 设置为 **command** 或 **full** 时才可以设置此属性。在基于 x86 的系统上，此属性没有任何特殊含义或行为。

```
example# eeprom security-password=
Changing PROM password:
New password:
Retype new password:
```

selftest-#megs

要测试的 RAM 兆字节数。如果 **diag-switch?** 为 **true**，则忽略此属性。缺省值为 1。

sir-sync?

指定 **true**，以在软件启动重置 (software-initiated reset, SIR) 陷阱后调用 OpenBoot PROM 的 **sync** 一词。缺省值为 **false**。（仅限 Sun Enterprise 10000。）

skip-vme-loopback?

如果为 **true**，POST 将不执行 VMEbus 回送测试。缺省值为 **false**。

st-targets

映射 SCSI 磁带单元（仅限 OpenBoot PROM 版本 1.x）。缺省值为 45670123，这表示单元 0 映射至目标 4，单元 1 映射至目标 5，依此类推。

sunmon-compat?

如果为 **true**，则显示 Restricted Monitor 提示符 (>)。缺省值为 **false**。

testarea

单字节临时字段，可用于读写测试。缺省值为 0。

tpe-link-test?

对于内置双绞线以太网，启用 10baseT 链接测试。缺省值为 **true**。

ttya-mode

TTYA（波特率、数据位数、奇偶校验、停止位数、握手）。缺省值为 9600、8、n、1、-。

按从左至右的顺序，这些字段为：

波特率:

110、300、1200、4800、9600...

Data bits（数据位数）：

5、6、7、8

parity:

n(none)、e(even)、o(odd)、m(mark)、s(space)

Stop bits（停止位数）：

1、1.5、2

Handshake（握手）：
 –(none)、h(hardware:rts/cts)、s(software:xon/xoff)

tttyb-mode

TTYB（波特率、数据位数、奇偶校验、停止位数、握手）。缺省值为 9600、8、n、1、-。

按从左至右的顺序，这些字段为：

波特率：

110、300、1200、4800、9600...

Data bits（数据位数）：

5、6、7、8

Stop bits（停止位数）：

1、1.5、2

parity:

n(none)、e(even)、o(odd)、m(mark)、s(space)

Handshake（握手）：

–(none)、h(hardware:rts/cts)、s(software:xon/xoff)

tttya-ignore-cd

如果为 true，操作系统忽略 TTYA 上的载体检测。缺省值为 true。

tttyb-ignore-cd

如果为 true，操作系统忽略 TTYB 上的载体检测。缺省值为 true。

tttya-rts-dtr-off

如果为 true，操作系统不在 TTYA 上插入 DTR 和 RTS。缺省值为 false。

tttyb-rts-dtr-off

如果为 true，操作系统不在 TTYB 上插入 DTR 和 RTS。缺省值为 false。

use-nvramrc?

如果为 true，则在系统启动时执行 NVRAMRC 中的命令。缺省值为 false。

verbosity

控制 PROM 消息的详细级别。值可以是 debug、max、normal、min 或 none 之一。缺省值为 normal。

version2?

如果为 true，混合 (1.x/2.x) PROM 在版本 2.x 中显示。缺省值为 true。

watchdog-reboot?

如果为 true，则在监视程序重置后重新引导。缺省值为 false。

watchdog-sync?

指定 true，以在监视程序重置陷阱后调用 OpenBoot PROM 的 sync 一词。缺省值为 false。（仅限 Sun Enterprise 10000。）

xir-sync?

指定 true，以在 XIR 陷阱后调用 OpenBoot PROM 的 sync 一词。缺省值为 false。（仅限 Sun Enterprise 10000。）

示例

示例 1 更改 RAM 的兆字节数。

以下示例演示了将系统测试的 RAM 兆字节数从 1 兆字节更改为 2 兆字节的方法。

```
example# eeprom selftest-#megs
selftest-#megs=1
```

```
example# eeprom selftest-#megs=2
```

```
example# eeprom selftest-#megs
selftest-#megs=2
```

示例 2 将 auto-boot? 参数设置为 true。

以下示例演示了将 auto-boot? 参数设置为 true 的方法。

```
example# eeprom auto-boot?=true
```

在用户模式下执行 eeprom 命令时，带有结尾问号 (?) 的参数需要括在双引号 (“”) 中，以免 shell 解释问号。在问号前加一个转义符 (\) 也可以防止 shell 解释问号。

```
example% eeprom "auto-boot?"=true
```

示例 3 使用 network-boot-arguments

要使用 DHCP 作为引导协议和 abcd.example.com 主机名来进行网络引导，请将 network-boot-arguments 中的这些值设置为：

```
example# eeprom network-boot-arguments="dhcp,hostname=abcd.example.com"
```

...然后使用以下命令引导：

```
ok boot net
```

请注意，从 PROM 命令行指定的网络引导参数会导致 network-boot-arguments 的内容被忽略。例如，如果将 network-boot-arguments 设置为如上所示，则 boot 命令：

```
ok boot net:dhcp
```

...导致使用 DHCP，但是在 network-boot-arguments 中指定的 hostname 在网络引导时不会使用。

示例 4 将系统控制台设置为辅助设备

以下命令将设备 /dev/term/a 分配为系统控制台设备。在使用 tip(1) 来建立到主机的 tip 连接之前进行赋值。

在 SPARC 计算机上：

示例 4 将系统控制台设置为辅助设备 (续)

```
# eeeprom output-device=/dev/term/a
```

在 x86 计算机上：

```
# eeeprom console=ttya
```

在 SPARC 计算机上，上述命令已足以用于将控制台分配至辅助设备。对于 x86 计算机，您可能还需要设置串行线的特征（请参见该计算机的 BIOS 文档）。此外，在一些 x86 计算机上，可以使用设备 a 以外的设备，如上所示。例如，如果存在第二个串行端口，可将控制台设置为 ttyb。

示例 5 指定 SPARC 系统引导至 kmdb

要指定 SPARC 计算机引导至 [kmdb\(1\)](#)，请输入以下命令：

```
# eeeprom boot-command="boot -k"
```

文件

```
/boot/solaris/bootenv.rc
```

用于在 x86 计算机上存储 eeeprom 值的文件。

```
/dev/openprom
```

设备文件

```
/usr/platform/platform-name/sbin/eeeprom
```

特定平台版本的 eeeprom。使用 `uname -i` 获取 *platform-name*。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见

[kmdb\(1\)](#)、[passwd\(1\)](#)、[sh\(1\)](#)、[svcs\(1\)](#)、[tip\(1\)](#)、[uname\(1\)](#)、[boot\(1M\)](#)、[kernel\(1M\)](#)、[init\(1M\)](#)、[svcadm\(1M\)](#)

《OpenBoot 3.x Command Reference Manual》

《ONC+ Developer's Guide》

引用名 efdaemon – embedded FCode interpreter daemon

用法概要 /usr/lib/efcode/sparcv9/efdaemon [-d]

描述 efdaemon, the embedded FCode interpreter daemon, invokes the embedded FCode interpreter when the daemon receives an interpretation request. A new session of the interpreter is started for each unique request by invoking the script /usr/lib/efcode/efcode.

efdaemon is used on selected platforms as part of the processing of some dynamic reconfiguration events.

选项 The following option is supported:

-d Set debug output. Log debug messages as LOG_DEBUG level messages by using syslog(). See [syslog\(3C\)](#).

文件	/dev/fcode	FCode interpreter pseudo device, which is a portal for receipt of FCode interpretation requests
	/usr/lib/efcode/efcode	Shell script that invokes the embedded FCode interpreter
	/usr/lib/efcode/interpreter	Embedded FCode interpreter
	/usr/lib/efcode/sparcv9/interpreter	Embedded FCode interpreter

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWefcx, SUNWefcux, system/embedded-fcode-interpreter, SUNWefclx

另请参见 [svcs\(1\)](#), [prtconf\(1M\)](#), [svcadm\(1M\)](#), [syslog\(3C\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The efdaemon service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/platform/sun4u/efdaemon:default

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名 `embedded_su` – allow an application to prompt for credentials and execute commands as the super user or another user

用法概要 `/usr/lib/embedded_su [-] [username [arg...]]`

描述 The `embedded_su` command allows an application to prompt the user for security credentials and then use those credentials to execute a program as another user or role (see [rbac\(5\)](#) for information on role-based access control). The default *username* is `root` (super user).

`embedded_su` is identical to [su\(1M\)](#), except that the user interaction is packaged in a form suitable for another program to interpret and display. Typically, `embedded_su` would be used to allow a graphical program to prompt for the super user password and execute a command as the super user, without requiring that the requesting program be run as the super user.

PROTOCOL `embedded_su` implements a simple protocol over standard input, standard output, and standard error. This protocol consists of three phases, roughly corresponding to PAM initialization, the PAM dialog, and PAM completion.

Phase 1: Initialization

After starting `embedded_su`, the application must send an initialization block on `embedded_su`'s standard input. This block is a text block, as described under “Text Blocks”. There are currently no initialization parameters defined; the application should send an empty block by sending a line consisting solely of a period (.).

Phase 2: Conversation

`embedded_su` then emits zero or more conversation blocks on its standard output. Each conversation block may require zero or more responses.

A conversation block starts with a line consisting of the word `CONV`, followed by whitespace, followed by the number of messages in the conversation block as a decimal integer. The number of messages may be followed by whitespace and additional data. This data, if present, must be ignored.

Each message consists of a line containing a header followed by a text block, as described under “Text Blocks”. A single newline is appended to each message, allowing the message to end with a line that does not end with a newline.

A message header line consists of a PAM message style name, as described in [pam_start\(3PAM\)](#). The message header values are:

<code>PAM_PROMPT_ECHO_OFF</code>	The application is to prompt the user for a value, with echoing disabled.
<code>PAM_PROMPT_ECHO_ON</code>	The application is to prompt the user for a value, with echoing enabled.
<code>PAM_ERROR_MSG</code>	The application is to display the message in a form appropriate for displaying an error.

PAM_TEXT_INFO The application is to display the message in a form appropriate for general information.

The PAM message style may be followed by whitespace and additional data. This data, if present, must be ignored.

After writing all of the messages in the conversation block, if any of them were PAM_PROMPT_ECHO_OFF or PAM_PROMPT_ECHO_ON, embedded_su waits for the response values. It expects the response values one per line, in the order the messages were given.

Phase 3: Completion

After zero or more conversation blocks, embedded_su emits a result block instead of a conversation block.

Upon success, embedded_su emits a single line containing the word “SUCCESS”. The word SUCCESS may be followed by whitespace and additional data. This data, if present, must be ignored.

Upon failure, embedded_su emits a single line containing the word “ERROR”, followed by a text block as described under “Text Bocks”. The text block gives an error message. The word ERROR may be followed by whitespace and additional data. This data, if present, must be ignored.

Text Blocks

Initialization blocks, message blocks, and error blocks are all text blocks. These are blocks of text that are terminated by a line containing a single period (.). Lines in the block that begin with a “.” have an extra “.” prepended to them.

Internationalization All messages are localized to the current locale; no further localization is required.

SECURITY embedded_su uses [pam\(3PAM\)](#) for authentication, account management, and session management. Its primary function is to export the PAM conversation mechanism to an unprivileged program. Like [su\(1M\)](#), the PAM configuration policy can be used to control embedded_su. The PAM service name used is “embedded_su”.

embedded_su is almost exactly equivalent to [su\(1M\)](#) for security purposes. The only exception is that it is slightly easier to use embedded_su in writing a malicious program that might trick a user into providing secret data. For those sites needing maximum security, potentially at the expense of application functionality, the EXAMPLES section shows how to disable embedded_su.

示例 In the following examples, left angle brackets (<<<) indicate a line written by embedded_su and read by the invoking application. Right angle brackets (>>>) indicate a line written by the application and read by embedded_su.

示例 1 Executing a command with the Correct Password

The following example shows an attempt to execute “somecommand” as “someuser”, with the correct password supplied:

```
/usr/lib/embedded_su someuser -c somecommand
>>>.
<<<CONV 1
<<<PAM_PROMPT_ECHO_OFF
<<<Password:
<<<.
>>>[ correct password ]
<<<SUCCESS
[ somecommand executes ]
```

示例 2 Executing a command with the Incorrect Password

The following example shows an attempt to execute “somecommand” as “someuser”, with the incorrect password supplied:

```
/usr/lib/embedded_su someuser -c somecommand
>>>.
<<<CONV 1
<<<PAM_PROMPT_ECHO_OFF
<<<Password:
<<<.
>>>[ incorrect password ]
[ delay ]
<<<ERROR
<<<embedded_su:Sorry
<<<.
[ exit ]
```

示例 3 Message Examples

A pam_message structure with *msg_style* equal to PAM_TEXT_INFO and msg equal to “foo” produces:

```
PAM_TEXT_INFO
foo
.
```

A pam_message structure with *msg_style* equal to PAM_ERROR_MESSAGE and msg equal to “bar\n” produces:

```
PAM_ERROR_MESSAGE
bar
[ blank line ]
.
```

示例3 Message Examples (续)

A pam_message structure with *msg_style* equal to PAM_ERROR_MESSAGE and msg equal to “aaa\nbbb” produces:

```
PAM_ERROR_MESSAGE
aaa
bbb
.
```

A pam_message structure with *msg_style* equal to PAM_TEXT_INFO and msg equal to "" produces:

```
PAM_TEXT_INFO
[ blank line ]
.
```

A pam_message structure with *msg_style* equal to PAM_TEXT_INFO and msg equal to NULL produces:

```
PAM_TEXT_INFO
.
```

示例4 Disabling embedded_su

To disable embedded_su, add a line to the /etc/pam.conf file similar to:

```
embedded_su auth requisite pam_deny.so.1
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[su\(1M\)](#), [pam\(3PAM\)](#), [pam_start\(3PAM\)](#), [attributes\(5\)](#), [rbac\(5\)](#)

引用名

etrn – start mail queue run

用法概要

etrn [-b] [-v] *server-host* [*client-hosts*]

描述

SMTP's ETRN command allows an SMTP client and server to interact, giving the server an opportunity to start the processing of its queues for messages to go to a given host. This is meant to be used in start-up conditions, as well as for mail nodes that have transient connections to their service providers.

The `etrn` utility initiates an SMTP session with the host *server-host* and sends one or more ETRN commands as follows: If no *client-hosts* are specified, `etrn` looks up every host name for which [sendmail\(1M\)](#) accepts email and, for each name, sends an ETRN command with that name as the argument. If any *client-hosts* are specified, `etrn` uses each of these as arguments for successive ETRN commands.

选项

The following options are supported:

- b System boot special case. Make sure localhost is accepting SMTP connections before initiating the SMTP session with *server-host*.

This option is useful because it prevents race conditions between [sendmail\(1M\)](#) accepting connections and *server-host* attempting to deliver queued mail. This check is performed automatically if no *client-hosts* are specified.
- v The normal mode of operation for `etrn` is to do all of its work silently. The `-v` option makes it verbose, which causes `etrn` to display its conversations with the remote SMTP server.

环境变量

No environment variables are used. However, at system start-up, `svc:/network/smtp:sendmail` reads `/etc/default/sendmail`. In this file, if the variable `ETRN_HOSTS` is set, `svc:/network/smtp:sendmail` parses this variable and invokes `etrn` appropriately. `ETRN_HOSTS` should be of the form:

```
"s1:c1.1,c1.2            s2:c2.1 s3:c3.1,c3.2,c3.3"
```

That is, white-space separated groups of *server:client* where *client* can be one or more comma-separated names. The *:client* part is optional. *server* is the name of the server to prod; a mail queue run is requested for each *client* name. This is comparable to running:

```
/usr/lib/sendmail -qR client
```

on the host *server*.

示例

示例1 Using `etrn`

Inserting the line:

```
ETRN_HOSTS="s1.domain.com:clnt.domain.com s2.domain.com:clnt.domain.com"
```

示例1 Using etrn (续)

in /etc/default/sendmail results in svc:/network/smtp:sendmail invoking etrn such that ETRN commands are sent to both s1.domain.com and s2.domain.com, with both having clnt.domain.com as the ETRN argument.

The line:

```
ETRN_HOSTS="server.domain.com:client1.domain.com,client2.domain.com"
```

results in two ETRN commands being sent to server.domain.com, one with the argument client1.domain.com, the other with the argument client2.domain.com.

The line:

```
ETRN_HOSTS="server1.domain.com server2.domain.com"
```

results in set of a ETRN commands being sent to both server1.domain.com and server2.domain.com; each set contains one ETRN command for each host name for which [sendmail\(1M\)](#) accepts email, with that host name as the argument.

文件

- /etc/mail/sendmail.cf sendmail configuration file
- /etc/default/sendmail Variables used by svc:/network/smtp:sendmail

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/smtp/sendmail
Interface Stability	Committed

另请参见

[sendmail\(1M\)](#), [attributes\(5\)](#)

RFC 1985

附注

Not all SMTP servers support ETRN.

引用名 fbconfig – frame buffer configuration utility

用法概要

```
fbconfig [-dev device_file] [-help]

fbconfig -list

fbconfig -gui

fbconfig -res \?

fbconfig [-dev device_file] [-file machine | system]
        [device_specific_options] [-res video-mode] [defaults]
        [-prconf] [-prdid[raw] [parsed]] [-propt]
```

描述

fbconfig is the generic command line interface to query and/or to configure a frame buffer device. Depending on the command line options, fbconfig can invoke a GUI program, or it can invoke a device-specific configuration program. The choice of device configuration program depends both on the device (specified with -dev) and on the currently configured X server.

The options recognized by fbconfig are shown in the first three command forms in the SYNOPSIS section. The remaining forms illustrate options that may be passed to a device-specific configuration program that performs the actual operations. The interpretation of these options will depend upon the specific configuration program that is invoked. The options shown are supported by most device configuration programs.

If the -dev option is omitted, the default frame buffer (/dev/fb or /dev/fb0) is assumed.

选项

The following options are supported by fbconfig:

-dev *device_file*

Specify the frame buffer device by its full pathname or simple filename. Pathnames of installed devices can be displayed using the -list option. If the -dev option is omitted, the default device is /dev/fb.

-gui

Invoke the Graphical User Interface (GUI). The GUI is available if the SUNWdcm package is installed. All other options are ignored.

The GUI can configure devices (as an alternative to the fbconfig command line) and can update the Xservers file without directly editing the file. The GUI allows the user that is logged in on the graphics device or devices to configure which graphics displays the window system should use, their screen layout (where they appear on the user's desktop), and screen properties (X attributes).

In addition, the GUI allows advanced users to create a new video format (resolution) that some graphics devices can select from fbconfig command line or from the device-dependent portion of the GUI. The GUI's online help explains all options and features.

-help

Displays the supported fbconfig command line options, along with a brief explanation of each. Also displays the -help text of the device-specific configuration program, if any. The frame buffer device can be specified using the -dev option, otherwise the default is used. Other fbconfig options are ignored. This is the default fbconfig option.

-list

Displays the pathnames of the installed frame buffer devices, the device model of each, and the configuration program that would be invoked for each device with the currently configured X server. Other fbconfig options are ignored.

Device File Name	Device Model	Config Program
-----	-----	-----
/dev/fbs/kfb0	XVR-2500	SUNWkfb_config
/dev/fbs/kfb1	XVR-2500	SUNWkfb_config
/dev/fbs/nfb0 [a b]	XVR-300	SUNWnfb_config
/dev/fbs/pfb0 [a b]	XVR-100	SUNWpfb_config

The following options are commonly supported by the device-specific configuration programs:

-defaults

Sets configuration options to their default values.

-file machine | system

Specifies which xorg.conf configuration file to use.

-prconf

Display the current configuration for the frame buffer and display device(s).

-predid [raw] [parsed]

Display the E-EDID (Enhanced Extended Display Identification Data) information obtained from the display device(s), which must be online, connected to the frame buffer. The output will be raw hexadecimal and/or human-readable text. The default is parsed.

-propt

Display the current software configuration.

-res \?

Display the video modes (resolutions) that can be used with the -res *video-mode* option.

Other supported options are determined by the device-specific configuration program:

device-specific-options

The syntax and descriptions of additional device-specific options are displayed in the -help output of fbconfig or the device-specific program. They are also contained in the man page for the device-specific program.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/graphics/fbconfig

另请参见 [fbconf_xorg\(1M\)](#), [attributes\(5\)](#)

Xserver(1), Xorg(1), Xsun(1), which are part of the OpenWindow, not the SunOS, man page collection.

Limitations Because of limitations in the m64 kernel driver and related software, fbconfig (with the -prconf option) is unable to distinguish between a current depth of 24 or 8+24. The -propt option returns the depth specified in the OWconfig file, which will be in effect following the next restart of the window system. The xwininfo utility, usually shipped in the package containing frame buffer software (such as SUNWxwplt), reports current depth of a specified window.

引用名	fbconf_xorg – configure frame buffer devices for Xorg
用法概要	<pre> /usr/lib/fbconfig/fbconf_xorg [-dev <i>device-file</i>] [-file <i>machine</i> <i>system</i> <i>config-path</i>] [-res <i>video-mode</i> [nocheck noconfirm]] [<i>device-specific-options</i>] [-defaults] [-prconf] [-predid [<i>raw</i>] [<i>parsed</i>]] [-propt] /usr/lib/fbconfig/fbconf_xorg [-dev <i>device-file</i>] [-prconf] [-propt] /usr/lib/fbconfig/fbconf_xorg [-dev <i>device-file</i>] [-help] [-res ?]</pre>
描述	The fbconf_xorg utility configures Frame Buffer devices and some of the X11 window system defaults for Xorg by updating the <code>xorg.conf</code> configuration file. Users should not normally need to run fbconf_xorg directly, as it is executed by fbconfig(1M) when the Xorg server is in use. The first fbconf_xorg command form shown in the SYNOPSIS section stores the specified option values in the <code>xorg.conf</code> configuration file. These settings are used to initialize the frame buffer device the next time the window system is run on that device. The persistence of these settings across window system sessions and system reboots is determined by the <code>xorg.conf</code> file. The second and third command forms display information, and do not alter the <code>xorg.conf</code> file. The presence of the <code>-help</code> and/or <code>-res ?</code> options causes any other option (not shown in the third form) to be ignored. You can configure only one frame buffer device at a time. To configure multiple devices, invoke the fbconf_xorg utility separately for each device. Only frame buffer device options can be specified with fbconf_xorg. Use the normal window system options to specify default depth (see svccfg(1M)), default visual class, and so forth. Specify these as device modifiers on the command line as specified in <code>Xserver(1)</code>. You can specify which <code>xorg.conf</code> file to open. By default, fbconf_xorg opens the machine-specific file, <code>/etc/X11/xorg.conf</code>. Use the <code>-file</code> option to specify an alternate file. For example, the system-global file, <code>/usr/lib/X11/xorg.conf</code>, can be opened instead. These standard <code>xorg.conf</code> files can be written only by the superuser or someone with the Desktop Configuration role.
选项	The following options are supported for all frame buffer devices: -defaults Sets configuration options for the specified device to their default values. This does not affect the <code>-res video mode</code> setting. See the device-specific portions of the DEFAULTS section below.

-dev *device-file*

Specifies the frame buffer device by either its full pathname or simple filename (for example, `/dev/fbs/efb0` or `efb0`). Pathnames of installed devices can be displayed using the `-list` option to [fbconfig\(1M\)](#). If the `-dev` option is omitted, the default device, `/dev/fb`, is used.

-file *machine | system | config-path*

Specifies which `xorg.conf` file to open. If `machine` is specified, the machine-specific `/etc/X11/xorg.conf` file is opened. If `system` is specified, the global `/usr/lib/X11/xorg.conf` file is opened. The absolute pathname of a configuration file can be used instead. If the specified file does not exist and is to be updated, it is created. The file system that contains the `xorg.conf` file must be writeable by someone with superuser-like privileges. This option has no effect unless other options are specified. The default is `machine`.

-help

Display the `fbconf_xorg` command line options that are supported in conjunction with the frame buffer device, along with a brief explanation of each option. The frame buffer device can be specified using the `-dev` option.

-prconf

Display the current configuration for the frame buffer device and attached display device(s). The frame buffer device can be specified using the `-dev` option.

The `-prconf` output might resemble:

Monitor/Resolution Information:

```

Monitor manufacturer: SUN
Product Code: 4
Serial Number: 12212555
Manufacture date: 2000, week 9
EDID Version: 1.1
Monitor dimensions: 36x29 cm
Default Gamma: 2.62
Monitor preferred resolution: SUNW_STD_1280x1024x60
Monitor supported resolutions from EDID: SUNW_STD_1280x1024x60,
SUNW_STD_1280x1024x76, 1152x900x66, VESA_STD_1280x1024x75,
VESA_STD_1280x1024x60, SUNW_STD_1152x900x66,
VESA_STD_720x400x70, VESA_STD_640x480x60,
VESA_STD_640x480x67, VESA_STD_640x480x72,
VESA_STD_640x480x75, VESA_STD_800x600x56,
VESA_STD_800x600x60, VESA_STD_800x600x72,
VESA_STD_800x600x75, VESA_STD_832x624x75,
VESA_STD_1024x768x60, VESA_STD_1024x768x70, 1024x768x75
Current resolution setting: FALLBACK_1152x900x66
```

FrameLock Configuration:

```
Slave Mode: Disabled
```

-predid [raw] [parsed]

Display the E-EDID (Enhanced Extended Display Identification Data) information obtained from the display device(s), which must be online, connected to the frame buffer. The frame buffer device can be specified using the `-dev` option. The output is raw hexadecimal and/or human-readable (parsed) text. The default is parsed.

The `-predid raw` output might resemble:

```
--- EDID Data for /dev/fbs/kfb0 ---

Block 0: EDID Base Block
0x00:  00 FF FF FF FF FF FF 00 04 43 06 F2 01 00 00 00
0x10:  01 11 01 04 0F 2B 20 78 2B 9C 68 A0 57 4A 9B 26
0x20:  12 48 4C FF FF 80 A9 59 A9 4F A9 4A A9 45 81 99
0x30:  81 80 61 59 45 59 48 3F 40 30 62 B0 32 40 40 C0
0x40:  13 00 AB 40 11 00 00 1E 00 00 00 FD 00 32 5A 1E
0x50:  6E 17 04 11 00 C8 90 00 50 3C 00 00 00 F7 00 0A
0x60:  F7 0F 03 87 C0 00 00 00 00 00 00 00 00 00 00 FC
0x70:  00 41 42 43 20 4C 43 44 32 31 0A 20 20 20 00 0B
```

-propt

Display all option settings for the frame buffer device, either as they currently are or as they are represented in the `xorg.conf` configuration file when `fbconf_xorg` completes. The device can be specified using the `-dev` option, and the file using the `-file` option.

The `-propt` output might resemble:

```
--- Graphics Configuration for /dev/fbs/efb0 ---

xorg.conf: machine -- /etc/X11/xorg.conf
Screen section: "efb0"
Device section: "efb0"
Monitor section: none

Video Mode: Not set

Screen Information:
DoubleWide: Disable
DoubleHigh: Disable
Clone: Disable
Offset/Overlap: [0, 0]
Outputs:      Direct

Visual Information:
Gamma Correction: Using default gamma value 2.22
```

-res ?

Display a list of video modes that can be used with the `-res video-mode` option.

The `?` argument might need to be escaped or placed in quotes (`\?`, `"?"`, or `'?'`), to protect it from misinterpretation by the shell.

The `-res ?` output might resemble:

Video modes accepted by the `-res` option:

```
AUTO                [1][2]
NONE                [1][2]
SUNW_STD_1920x1200x75
SUNW_STD_1920x1200x70
SUNW_DIG_1920x1200x60
SUNW_STD_1920x1080x72
SUNW_DIG_1920x1080x60
...
...
VESA_STD_640x480x75  [1]
VESA_STD_640x480x72  [1]
VESA_STD_640x480x60  [1]
```

[1] Resolution is supported by monitor

[2] Preferred resolution for monitor

Abbreviations such as “1280x1024x75” might also be used.

`-res video-mode` [`nocheck` | `noconfirm`]

Set the video mode for the display device that is connected to the frame buffer device.

A list of video modes can be displayed using the `-res ?` option.

The basic format of a video-mode is *widthxheightxrate*, where:

- *width* is the screen width in pixels.
- *height* is the screen height in pixels.
- *rate* is the vertical frequency of the screen refresh.

A video-mode argument might have an @ (at sign) instead of x preceding the refresh rate. For instance, 1280x1024x76 and 1280x1024@76 are equivalent.

A video-mode name might carry additional information, as with SUNW_STD_1280x1024x76.

The `-res` argument, `auto`, represents the video mode that is currently programmed into the device. The argument, `none`, is a synonym for `auto`.

Note that some video modes might be inappropriate for certain frame buffer devices and/or display devices.

The `-res` option accepts suboption keywords following the video-mode specification.

`nocheck`

The `nocheck` suboption causes the video-mode argument to be accepted, regardless of whether it is supported by the currently attached monitor, whether it is known within

the current configuration, and so forth. Note that using an unchecked, inappropriate video mode can leave the system without usable video output. This suboption is useful if a different monitor is to be connected to the frame buffer device. This suboption also implies `noconfirm`.

`noconfirm`

If the video-mode argument is unable to be validated, the default action is to display a warning message and ask the user whether to continue. The `noconfirm` suboption suppresses this confirmation request. This suboption is useful when `fbconf_xorg` is being run from a shell script.

The following device-specific options are supported for certain frame buffer devices. Unless specified otherwise, these options do not take effect until the user logs out and back in.

`-deflinear true | false`

This option selects the default X visual. Two types of visuals are supported, linear and nonlinear. Linear visuals are gamma corrected. Nonlinear visuals are not.

If the value of this option is `true`, the default visual is set to default depth 24 and the default class is `TrueColor` with gamma correction enabled. If `false`, a nonlinear visual that satisfies the other default visual selection options, such as the default depth and default class, is chosen as the default visual.

The `-deflinear`, `-defoverlay`, and `-deftransparent` options each select the default X visual. Only one of these might be enabled at a time. Enabling one causes the others to be disabled.

`-defoverlay true | false`

This option selects the default X visual. Some devices might provide an 8-bit `PseudoColor` visual whose pixels are disjoint from the rest of the visuals. This is called the overlay visual. Windows created in this visual does not damage windows created in other visuals. The converse, however, is not true: Windows created in other visuals damage overlay windows.

If the value of this option is `true`, the overlay visual is the default visual. The default depth is 8-bit and the default class is `PseudoColor`. If `false`, the non-overlay visual that satisfies the other default visual selection options, such as the default depth and the default class, is chosen as the default visual.

The `-deflinear`, `-defoverlay`, and `-deftransparent` options each select the default X visual. Only one of these might be enabled at a time. Enabling one causes the others to be disabled.

`-deftransparent true | false`

This option selects the default X visual. Some devices might provide an 8-bit `PseudoColor` visual whose pixels are disjoint from the rest of the visuals. This is called the overlay visual. Windows created in this visual does not damage windows created in other visuals.

If the value of this option is `true`, the overlay visual used as the default is a transparent overlay visual. A visual with transparency supports a colormap with 255 colors and one transparent pixel. The default depth is 8-bit and the default class is `PseudoColor`. If `false`, the nonoverlay visual that satisfies the other default visual selection options, such as the default depth and the default class, is chosen as the default visual.

The `-deflinear`, `-defoverlay`, and `-deftransparent` options each select the default X visual. Only one of these might be enabled at a time. Enabling one causes the others to be disabled.

`-doublehigh enable | disable`

Configures the two outputs of the frame buffer device into one vertical virtual display. The default is `disable`. The `-doublewide` and `-doublehigh` options are mutually exclusive. Enabling one causes the other to be disabled.

`-doublewide enable | disable`

Configures the two outputs of the frame buffer device into one horizontal virtual display. The default is `disable`. The `-doublewide` and `-doublehigh` options are mutually exclusive. Enabling one causes the other to be disabled.

`-g gamma-correction-value`

Sets the gamma correction value. All linear visuals provide gamma correction. The gamma correction value should be in the range, 0.1 to 10.0. The default is 2.22. This option can be used while the window system is running. Changing the gamma correction value affects all of the windows displayed by linear visuals.

`-gfile gamma-correction-file`

Loads the gamma correction table from the file specified by *gamma-correction-file*. This text file specifies the gamma correction values for the R, G, and B channels. Three consecutive values form an RGB triplet. For a `kfb` device, there must be exactly 256 RGB triplets. A value might be represented in hexadecimal, decimal, or octal format (for example, `0x3FF`, `1023`, or `01777`, respectively). Values are separated by one or more whitespace or new line characters. Comments begin with a hash sign character (`#`) and end at the end of the line.

You can load the gamma correction table with this option while the window system is running. The new gamma correction affects all the windows being displayed using the linear visuals. When gamma correction is done using a user-specified table, the gamma correction value (`-g`) is undefined. By default, the window system assumes a gamma correction value of 2.22 and loads the gamma table it creates corresponding to this value.

The following is an example of a *gamma-correction-file* file:

```
# Gamma Correction Table
0x00 0x00 0x00
0x01 0x01 0x01
0x02 0x02 0x02
... ..
```

```

    ...    ...    ...
    0xFF 0xFF 0xFF

```

-multisample *available | disable | forceon*

If set to *disable*, no multisample is possible. If set to *available*, multisample is possible but is selected on a per-window basis using a library interface. If set to *forceon*, all Sun OpenGL windows are rendered using multisampling. To query the number of samples used, specify the **-propt** option.

-offset *x-val y-val*

Adjusts the position of the specified stream by the value specified. This option is only implemented in **-doublewide** and **-doublehigh** modes. For **-doublewide**, use the *x-val* to position the rightmost stream. Negative is left (overlaps with the left stream). For **-doublehigh**, use the *y-val* to position the bottom stream. Negative is up (overlaps with top stream). The default is *[0, 0]*.

-samples *1 | 2 | 4 | 8 | 16*

Requests the number of samples to compute per display pixel. The requested number of samples per pixel is used if **-multisample** is not disabled and resources exist for the request. To query the number of samples used, specify the **-propt** option or run the **xglnfo** utility. The **xglnfo** utility can return the number of multisamples after you specify the option **-multisample available**. The default is 4.

-slave *disable | multiview*

If you set the **multiview** argument for the **-slave** option, the device synchronizes video with a master through the multiview genlock ribbon cable. The system should be powered off whenever connecting or disconnecting this cable. Both devices should be running the same resolution and the option should be issued when the window system is running. The default is *disable*.

Defaults

Certain options have implied default arguments. The default argument is used when the option is not present on the **fbconf_xorg** command line. For instance, a default argument for **-dev** is */dev/fb*.

Options that set configuration state do not have implied defaults. The **-res** option is one example. If a configuration option is omitted from the **fbconf_xorg** command line, the corresponding **xorg.conf** configuration setting remains unchanged. The exception is that if configuration options are mutually exclusive, setting one automatically unsets each of the others. An example is **-deflinear**, **-defoverlay**, and **-deftransparent**.

If a configuration setting is not present in the configuration file when the window system is run, a default value is used. For instance, the default state associated with **-res** is *auto*. A setting might not be present in the file, or the file itself might not exist, until **fbconf_xorg** has been invoked with the corresponding command line option.

The **-defaults** option sets the default values for most configuration settings.

Options and their defaults are shown below.

Device-independent defaults:

Option	Default Argument
-dev	/dev/fb
-file	machine
Option	Default State
-res	auto

示例

示例 1 Switching the Resolution of the Monitor Type

The following example sets the video mode for the monitor on the /dev/fbs/efb0 device to 1280 x 1024 at 76 Hz:

```
example% fbconf_xorg -dev efb0 -res 1280x1024x76
```

退出状态

The following exit values are returned:

- 0
Execution completed successfully.
- 1
Invalid command line usage.
- 2
An error occurred.

文件

/dev/fb	Symbolic link to the default frame buffer device.
/dev/fbs/efbnn	Device special file for an efb frame buffer
/usr/lib/fbconfig/SunModes_xorg.conf	Video mode definitions included in new configuration files.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/graphics/fbconfig/fbconfig-kfb
Interface Stability	Uncommitted

另请参见

[fbconfig\(1M\)](#), [svccfg\(1M\)](#), [attributes\(5\)](#), [efb\(7D\)](#), [fbio\(7I\)](#)

See the Xorg(1) and Xserver(1) man pages in the X Server man page collection.

引用名

用法概要

fcinfo, fcdm – Fibre Channel HBA Port Command Line Interface

fcinfo hba-port [-lite] [HBA_port_WWN]...

fcadm hba-port [-lite] [HBA_port_WWN]...

fcinfo remote-port [-ls] [-p HBA_port_WWN] [REMOTE_port_WWN]...

fcadm remote-port [-ls] [-p HBA_port_WWN] [REMOTE_port_WWN]...

fcinfo logical-unit | lu [-v] [OS device_path]

fcadm logical-unit | lu [-v] [OS device_path]

fcadm create-npiv-port -p Virtual_Port_WWN [-n Virtual_Node_WWN] PhysicalPort_port_WWN

fcadm delete-npiv-port -p Virtual_Port_WWN [-n Virtual_Node_WWN] PhysicalPort_port_WWN

fcadm create-fcoe-port [-tf] -p Port_WWN [-n Node_WWN] MAC_Interface

fcadm delete-fcoe-port MAC_Interface

fcadm list-fcoe-ports MAC_Interface

fcadm force-lip Port_WWN

fcinfo [-V]

fcadm [-V]

fcinfo [-?]

fcadm [-?]

描述

fcinfo and fcdm are command line interfaces that collect administrative information on fibre channel host bus adapter (HBA) ports on a host. They also collect data on any fibre channel targets that might be connected to those ports in a Storage Area Network (SAN).

A port can be either in initiator mode or target mode. The same FC attribute information and remote port information are applied to both the initiator and the target modes port except that SCSI-related operations described below do not apply for the target mode port.

SUBCOMMANDS

The following subcommands are supported by both fcinfo and fcdm:

hba-port

Lists information for the HBA port referenced by the specified HBA_port_WWN. If HBA_port_WWN is not specified, all initiator mode and target mode fibre channel HBA ports on the host will be listed.

- `remote-port` Lists the *remote-port* information for those remote ports that are specified. If no *REMOTE_port_WWN* is specified, all remote ports that are visible through *HBA_port_WWN* are listed.
- `logical-unit | lu` Lists the logical unit referenced by the specified *device_path*. If *device_path* is not specified, all fibre channel logical units will be listed. This subcommand applies only to the initiator mode.

The following subcommands, which administer N_Port ID Virtualization (NPIV) HBA ports, are supported only by `fcadm`:

`create-npiv-port`

Create an NPIV virtual port on the specified adapter.

`delete-npiv-port`

Delete an NPIV virtual port. This delete only ports create by `fcadm`.

The following subcommands, which administer Fibre Channel over Ethernet (FCoE) ports, are supported only by `fcadm`:

`create-fcoe-port`

Creates a FCoE port associated with the given MAC interface and the given *Port_WWN/Node_WWN*. If `-p` or `-n` is not specified, the Port WWN or Node WWN will be generated automatically. If the specified MAC interface does not support multiple unicast address, no FCoE port will be created and you receive a message indicating that you can add an `-f` option to force enabling promiscuous mode on the specified MAC interface to enable FCoE.

`delete-fcoe-port`

Deletes the FCoE port associated with the specified MAC interface.

`list-fcoe-ports`

List information of FCoE ports.

`force-lip`

Force the link to reinitialize. When issuing this subcommand on the target port side, causes a reset of the target port. When issuing it from the host port side, resets the host port. When an FC switch is connected, other FC ports in the SAN will get a RSCN (Remote State Change Notification). Furthermore, other initiators will always rediscover the port after this, and the FC login session will be established or reused. Also, I/Os will be continued. This command is disruptive to I/Os, but it is not destructive, as it does not cause any data loss.

选项

The following options are supported:

`-e, --fcoe`

Lists the information for all FCoE ports.

-f, --fcoe-promiscuous

Used with `create-fcoe-port` to create an FCoE port on a MAC interface associated with a NIC that does not support multiple unicast address. Promiscuous mode will be enabled for the specified MAC interface.

-l, --linkstat

Lists the link error statistics information for the port referenced by the specified *HBA_port_WWN* or *REMOTE_port_WWN*.

-n *HBA_node_WWN*, --node *HBA_node_WWN*

When used with NPIV options, specify a virtual node WWN. If used with `create-npiv-port`, it can be omitted, and a random based WWN will be used. It is mandatory for `delete-npiv-port`.

When used with `create-fcoe-port` subcommand, specify the node WWN for the FCoE port. It can be omitted, in which case a WWN will be generated based on the MAC address of the specified MAC interface.

-p *HBA_port_WWN*, --port *HBA_port_WWN*

Retrieve remote port information from the *HBA_port_WWN* of the local HBA port on the host. When used with the `remote-port` subcommand, the `-p` option is mandatory.

When used with NPIV options, specify a virtual port WWN. If used with `create-npiv-port`, it can be omitted, and a random based WWN will be used. It is mandatory for `delete-npiv-port`.

When used with `create-fcoe-port` subcommand, specify the port WWN for the FCoE port. It can be omitted, in which case a WWN will be generated based on the MAC address of the specified MAC interface.

-s, --scsi-target

Lists the SCSI target information for all remote ports the user has asked for. The `-p`, --port option must always be specified and must be a valid HBA port on the host. This HBA port will be used as the initiator for which to retrieve the SCSI level target information. Note that this will only function on remote port fibre channel World-Wide Names that support an FC4 type of SCSI. This option applies only to an initiator mode port.

-t [*HBA_node_WWN*], --target [*HBA_node_WWN*]

Lists the information for the port with the target mode referenced by the specified *HBA_node_WWN* or all target mode ports discovered in the host.

When used with `create-fcoe-port`, create a FCoE target mode port.

-v, --verbose

When used with the `logical-unit` subcommand, the `-v` displays additional information for the logical unit, including SCSI vendor and product information and device type as well as the FC World-Wide names for the local and remote fibre channel ports to which this device is attached.

-V, --version
Displays the version information.

-, --help
Displays the usage information.

示例

示例1 Listing All HBA Ports

The following command lists all initiator mode fibre channel HBA ports on the host:

```
# fcinfo hba-port
```

```
HBA Port WWN: 210000e08b074cb5
  Port Mode: Initiator
  OS Device Name: /dev/cfg/c1
  Manufacturer: QLogic Corp.
  Model: 375-3108-xx
  Firmware Version: 3.3.116
  FCode/BIOS Version: 1.13.08
  Serial Number: not available
  Driver Name: qlc
  Driver Version: 20070212-2.19
  Type: N-port
  State: online
  Supported Speeds: 1Gb 2Gb
  Current Speed: 2Gb
  Node WWN: 200000e08b074cb5
NPIV Port List:
  Virtual Port 1:
    Port WWN: 200000e08b074cb1
    Node WWN: 200000e08b074cb3
HBA Port WWN: 210100e08b274cb5
  Port Mode: Initiator
  OS Device Name: /dev/cfg/c2
  Manufacturer: QLogic Corp.
  Model: 375-3108-xx
  Firmware Version: 3.3.116
  FCode/BIOS Version: 1.13.08
  Serial Number: not available
  Driver Name: qlc
  Driver Version: 20070212-2.19
  Type: N-port
  State: online
  Supported Speeds: 1Gb 2Gb
  Current Speed: 2Gb
  Node WWN: 200100e08b274cb5
HBA Port WWN: 210000e08b072ab5
  Port Mode: Initiator
  OS Device Name: /dev/cfg/c3
```

示例1 Listing All HBA Ports (续)

```

Manufacturer: QLogic Corp.
Firmware Version: 3.3.116
FCode/BIOS Version: 1.13.08
Model: 375-3108-xx
Serial Number: not available
Driver Name: qlc
Driver Version: 20070212-2.19
Type: L-port
State: online
Supported Speeds: 1Gb 2Gb
Current Speed: 2Gb
Node WWN: 200000e08b072ab5
HBA Port WWN: 210100e08b272ab5
Port Mode: Initiator
OS Device Name: /dev/cfg/c4
Manufacturer: QLogic Corp.
Model: 375-3108-xx
Firmware Version: 3.3.116
FCode/BIOS Version: 1.13.08
Serial Number: 0402F00-0549112808
Driver Name: qlc
Driver Version: 20070212-2.19
Type: N-port
State: online
Supported Speeds: 1Gb 2Gb
Current Speed: 2Gb
Node WWN: 200100e08b272ab5

```

示例2 Listing Target Mode HBA Ports

The following command lists all target mode fibre channel HBA ports on the host:

```

# fcinfo hba-port -t
HBA Port WWN: 210100e08bb09221
Port Mode: Target
Port ID: 10700
OS Device Name: Not Applicable
Manufacturer: QLogic Corp.
Model: d30ac7e0
Firmware Version: 4.0.109
FCode/BIOS Version: N/A
Type: F-port
State: online
Supported Speeds: not established
Current Speed: 2Gb
Node WWN: 200100e08bb09221
HBA Port WWN: 210000e08b909221

```

示例 2 Listing Target Mode HBA Ports (续)

```
Port Mode: Target
Port ID: 10900
OS Device Name: Not Applicable
Manufacturer: QLogic Corp.
Model: d37ad1e0
Firmware Version: 4.0.109
FCode/BIOS Version: N/A
Type: F-port
State: online
Supported Speeds: not established
Current Speed: 2Gb
Node WWN: 200000e08b909221
HBA Port WWN: 200000144fc2d508
Port Mode: Target
Port ID: 9a0025
OS Device Name: Not Applicable
Manufacturer: Sun Microsystems, Inc.
Model: FCoE Virtual FC HBA
Firmware Version: N/A
FCode/BIOS Version: N/A
Serial Number: N/A
Driver Name: COMSTAR FCOET
Driver Version: 1.0
Type: F-port
State: online
Supported Speeds: 1Gb 10Gb
Current Speed: 10Gb
Node WWN: 100000144fc2d508
HBA Port WWN: 200000144fc2d509
Port Mode: Target
Port ID: 9a0023
OS Device Name: Not Applicable
Manufacturer: Sun Microsystems, Inc.
Model: FCoE Virtual FC HBA
Firmware Version: N/A
FCode/BIOS Version: N/A
Serial Number: N/A
Driver Name: COMSTAR FCOET
Driver Version: 1.0
Type: F-port
State: online
Supported Speeds: 1Gb 10Gb
Current Speed: 10Gb
Node WWN: 100000144fc2d509
```

示例 3 Listing HBA Ports and Link Statistics

The following command lists information for the HBA ports and the link statistics for those ports:

```
# fcinfo hba-port -l 210000e08b074cb5 210100e08b274cb5
```

```
HBA Port WWN: 210000e08b074cb5
  Port Mode: Initiator
  OS Device Name: /dev/cfg/c1
  Manufacturer: QLogic Corp.
  Model: 375-3108-xx
  Firmware Version: 3.3.116
  FCode/BIOS Version: 1.13.08
  Serial Number: 0402F00-0549112808
  Driver Name: qlc
  Driver Version: 20070212-2.19
  Type: N-port
  State: online
  Supported Speeds: 1Gb 2Gb
  Current Speed: 2Gb
  Node WWN: 200000e08b074cb5
  Link Error Statistics:
    Link Failure Count: 0
    Loss of Sync Count: 0
    Loss of Signal Count: 0
    Primitive Seq Protocol Error Count: 0
    Invalid Tx Word Count: 0
    Invalid CRC Count: 0

HBA Port WWN: 210100e08b274cb5
  Port Mode: Initiator
  OS Device Name: /dev/cfg/c2
  Manufacturer: QLogic Corp.
  Model: 375-3108-xx
  Firmware Version: 3.3.116
  FCode/BIOS Version: 1.13.08
  Serial Number: 0402F00-0549112808
  Driver Name: qlc
  Driver Version: 20070212-2.19
  Type: N-port
  State: online
  Supported Speeds: 1Gb 2Gb
  Current Speed: 2Gb
  Node WWN: 200100e08b274cb5
  Link Error Statistics:
    Link Failure Count: 0
    Loss of Sync Count: 0
    Loss of Signal Count: 0
    Primitive Seq Protocol Error Count: 0
```

示例 3 Listing HBA Ports and Link Statistics (续)

```
Invalid Tx Word Count: 0
Invalid CRC Count: 0
```

示例 4 Listing All Remote Ports

The following command lists all remote ports that are visible through the given HBA port:

```
# fcinfo remote-port -p 210100e08b274cb5
```

```
Remote Port WWN: 50020f230000b4af
  Active FC4 Types: SCSI
  SCSI Target: yes
  Port Symbolic Name: unknown
  Node WWN: 50020f200000b4af
Remote Port WWN: 210000e08b07daa6
  Active FC4 Types: SCSI
  SCSI Target: no
  Port Symbolic Name: unknown
  Node WWN: 200000e08b07daa6
Remote Port WWN: 20030003ba27c788
  Active FC4 Types: SCSI
  SCSI Target: yes
  Port Symbolic Name: unknown
  Node WWN: 10000003ba27c788
Remote Port WWN: 210000e08b096a60
  Active FC4 Types: SCSI,IP
  SCSI Target: no
  Port Symbolic Name: unknown
  Node WWN: 200000e08b096a60
```

示例 5 Listing Remote Ports and Link Statistics

The following command lists information for the remote ports and the link statistics for those ports:

```
# fcinfo remote-port -l -p 210100e08b272ab5
```

```
Remote Port WWN: 210100e08b296a60
  Active FC4 Types: SCSI,IP
  SCSI Target: no
  Port Symbolic Name: unknown
  Node WWN: 200100e08b296a60
  Link Error Statistics:
    Link Failure Count: 0
    Loss of Sync Count: 0
    Loss of Signal Count: 0
    Primitive Seq Protocol Error Count: 0
    Invalid Tx Word Count: 0
```


示例 5 Listing Remote Ports and Link Statistics (续)

```

Invalid CRC Count: 0
Remote Port WWN: 20030003ba27d56d
Active FC4 Types: SCSI
SCSI Target: yes
Port Symbolic Name: unknown
Node WWN: 10000003ba27d56d
Link Error Statistics:
    Link Failure Count: 0
    Loss of Sync Count: 4765165
    Loss of Signal Count: 4765165
    Primitive Seq Protocol Error Count: 0
    Invalid Tx Word Count: 0
    Invalid CRC Count: 0
Remote Port WWN: 210100e08b27f7a6
Active FC4 Types: SCSI
SCSI Target: no
Port Symbolic Name: unknown
Node WWN: 200100e08b27f7a6
Link Error Statistics:
    Link Failure Count: 0
    Loss of Sync Count: 0
    Loss of Signal Count: 0
    Primitive Seq Protocol Error Count: 0
    Invalid Tx Word Count: 0
    Invalid CRC Count: 0
Remote Port WWN: 50020f230000b897
Active FC4 Types: SCSI
SCSI Target: yes
Port Symbolic Name: unknown
Node WWN: 50020f200000b897
Link Error Statistics:
    Link Failure Count: 0
    Loss of Sync Count: 7
    Loss of Signal Count: 7
    Primitive Seq Protocol Error Count: 0
    Invalid Tx Word Count: 0
    Invalid CRC Count: 0
Remote Port WWN: 210100e08b27daa6
Active FC4 Types: SCSI
SCSI Target: no
Port Symbolic Name: unknown
Node WWN: 200100e08b27daa6
Link Error Statistics:
    Link Failure Count: 0
    Loss of Sync Count: 0
    Loss of Signal Count: 0

```

示例 5 Listing Remote Ports and Link Statistics (续)

```

        Primitive Seq Protocol Error Count: 0
        Invalid Tx Word Count: 0
        Invalid CRC Count: 0
Remote Port WWN: 210000e08b074cb5
    Active FC4 Types: SCSI,IP
    SCSI Target: no
    Port Symbolic Name: unknown
    Node WWN: 200000e08b074cb5
    Link Error Statistics:
        Link Failure Count: 0
        Loss of Sync Count: 0
        Loss of Signal Count: 0
        Primitive Seq Protocol Error Count: 0
        Invalid Tx Word Count: 0
        Invalid CRC Count: 0
Remote Port WWN: 210100e08b296060
    Active FC4 Types: SCSI
    SCSI Target: no
    Port Symbolic Name: unknown
    Node WWN: 200100e08b296060
    Link Error Statistics:
        Link Failure Count: 0
        Loss of Sync Count: 0
        Loss of Signal Count: 0
        Primitive Seq Protocol Error Count: 0
        Invalid Tx Word Count: 0
        Invalid CRC Count: 0

```

示例 6 Listing All SCSI Targets and Link Statistics

The following command lists all remote ports as well as the link statistics and *scsi-target* information:

```
# fcinfo remote-port -sl -p 210100e08b272ab5
```

```

Remote Port WWN: 210100e08b296a60
    Active FC4 Types: SCSI,IP
    SCSI Target: no
    Port Symbolic Name: unknown
    Node WWN: 200100e08b296a60
    Link Error Statistics:
        Link Failure Count: 0
        Loss of Sync Count: 0
        Loss of Signal Count: 0
        Primitive Seq Protocol Error Count: 0
        Invalid Tx Word Count: 0
        Invalid CRC Count: 0

```

示例 6 Listing All SCSI Targets and Link Statistics (续)

```

Remote Port WWN: 20030003ba27d56d
  Active FC4 Types: SCSI
  SCSI Target: yes
  Port Symbolic Name: unknown
  Node WWN: 10000003ba27d56d
  Link Error Statistics:
    Link Failure Count: 0
    Loss of Sync Count: 4765165
    Loss of Signal Count: 4765165
    Primitive Seq Protocol Error Count: 0
    Invalid Tx Word Count: 0
    Invalid CRC Count: 0
  LUN: 0
    Vendor: SUN
    Product: T4
    OS Device Name: /dev/rdisk/c4t20030003BA27D56Dd0s2
  LUN: 1
    Vendor: SUN
    Product: T4
    OS Device Name: /dev/rdisk/c4t20030003BA27D56Dd1s2
Remote Port WWN: 210100e08b27f7a6
  Active FC4 Types: SCSI
  SCSI Target: no
  Port Symbolic Name: unknown
  Node WWN: 200100e08b27f7a6
  Link Error Statistics:
    Link Failure Count: 0
    Loss of Sync Count: 0
    Loss of Signal Count: 0
    Primitive Seq Protocol Error Count: 0
    Invalid Tx Word Count: 0
    Invalid CRC Count: 0
Remote Port WWN: 50020f230000b897
  Active FC4 Types: SCSI
  SCSI Target: yes
  Port Symbolic Name: unknown
  Node WWN: 50020f200000b897
  Link Error Statistics:
    Link Failure Count: 0
    Loss of Sync Count: 7
    Loss of Signal Count: 7
    Primitive Seq Protocol Error Count: 0
    Invalid Tx Word Count: 0
    Invalid CRC Count: 0
  LUN: 0
    Vendor: SUN

```

示例 6 Listing All SCSI Targets and Link Statistics (续)

```
Product: T300
OS Device Name: Unknown
LUN: 1
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d1s2
LUN: 2
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d2s2
LUN: 3
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d3s2
LUN: 4
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d4s2
LUN: 5
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d5s2
LUN: 6
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d6s2
LUN: 7
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d7s2
LUN: 8
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d8s2
LUN: 9
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d9s2
LUN: 10
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d10s2
LUN: 11
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d11s2
```

示例 6 Listing All SCSI Targets and Link Statistics (续)

```

LUN: 12
  Vendor: SUN
  Product: T300
  OS Device Name: /dev/rdisk/c4t50020F230000B897d12s2
LUN: 13
  Vendor: SUN
  Product: T300
  OS Device Name: /dev/rdisk/c4t50020F230000B897d13s2
LUN: 14
  Vendor: SUN
  Product: T300
  OS Device Name: /dev/rdisk/c4t50020F230000B897d14s2
LUN: 15
  Vendor: SUN
  Product: T300
  OS Device Name: /dev/rdisk/c4t50020F230000B897d15s2
Remote Port WWN: 210100e08b27daa6
  Active FC4 Types: SCSI
  SCSI Target: no
  Port Symbolic Name: unknown
  Node WWN: 200100e08b27daa6
  Link Error Statistics:
    Link Failure Count: 0
    Loss of Sync Count: 0
    Loss of Signal Count: 0
    Primitive Seq Protocol Error Count: 0
    Invalid Tx Word Count: 0
    Invalid CRC Count: 0
Remote Port WWN: 210000e08b074cb5
  Active FC4 Types: SCSI,IP
  SCSI Target: no
  Port Symbolic Name: unknown
  Node WWN: 200000e08b074cb5
  Link Error Statistics:
    Link Failure Count: 0
    Loss of Sync Count: 0
    Loss of Signal Count: 0
    Primitive Seq Protocol Error Count: 0
    Invalid Tx Word Count: 0
    Invalid CRC Count: 0
Remote Port WWN: 210100e08b296060
  Active FC4 Types: SCSI
  SCSI Target: no
  Port Symbolic Name: unknown
  Node WWN: 200100e08b296060
  Link Error Statistics:

```

示例 6 Listing All SCSI Targets and Link Statistics (续)

```

Link Failure Count: 0
Loss of Sync Count: 0
Loss of Signal Count: 0
Primitive Seq Protocol Error Count: 0
Invalid Tx Word Count: 0
Invalid CRC Count: 0

```

示例 7 Listing SCSI Target Information

The following command lists all remote ports as well as the *scsi-target* information:

```
# fcinfo remote-port -s -p 210100e08b272ab5
```

```

Remote Port WWN: 210100e08b296a60
  Active FC4 Types: SCSI,IP
  SCSI Target: no
  Port Symbolic Name: unknown
  Node WWN: 200100e08b296060
Remote Port WWN: 20030003ba27d56d
  Active FC4 Types: SCSI
  SCSI Target: yes
  Port Symbolic Name: unknown
  Node WWN: 10000003ba27d56d
  LUN: 0
    Vendor: SUN
    Product: T4
    OS Device Name: /dev/rdisk/c4t20030003BA27D56Dd0s2
  LUN: 1
    Vendor: SUN
    Product: T4
    OS Device Name: /dev/rdisk/c4t20030003BA27D56Dd1s2
Remote Port WWN: 210100e08b27f7a6
  Active FC4 Types: SCSI
  SCSI Target: no
  Port Symbolic Name: unknown
  Node WWN: 200100e08b27f7a6
Remote Port WWN: 50020f230000b897
  Active FC4 Types: SCSI
  SCSI Target: yes
  Port Symbolic Name: unknown
  Node WWN: 50020f200000b897
  LUN: 0
    Vendor: SUN
    Product: T300
    OS Device Name: Unknown
  LUN: 1
    Vendor: SUN

```

示例7 Listing SCSI Target Information (续)

```
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d1s2
LUN: 2
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d2s2
LUN: 3
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d3s2
LUN: 4
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d4s2
LUN: 5
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d5s2
LUN: 6
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d6s2
LUN: 7
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d7s2
LUN: 8
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d8s2
LUN: 9
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d9s2
LUN: 10
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d10s2
LUN: 11
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d11s2
LUN: 12
Vendor: SUN
Product: T300
OS Device Name: /dev/rdisk/c4t50020F230000B897d12s2
```

示例 7 Listing SCSI Target Information (续)

```

LUN: 13
  Vendor: SUN
  Product: T300
  OS Device Name: /dev/rdisk/c4t50020F230000B897d13s2
LUN: 14
  Vendor: SUN
  Product: T300
  OS Device Name: /dev/rdisk/c4t50020F230000B897d14s2
LUN: 15
  Vendor: SUN
  Product: T300
  OS Device Name: /dev/rdisk/c4t50020F230000B897d15s2
Remote Port WWN: 210100e08b27daa6
  Active FC4 Types: SCSI
  SCSI Target: no
  Port Symbolic Name: unknown
  Node WWN: 200100e08b27daa6
Remote Port WWN: 210000e08b074cb5
  Active FC4 Types: SCSI,IP
  SCSI Target: no
  Port Symbolic Name: unknown
  Node WWN: 200000e08b074cb5
Remote Port WWN: 210100e08b296060
  Active FC4 Types: SCSI
  SCSI Target: no
  Port Symbolic Name: unknown
  Node WWN: 200100e08b296060

```

示例 8 Listing the Logical Unit

The following command lists the logical unit:

```
# fcinfo logical-unit
```

```

/dev/rdisk/c0t600C0FF000000000036223AE73EB705d0s2
/dev/rdisk/c0t600C0FF0000000000362272539E5B03d0s2
/dev/rmt/0n

```

示例 9 Displaying Additional Information for the Logical Unit

The following command displays additional information about the logical unit using the -v option for device /dev/rmt/0n:

```
# fcinfo lu -v /dev/rmt/0n
```

```

OS Device Name: /dev/rmt/0n
  HBA Port WWN: 210000e07c030b91
    Remote Port WWN: 21010003b7830a6

```


示例 9 Displaying Additional Information for the Logical Unit (续)

```

LUN: 0
Vendor: STK
Product: 9940A
Device Type: Tape device

```

The following command displays additional information about the logical unit using the -v option for device /dev/rdisk/c0t600C0FF0000000000036223AE73EB705d0s2

```

# fcinfo logical-unit -v \
/dev/rdisk/c0t600C0FF0000000000036223AE73EB705d0s2

OS Device Name: /dev/rdisk/c0t600C0FF0000000000036223AE73EB705d0s2
HBA Port WWN: 210100e08b27a8a1
Remote Port WWN: 256000c0ffc03622
LUN: 0
Remote Port WWN: 216000c0ff803622
LUN: 0
HBA Port WWN: 210000e08b07a8a1
Remote Port WWN: 256000c0ffc03622
LUN: 0
Remote Port WWN: 216000c0ff803622
LUN: 0
Vendor: SUN
Product: StorEdge 3510
Device Type: Disk device

```

示例 10 Adding an NIPV Port

The following command adds an NPIV port to the HBA with a port name:

```
210000e08b170f1c
```

The NPIV port has a port name of 2000000000000001 and a node name of 2100000000000001.

```
# fcadm create-npiv-port -p 2000000000000001 -n 2100000000000001 \
210000e08b170f1c
```

```

Created NPIV Port:
HBA Port WWN: 2000000000000001
Node WWN: 2100000000000001
Physical HBA Port WWN: 210000e08b170f1c

```

示例 11 Adding an NIPV Port with Random WWN

The following command adds an NPIV port to the HBA with a randomly assigned port name of 210000e08b170f1c.

示例 11 Adding an NIPV Port with Random WWN (续)

```
# fcadm create-npiv-port 210000e08b170f1c

Created NPIV Port:
  HBA Port WWN: 2038295824942801
  Node WWN: 2100392849248001
  Physical HBA Port WWN: 210000e08b170f1c
```

示例 12 Deleting an NIPV Port

The following command deletes an NPIV port.

```
# fcadm delete-npiv-port -p 2000000000000001 -w 2100000000000001 \
210000e08b170f1c
```

示例 13 Creating an FCoE Target Port

The following command creates an FCoE port associated with the MAC interface nxge0.

```
# fcadm create-fcoe-port -t nxge0
```

示例 14 Deleting an FCoE Port

The following command deletes the FCoE port associated with the MAC interface nxge0.

```
# fcadm delete-fcoe-port -t nxge0
```

示例 15 Listing Information for an FCoE Port

The following command lists information for FCoE ports.

```
# fcadm list-fcoe-ports
HBA Port WWN: 200000144fc2d508
  Port Type: Target
  MAC Name: nxge0
  MTU Size: 9194
  MAC Factory Address: 00144fc2d508
  MAC Current Address: 00144fc2d508
HBA Port WWN: 200000144fc2d509
  Port Type: Target
  MAC Name: nxge1
  MTU Size: 9194
  MAC Factory Address: 00144fc2d509
  MAC Current Address: 00144fc2d509
```

示例 16 Reinitializing the Link of an FC Port

The following command forces the link connected with the port 200000144fc2d508 to reinitialize.

```
# fcadm force-lip 200000144fc2d508
```

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/storage/fc-utilities
Interface Stability	Committed

另请参见 [attributes\(5\)](#)

Error Messages Errors that can occur in addition to the errors normally associated with system administration commands:

- *HBA_port_WWN*: not found
- *Remote_port_WWN*: not found
- *HBA_port_WWN*: NPIV not supported on this HBA

引用名 fdetach – detach a name from a STREAMS-based file descriptor

用法概要 fdetach *path*

描述 The fdetach command detaches a STREAMS-based file descriptor from a name in the file system. Use the *path* operand to specify the path name of the object in the file system name space, which was previously attached. See [fattach\(3C\)](#).

The user must be the owner of the file or a user with the appropriate privileges. All subsequent operations on *path* will operate on the underlying file system entry and not on the STREAMS file. The permissions and status of the entry are restored to the state they were in before the STREAMS file was attached to the entry.

操作数 The following operands are supported:

path Specifies the path name of the object in the file system name space, which was previously attached.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [fattach\(3C\)](#), [fdetach\(3C\)](#), [attributes\(5\)](#), [streamio\(7I\)](#)

《STREAMS Programming Guide》

引用名

fdisk – create or modify fixed disk partition table

用法概要

```
fdisk [-o offset] [-s size] [-P fill_patt] [-S geom_file]
      [-w | -r | -d | -n | -I | -B | -t | -T | -g | -G | -R | -E]
      [--F fdisk_file] [ [-v] -W {fdisk_file | -}]
      [-h] [-b masterboot]
      [-A id : act : bhead : bsect : bcyl : ehhead : esect :
          ecyl : rsect : numsect]
      [-D id : act : bhead : bsect : bcyl : ehhead : esect :
          ecyl : rsect : numsect] rdevice
```

描述

This command is used to do the following:

- Create and modify an fdisk partition table on x86 systems
- Create and modify an fdisk partition table on removable media on SPARC or x86 systems
- Install the master boot record that is put in the first sector of the fixed disk on x86 systems only

This table is used by the first-stage bootstrap (or firmware) to identify parts of the disk reserved for different operating systems, and to identify the partition containing the second-stage bootstrap (the *active* Solaris partition). The *rdevice* argument must be used to specify the raw device associated with the fixed disk, for example, `/dev/rdisk/c0t0d0p0`.

The program can operate in three different modes. The first is interactive mode. In interactive mode, the program displays the partition table as it exists on the disk, and then presents a menu allowing the user to modify the table. The menu, questions, warnings, and error messages are intended to be self-explanatory.

In interactive mode, if there is no partition table on the disk, the user is given the options of creating a default partitioning or specifying the initial table values. The default partitioning allocates the entire disk for the Solaris system and makes the Solaris system partition active. In either case, when the initial table is created, fdisk also writes out the first-stage bootstrap (x86 only) code along with the partition table. In this mode, (x86 only) when creating an entry for a non-EFI partition on a disk that is larger than 2 TB (terabytes), fdisk warns that the maximum size of the partition is 2 TB. Under these conditions percentages displayed by fdisk are based on 2 TB.

The second mode of operation is used for automated entry addition, entry deletion, or replacement of the entire fdisk table. This mode can add or delete an entry described on the command line. In this mode the entire fdisk table can be read in from a file replacing the original table. fdisk can also be used to create this file. There is a command line option that will cause fdisk to replace any fdisk table with the default of the whole disk for the Solaris system.

The third mode of operation is used for disk diagnostics. In this mode, a section of the disk can be filled with a user-specified pattern and mode sections of the disk can also be read or written.

注 – The third mode of operation is not currently supported for extended partitions

When `fdisk` creates a partition, the space is allocated in the `fdisk` partition table, but the allocated disk space is not initialized. `newfs(1M)` is required to create and write file system metadata to the new partition, and `format(1M)` is required to write the VTOC or EFI/GPT metadata.

Menu Options

The menu options for interactive mode given by the `fdisk` program are:

Create a partition

This option allows the user to create a new partition. The maximum number of partitions is 4. The program will ask for the type of the partition (SOLARIS, MS-DOS, UNIX, or other). It will then ask for the size of the partition as a percentage of the disk. The user may also enter the letter `c` at this point, in which case the program will ask for the starting cylinder number and size of the partition in cylinders. If a `c` is not entered, the program will determine the starting cylinder number where the partition will fit. In either case, if the partition would overlap an existing partition or will not fit, a message is displayed and the program returns to the original menu.

Change Active (Boot from) partition

This option allows the user to specify the partition where the first-stage bootstrap will look for the second-stage bootstrap, otherwise known as the *active* partition.

Delete a partition

This option allows the user to delete a previously created partition. Note that this will destroy all data in that partition.

Change between Solaris and Solaris2 Partition IDs

This option allows the user to switch between the current `fdisk` operating system partition identifier and the previous one. This does not affect any data in the disk partition and is provided for compatibility with older software.

Edit/View extended partitions

This option provides the extended partition menu to the user. Use the extended partition menu to add and delete logical drives, change the `sysid` of the logical drives, and display logical drive information. To commit the changes made in the extended partition, you must return to the main menu using the extended partition submenu option `r`. There is also an option to display the list of options that the extended partition submenu supports. Given below is the list:

a Add a logical drive.

Use this submenu option to add a logical drive. There are three pieces of information that are required: The beginning cylinder, the size (in cylinders or in human readable form - KB, MB, or GB), and the partition ID. While specifying the partition ID, there is an option (`I`) that you can use to list the supported partitions.

d Delete a logical drive.

Use this submenu option to delete a logical drive. The only input required is the number of the logical drive that is to be deleted.

h Display the help menu.

This submenu option displays the supported operations in the extended partition submenu.

i Change the id of the logical drive.

Use this submenu option to change the system ID of the existing logical drives. A list of supported system IDs is displayed when you use the I option when in this submenu.

p Display the logical drive layout.

Displays the logical drive information to stdout. This output reflects any changes made during the current run of the `fdisk` program. The changes are not committed to the disk until return to the main menu (using the submenu `r`) and choose the option to commit the changes to the disk.

r Return to the main `fdisk` menu.

Exit the extended partition submenu and return to the main menu.

Note the dynamic nature of the numbering of extended partitions. For example, consider a Solaris system with the partitions `p1`, `p2`, `p3`, and `p4`. Following creation of an extended partition, the same system has a logical device node, `p5`, and successive nodes numbered consecutively up to a maximum of `p36`. If one logical drive is deleted, say, `p8`, then all nodes following `p8` (`p9` up to `p36`) move up one in the list of partitions, so that `p9` becomes `p8`, `p10` becomes `p9`, and so forth.

Use the following options to include your modifications to the partition table at this time or to cancel the session without modifying the table:

Exit This option writes the new version of the table created during this session with `fdisk` out to the fixed disk, and exits the program.

Cancel This option exits without modifying the partition table.

选项

The following options apply to `fdisk`:

-A *id:act:bhead:bsect:bcyl:ehhead:esect:ecyl:rsect:numsect*

Add a partition as described by the argument (see the `-F` option below for the format). Use of this option will zero out the VTOC on the Solaris partition if the `fdisk` table changes.

-b *master_boot*

Specify the file `master_boot` as the master boot program. The default master boot program is `/usr/lib/fs/ufs/mboot`.

-B

Default to one Solaris partition that uses the whole disk. On an x86 machine, if the disk is larger than 2 TB (terabytes), the default size of the Solaris partition will be limited to 2 TB.

-d

Turn on verbose *debug* mode. This will cause *fdisk* to print its state on *stderr* as it is used. The output from this option should not be used with -F.

-D *id:act:bhead:bsect:bcyl:ehhead:esect:ecyl:rsect:numsect*

Delete a partition as described by the argument (see the -F option below for the format). Note that the argument must be an exact match or the entry will not be deleted! Use of this option will zero out the VTOC on the Solaris partition if the *fdisk* table changes.

-E

Create an EFI partition that uses the entire disk.

-F *fdisk_file*

Use *fdisk* file *fdisk_file* to initialize table. Use of this option will zero out the VTOC on the Solaris partition if the *fdisk* table changes.

The *fdisk_file* contains four specification lines for the primary partitions followed by specification lines for the logical drives. You must have four lines for the primary partitions if there is at least one logical drive. In this case, if the number of primary partitions to be configured is less than four, the remaining lines should be filled with zeros.

Each line is composed of entries that are position-dependent, are separated by whitespace or colons, and have the following format:

id act bhead bsect bcyl ehhead esect ecyl rsect numsect

...where the entries have the following values:

<i>id</i>	This is the type of partition and the correct numeric values may be found in <i>fdisk.h</i> .
<i>act</i>	This is the active partition flag; 0 means not active and 128 means active. For logical drives, this flag will always be set to 0 even if specified as 128 by the user.
<i>bhead</i>	This is the head where the partition starts. If this is set to 0, <i>fdisk</i> will correctly fill this in from other information.
<i>bsect</i>	This is the sector where the partition starts. If this is set to 0, <i>fdisk</i> will correctly fill this in from other information.
<i>bcyl</i>	This is the cylinder where the partition starts. If this is set to 0, <i>fdisk</i> will correctly fill this in from other information.
<i>ehhead</i>	This is the head where the partition ends. If this is set to 0, <i>fdisk</i> will correctly fill this in from other information.

- esect* This is the sector where the partition ends. If this is set to 0, `fdisk` will correctly fill this in from other information.
- ecyl* This is the cylinder where the partition ends. If this is set to 0, `fdisk` will correctly fill this in from other information.
- rsct* The relative sector from the beginning of the disk where the partition starts. This must be specified and can be used by `fdisk` to fill in other fields. For logical drives, you must make sure that there are at least 63 free sectors before the *rsct* specified for a logical drive.
- numsect* The size in sectors of this disk partition. This must be specified and can be used by `fdisk` to fill in other fields.
- g** Get the label geometry for disk and display on stdout (see the **-S** option for the format).
- G** Get the physical geometry for disk and display on stdout (see the **-S** option for the format).
- h** Issue verbose message; message will list all options and supply an explanation for each.
- I** Forgo device checks. This is used to generate a file image of what would go on a disk without using the device. Note that you must use **-S** with this option (see above).
- n** Don't update `fdisk` table unless explicitly specified by another option. If no other options are used, **-n** will only write the master boot record to the disk. In addition, note that `fdisk` will not come up in interactive mode if the **-n** option is specified.
- o *offset*** Block offset from start of disk. This option is used for **-P**, **-r**, and **-w**. Zero is assumed when this option is not used.
- P *fill_patt*** Fill disk with pattern *fill_patt*. *fill_patt* can be decimal or hex and is used as number for constant long word pattern. If *fill_patt* is #, then pattern is block # for each block. Pattern is put in each block as long words and fills each block (see **-o** and **-s**).
- r** Read from disk and write to stdout. See **-o** and **-s**, which specify the starting point and size of the operation.
- R** Treat disk as read-only. This is for testing purposes.
- s *size*** Number of blocks to perform operation on (see **-o**).

-S *geom_file*

Set the label geometry to the content of the *geom_file*. The *geom_file* contains one specification line. Each line is delimited by a new-line character (`\n`). If the first character of a line is an asterisk (*), the line is treated as a comment. Each line is composed of entries that are position-dependent, are separated by white space, and have the following format:

pcyl ncyl acyl bcyl nheads nsectors sectsiz

where the entries have the following values:

pcyl This is the number of physical cylinders for the drive.

ncyl This is the number of usable cylinders for the drive.

<i>acyl</i>	This is the number of alt cylinders for the drive.

b cyl This is the number of offset cylinders for the drive (should be zero).

nheads The number of heads for this drive.

nsectors The number of sectors per track.

sectsiz The size in bytes of a sector.

-t

Adjust incorrect slice table entries so that they will not cross partition table boundaries.

-T

Remove incorrect slice table entries that span partition table boundaries.

-V

Output the HBA (virtual) geometry dimensions. This option must be used in conjunction with the -w flag. This option will work for platforms which support virtual geometry. (x86 only)

-W

Write to disk and read from stdin. See -o and -s, which specify the starting point and size of the operation.

-W-

Output the disk table to stdout.

-W fdisk_file

Create an fdisk file *fdisk_file* from disk table. This can be used with the -F option above.

文件

/dev/rdisk/c0t0d0p0	Raw device associated with the fixed disk.

<code>/usr/lib/fs/ufs/mboot</code>	Default master boot program.

属性

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86 and SPARC
Availability	system/core-os

另请参见

[uname\(1\)](#), [fmthard\(1M\)](#), [format\(1M\)](#), [newfs\(1M\)](#), [parted\(1M\)](#), [prtvto\(1M\)](#), [attributes\(5\)](#)

诊断

Most messages will be self-explanatory. The following may appear immediately after starting the program:

Fdisk: cannot open <device>
This indicates that the device name argument is not valid.

Fdisk: unable to get device parameters for device <device>
This indicates a problem with the configuration of the fixed disk, or an error in the fixed disk driver.

Fdisk: error reading partition table
This indicates that some error occurred when trying initially to read the fixed disk. This could be a problem with the fixed disk controller or driver, or with the configuration of the fixed disk.

Fdisk: error writing boot record
This indicates that some error occurred when trying to write the new partition table out to the fixed disk. This could be a problem with the fixed disk controller, the disk itself, the driver, or the configuration of the fixed disk.

引用名	ff – list file names and statistics for a file system	
用法概要	ff [-F <i>FSType</i>] [-V] [<i>generic_options</i>] [-o <i>specific_options</i>] <i>special</i> ...	
描述	ff prints the pathnames and inode numbers of files in the file system which resides on the special device <i>special</i>. Other information about the files may be printed using options described below. Selection criteria may be used to instruct ff to only print information for certain files. If no selection criteria are specified, information for all files considered will be printed (the default); the -i option may be used to limit files to those whose inodes are specified. Output is sorted in ascending inode number order. The default line produced by ff is: <i>path-name</i> i-number The maximum information the command will provide is: <i>path-name</i> i-number size uid	
选项	-F	Specify the <i>FSType</i> on which to operate. The <i>FSType</i> should either be specified here or be determinable from <i>/etc/vfstab</i> by matching the <i>special</i> with an entry in the table, or by consulting <i>/etc/default/fs</i> .
	-V	Echo the complete command line, but do not execute the command. The command line is generated by using the options and arguments provided by the user and adding to them information derived from <i>/etc/vfstab</i> . This option may be used to verify and validate the command line.
	<i>generic_options</i>	Options that are supported by most <i>FSType</i> -specific modules of the command. The following options are available:
	-I	Do not print the i-node number after each path name.
	-l	Generate a supplementary list of all path names for multiply-linked files.
	-p <i>prefix</i>	The specified <i>prefix</i> will be added to each generated path name. The default is '.' (dot).
	-s	Print the file size, in bytes, after each path name.
	-u	Print the owner's login name after each path name.
	-a -n	Select if the file has been accessed in <i>n</i> days.
	-m -n	Select if the file has been written or created in <i>n</i> days.
	-c -n	Select if file's status has been changed in <i>n</i> days.
	-n <i>file</i>	Select if the file has been modified more recently than the argument <i>file</i> .

-i i-node-list

Generate names for only those i-nodes specified in *i-node-list*. *i-node-list* is a list of numbers separated by commas (with no intervening spaces).

-o

Specify *FSType*-specific options in a comma separated (without spaces) list of suboptions and keyword-attribute pairs for interpretation by the *FSType*-specific module of the command.

操作数

special

A special device.

用法

See [largefile\(5\)](#) for the description of the behavior of `ff` when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).

文件

/etc/default/fs

default local file system type. Default values can be set for the following flags in `/etc/default/fs`. For example: `LOCAL=ufs`

LOCAL

The default partition for a command if no *FSType* is specified.

/etc/vfstab

list of default parameters for each file system

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[find\(1\)](#), [ncheck\(1M\)](#), [stat\(2\)](#), [vfstab\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#) Manual pages for the *FSType*-specific modules of `ff`.

附注

This command may not be supported for all *FSTypes*.

The `-a`, `-m`, and `-c` flags examine the `st_atime`, `st_mtime`, and `st_ctime` fields of the `stat` structure respectively. (See [stat\(2\)](#).)

系统管理命令－第 1 单元

585

- 引用名

ff_ufs – list file names and statistics for a ufs file system
- 用法概要

ff -F ufs [*generic_options*] [-o a,m,s] *special*...
- 描述

ff prints the pathnames and inode numbers of files in the file system which resides on the special device *special*.

See [ff\(1M\)](#) for information regarding the ff command. See **OPTIONS** for information regarding the ufs-specific options.
- 选项

The following options are supported:

 - o Specify ufs file system specific options. The following options available are:
 - a Print the '.' and '..' directory entries.
 - m Print mode information. This option must be specified in conjunction with the -i *i-node-list* option (see [ff\(1M\)](#)).
 - s Print only special files and files with set-user-ID mode.
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[find\(1\)](#), [ff\(1M\)](#), [ncheck\(1M\)](#), [attributes\(5\)](#)

引用名	fiocompress – file compression utility
用法概要	<pre>/usr/sbin/fiocompress -c [-m] [-b block_size] input_file output_file /usr/sbin/fiocompress -d input_file output_file</pre>
描述	The <code>fiocompress</code> utility is a file compression tool that works together with the dcfs(7FS) file system to perform per-file compression. You can use <code>fiocompress</code> to decompress a compressed file or mark a compressed file as compressed, causing automatic decompression on read. The primary use of <code>fiocompress</code> is to compress files in the boot archive. Note that this utility is not a Committed interface. See attributes(5).
选项	The following options are supported: <code>-b block_size</code> Specify a block size for compression. The default block size is 8192. <code>-c</code> Compress the specified file. <code>-d</code> Decompress the specified file. <code>-m</code> Mark the compressed file for automatic decompression on read. Can be used only in conjunction with <code>-c</code>.
退出状态	0 The command completed successfully. –1 The command exited due to an error.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Private

另请参见 [boot\(1M\)](#), [bootadm\(1M\)](#), [dcfs\(7FS\)](#), [ufs\(7FS\)](#), [attributes\(5\)](#)

附注 This compression/decompression utility works only with files stored in a UFS file system.

There is no obvious way to determine whether a given file is compressed, other than copying the file and comparing the number of disk blocks of the copy against the original.

引用名

flowadm – 为协议、服务、容器以及虚拟机管理带宽资源控制

用法概要

```
flowadm show-flow [-P] [[-p] -o field[,...]] [-l link] [flow]

flowadm add-flow [-t] [-R root-dir] -l link -a attr=value[,...]
    [-p prop=value[,...]] flow

flowadm remove-flow [-t] [-R root-dir] {-l link | flow}

flowadm set-flowprop [-t] [-R root-dir] -p prop=value[,...] flow
flowadm reset-flowprop [-t] [-R root-dir] [-p prop[,...]] flow
flowadm show-flowprop [-P] [[-c] -o field[,...]] [-l link]
    [-p prop[,...]] [flow]
```

描述

flowadm 命令用于为特定链接上的某类通信流量创建、修改、删除和显示网络带宽以及关联资源。

flowadm 命令允许用户针对传输、服务或子网管理网络带宽资源。将此服务指定为传输和本地端口的组合。子网由其 IP 地址和子网掩码指定。该命令可在任何类型的数
据链接上使用，包括物理链接、虚拟 NIC 以及链路聚合。

流定义为基于层 3 和层 4 头的属性集合，可用于标识协议、服务或虚拟机。当基于流属性标识流时，请分隔内核资源（包括层 2、3 和 4 队列）、其处理线程以及其他针对其唯一创建的资源，以便其他通信流量不会或尽可能小地影响到它。

传入和外发包以非常快速的可伸缩方式与流匹配，因此，可以在性能影响最小的情况下执行限制。

flowadm 命令可用于标识流，而无需强加任何带宽资源控制。这将导致通信流量类型获取其自己的资源和队列，以便其和其余网络通信流量隔离，获得更易监测和更具确定性的行为。

流可以在全局区域和非全局区域中创建、修改和删除。区域管理员仅可以在其区域（全局或非全局）中创建流。然而，在全局区域中创建的流可以迁移至非全局区域，以下段落中有述。管理员可以仅从其中创建流的区域（全局或非全局）内修改或删除流。从全局区域，用户可以查看系统上全局区域和任何非全局区域内的所有流。从非全局区域，用户仅可以查看此区域内的那些流。

管理员在全局区域中创建了流之后，与该流关联的数据链接可以分配至非全局区域。在此情况下，关联流也分配至同一非全局区域。当该非全局区域停止后，数据链接及其关联流将返回至全局区域。

不同的区域名区分同一名称的流。例如，用户可以拥有三个名为 fastpak 的流，如果每个 fastpak 位于不同的区域。例如，zone1/fastpak、zone2/fastpak 和 zone3/fastpak 都是有效的区域名。

flowadm 将作为具有相应选项的子命令集合进行实现。选项在每个子命令的上下文中进行了说明。

子命令

支持以下子命令：

flowadm show-flow [-pP] [-o *field*[,...]] [-l *link*] [*flow*]

显示所有流、某个链接上的所有流或指定的 *flow* 的流配置信息（缺省值）或统计信息。

-o *field*[,...]

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名必须是以下所列字段之一，或为特殊值 **all**，以显示所有字段。对于找到的每个流，可显示以下字段：

flow

流的名称。

link

流所在的链接名。

ipaddr

流的 IP 地址。此地址可以是本地的，也可以是远程的，具体取决于流的定义。

transport

针对要使用协议的层名。

port

流的本地服务端口。

dsfield

流和掩码的区分服务值，与 **DSFIELD** 值结合使用，以指定 IP 数据包头的区分服务字段中受关注的位。

-p, --parseable

使用稳定的可供计算机解析的格式显示。

-P, --persistent

显示持久流属性信息。

-l *link*, --link=*link* | *flow*

显示命名链接上所有流的信息或命名流的信息。

flowadm add-flow [-t] [-R *root-dir*] -l *link* -a *attr=value*[,...] -p *prop=value*[,...] *flow*

将流添加到系统。流由其流属性标识。

作为标识特定流的一部分，可以限制它的带宽资源。如果未指定带宽限制，通信流量仍将获取其唯一的层 2、3 和 4 队列和处理线程，包括 NIC 硬件资源（受支持时），以便选定的通信流量可与其他流量分隔开来，并在流动过程中受到来自其他通信流量的影响最小。

-t, --temporary

更改是临时的，在重新引导期间不会持续存在。持续存在是缺省值。

-R *root-dir*、--root-dir=*root-dir*
 指定备用根目录，其中 **flowadm** 应应用持久创建。

-l *link*、--link=*link*
 指定将向其添加流的链接。

-a *attr=value*[,...], --attr=*value*
 要设置为指定值的以逗号分隔的属性列表。

-p *prop=value*[,...], --prop=*value*[,...]
 要设置为指定值的以逗号分隔的属性列表。

flowadm remove-flow [-t] [-R *root-dir*] -l [*link* | *flow*]
 删除由其链接或名称标识的现有流。

-t、--temporary
 更改是临时的，在重新引导期间不会持续存在。持续存在是缺省值。

-R *root-dir*、--root-dir=*root-dir*
 指定备用根目录，其中 **flowadm** 应应用持久删除。

-l *link* | *flow*, --link=*link* | *flow*
 如果指定一个链接，则从该链接删除所有流。如果指定单个流，则仅删除该流。

flowadm set-flowprop [-t] [-R *root-dir*] -p *prop=value*[,...] *flow*
 设置按名称指定的流上的一个或多个属性的值。使用 **show-flow** 子命令可以检索完整的属性列表。

-t、--temporary
 更改是临时的，在重新引导期间不会持续存在。持续存在是缺省值。

-R *root-dir*、--root-dir=*root-dir*
 指定备用根目录，其中 **flowadm** 应应用持久属性设置。

-p *prop=value*[,...], --prop=*value*[,...]
 要设置为指定值的以逗号分隔的属性列表。

flowadm reset-flowprop [-t] [-R *root-dir*] -p [*prop=value*[,...]] *flow*
 将指定流上的一个或多个属性重置为其缺省值。如果未指定属性，则重置所有属性。请参见 **show-flowprop** 子命令了解属性说明，其中包含它们的缺省值。

-t、--temporary
 指定重置为临时重置。临时重置将持续到下一次重新引导。

-R *root-dir*、--root-dir=*root-dir*
 指定备用根目录，其中 **flowadm** 应应用持久属性设置。

-p *prop=value*[,...], --prop=*value*[,...]
 要重置的以逗号分隔的属性列表。

flowadm show-flowprop [-cP] [-l *link*] [-p *prop*[,...]] [*flow*]
 显示所有流、指定链接上的流或指定流的一个或多个属性的当前值或持久值。

缺省情况下，将显示当前值。如果未指定属性，所有可用流属性都将显示。对于每个属性，将显示下列字段：

FLOW (流)

流的名称。

PROPERTY (属性)

属性的名称。

VALUE (值)

当前（或持久性）属性值。如果该值未设置，则显示为 --（双连字符），如果该值未知，则显示为？（问号）。尚未设置或尚未重置的持久性值将显示为 --，并且将使用系统的 DEFAULT 值（如果有）。

DEFAULT (缺省)

属性的缺省值。如果属性无缺省值，则显示 --（双连字符）。

POSSIBLE (可能)

属性可具有的值的逗号分隔列表。如果这些值跨越一定的数字范围，最小值和最大值可能会显示为缩写。如果可能值未知或无限制，则显示 --（双连字符）。

流属性在以下“流属性”部分中有述。

-c, --parseable

使用稳定的可供计算机解析的格式显示。

-P, --persistent

显示持久流属性信息。

-p *prop*[,...], --prop=*prop*[,...]

要显示的属性的逗号分隔列表。

流属性

在 flowadm 命令中标识流的流操作数是一个来自以下列表的以逗号分隔的一个或多个关键字、值对列表。

local_ip[/*prefix_len*]

按本地 IP 地址标识网络流。*value* 必须是采用点分十进制记法的 IPv4 地址或采用冒号分隔记法的 IPv6 地址。*prefix_len* 是可选项。

如果指定了 *prefix_len*，它将描述子网地址的网络掩码，并与 ifconfig(1M) 和 route(1M) 地址遵循相同的表示法约定。如果未指定，给定的 IP 地址将被视为主机地址，其 IPv4 地址的缺省前缀长度是 /32，对于 IPv6，为 /128。

remote_ip[/*prefix_len*]

按远程 IP 地址标识网络流。语法与 local_ip 属性相同

transport={tcp|udp|sctp|icmp|icmpv6}

标识要使用的层 4 协议。通常结合 local_port 或 remote_port 使用，以标识需要特别关注的本地或远程服务。

`local_port`

标识由本地端口指定的服务。

`remote_port`

标识由远程端口指定的服务。

`dsfield[:dsfield_mask]`

标识 8 位区分服务字段（如 RFC 2474 中所定义）。

可选的 `dsfield_mask` 用于在与 `dsfield` 值比较时，指定区分服务字段中受关注的位。位位置中的 0 表示该位值需要忽略，1 则表示不要忽略。掩码范围可以从 0x01 到 0xff。如果未指定 `dsfield_mask`，将使用缺省掩码 0xff。`dsfield` 值和掩码必须为十六进制。

以下类型的属性组合受支持：

```
local_ip[/prefixlen]=address
remote_ip[/prefixlen]=address
transport={tcp|udp|sctp|icmp|icmpv6}
transport={tcp|udp|sctp},local_port=port
transport={tcp|udp|sctp},remote_port=port
dsfield=val[:dsfield_mask]
```

在给定链接上，上述组合是互斥的。尝试创建不同组合的流将失败。

限制

此处存在个别流限制与各区域流限制。

个别流限制

个别流上的限制无需了解其他已添加至链接的流。

针对每个流，属性只能列一次。例如，以下命令无效：

```
# flowadm add-flow -l vnic1 -a local_port=80,local_port=8080 httpflow
```

`transport` 和 `local_port` 或者 `transport` 和 `remote_port`：

TCP、UDP 或 SCTP 流可通过本地端口或远程端口进行指定。不允许指定端口的 ICMP 或 ICMPv6 流。

如果指定了 `local_port` 或 `remote_port`，传输必须是 TCP、UDP 或 SCTP。

以下命令有效：

```
# flowadm add-flow -l e1000g0 -a transport=udp udpflow
# flowadm add-flow -l e1000g0 -a transport=tcp,local_port=80 \
udp80flow
```

以下命令无效：

```
# flowadm add-flow -l e1000g0 -a remote_port=25 flow25
# flowadm add-flow -l e1000g0 -a transport=icmpv6,remote_port=16 \
flow16
```

各区域流限制

在一个区域内，两个流不能具有相同的名称。添加具有指定链接的流后，显示、修改或删除该流时无需此链接。

流属性

以下流属性均受支持。请注意，将给定属性设置为给定值的功能取决于驱动程序和硬件。

maxbw
设置流的全双工带宽。此带宽指定为一个带有某个容量级后缀（分别表示 Kbps、Mbps 和 Gbps 的 K、M 或 G）的整数。如果未指定单位，输入值将以 Mbps 为单位进行读取。缺省值是无带宽限制。

示例

示例1 围绕重点端口创建策略

以下命令将围绕 HTTPS 服务器上的传入 HTTPS 通信流量创建策略，以便 HTTPS 获取专用 NIC 硬件和内核 TCP/IP 资源。稍后可使用指定的名称 `https-1` 修改或删除该策略。

```
# flowadm add-flow -l bge0 -a transport=TCP,local_port=443 https-1
# flowadm show-flow -l bge0
FLOW          LINK          IP ADDR          PROTO  PORT  RPORT  DSFLD
https1        bge0          --              tcp    443   --     --
```

示例2 修改现有策略以添加带宽资源控制

以下命令将修改之前示例中的 `https-1` 策略。该命令可以添加带宽控制。

```
# flowadm set-flowprop -p maxbw=500M https-1
# flowadm show-flow https-1
FLOW          LINK          IP ADDR          PROTO  PORT  RPORT  DSFLD
https1        bge0          --              tcp    443   --     --

# flowadm show-flowprop https-1
FLOW          PROPERTY  VALUE  DEFAULT  POSSIBLE
https-1       maxbw    500    --       --
```

示例3 限制 UDP 带宽使用

以下命令将针对 UDP 协议创建策略，以便其无法使用超过 100Mbps 的可用带宽。将流命名为 `limit-udp-1`。

```
# flowadm add-flow -l bge0 -a transport=UDP -p maxbw=100M, \
limit-udp-1
```

示例4 设置策略，利用 dsfield 属性

以下命令将针对具有 500 Mbps 的带宽的 EF PHB（101110 的 DSCP 值来自 RFC 2598）设置策略。该流的 `dsfield` 值将是 `0x2e`（101110），而 `dsfield_mask` 为 `0xfc`（因为我们希望忽略两个最低有效位）。

示例 4 设置策略，利用 dsfield 属性 (续)

```
# flowadm add-flow -l bge0 -a dsfield=0x2e:0xfc \  
-p maxbw=500M efphb-flow
```

示例 5 查看多个区域中的流

以下命令将显示具有相同名称的两个流。第一个流位于全局区域中，第二个位于区域 zone1 中。该命令从全局区域调用，使用户能够查看系统上的所有流

```
# flowadm show-flow  
FLOW      LINK      IPADDR      PROTO  LPORT  RPORT  DSFLD  
tcpflow    e1000g2    --          tcp    --     --     --  
zone1/tcpflow e1000g1    --          tcp    --     --     --
```

退出状态

0
所有操作都成功执行。

>0
出现错误。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/network
接口稳定性	Committed（已确定）

另请参见

[acctadm\(1M\)](#)、[dladm\(1M\)](#)、[flowstat\(1M\)](#)、[ifconfig\(1M\)](#)、[prstat\(1M\)](#)、[route\(1M\)](#)、[attributes\(5\)](#)

附注

show-usage 子命令（在以前版本的 flowadm 中存在）已替换为 [flowstat\(1M\)](#) -h 命令。

引用名	flowstat – 报告流统计信息
用法概要	<pre>flowstat [-r -t] [-i interval] [-l link] [flow] flowstat [-S] [-A] [-i interval] [-p] [-o field[,...]] [-u R K M G T P] [link] [-l link] [flow] flowstat -h [-a] -f filename [-d] [-F format] [-s time] [-e time] [flow]</pre>
描述	flowstat 命令报告有关用户定义的流的运行时统计信息。flowadm show-flow 为该命令提供流名称信息。
选项	flowstat 命令具有以下选项和操作数，这些选项和操作数是下面“子命令”下显示的多种命令形式中的常见形式。
子命令	flowstat 支持以下命令形式。 <pre>flowstat [-r -t] [-i interval] [-l link] [flow]</pre> 该命令形式反复检查所有流并报告统计信息。输出按流使用情况降序排列。如果未指定任何流，系统会显示所有流的统计信息。 -r 仅显示接收端统计信息。包括接收的字节和包、删除等。请参见完整列出的示例。 -t 仅显示传输端统计信息。包括发送的字节和包、删除等。请参见下面的示例。 -i interval 指定刷新统计信息的间隔（以秒为单位）。缺省间隔为 1 秒。 -l link flow 显示指定链接上所有流的统计信息，或者指定流的统计信息。 <pre>flowstat [-S] [-A] [-i interval] [-p] [-o field[,...]] [-u R K M G T P] [-l link] [flow]</pre> 该命令形式允许您指定要显示的统计信息。 -A 转储该流的所有统计信息字段。该命令的输出统计信息包含所有其他 flowstat 命令所报告的所有统计信息。 -i interval 指定刷新统计信息的间隔（以秒为单位）。缺省间隔为 1 秒。 {-l link} flow 显示指定链接上所有流的统计信息，或者指定流的统计信息。 -o field[,...] 显示要显示的输出字段的逗号分隔列表，不区分大小写。字段名称必须是下面列出的字段之一，或特殊值 all 以显示所有支持的字段。 支持的 RX 字段的列表：

- flow
- rbytes
- ipkts
- idrops

TX 字段的列表：

- flow
- obytes
- opkts
- odrops

-p

以稳定的机器可解析格式显示输出。

-S

按流连续显示网络的使用情况的方式与 [prstat\(1M\)](#) 按进程显示 CPU 使用情况的方式类似。

-u R|K|M|G|T|P

如果使用，允许选择显示所有统计信息时采用的单位，例如 **R**：原始数、**K**：千位、**M**：兆位、**T**：兆兆位、**P**：千兆位。如果未使用，可使用不同的单位（根据需要）显示统计信息。

flowstat -h [-a] -f filename [-d] [-F format] [-s time] [-e time] [flow]

显示存储的扩展记帐文件中的网络使用历史。使用该语法要求网络帐户已事先配置并通过 [acctadm\(1M\)](#) 启用。缺省输出是启用扩展记帐的整个周期中现有链路的网络使用情况的摘要。

-a

显示启用扩展记帐的指定周期内的所有历史网络使用情况。这包括已删除的流的使用信息。

-f filename

指定从中读取网络流量使用历史记录 of 扩展记帐记录的文件。

-d

显示具有日志信息的日期。日期格式为 *mm/dd/yyyy*。

-F format

指定网络流量使用历史记录信息的输出格式。gnuplot 是唯一支持的格式。

-s time

-e time

指定数据显示的开始时间和停止时间。时间格式为 *MM/DD/YYYY*，*hh:mm:ss*。hh 采用 24 小时制表示法。

操作数

flowstat 命令形式具有单个可选的操作数。

flow

如果指定该选项，仅对命名的流进行报告。否则，将对所有流进行报告。流名称的格式为 *zonename/flowname*。不包含 *zonename* 修饰符的 *flowname* 被视为处于全局区域。

示例

示例1 显示统计信息

要显示所有流的统计信息，请输入以下命令。使用相应单位的 3 位数显示统计信息。缺省时间间隔为 1 秒。

```
# flowstat -i 1
FLOW  IPKTS  RBYTES  IDROPS  OPKTS  OBYTES  ODROPS
flow1 528.54K 787.39M      0 179.39K 11.85M      0
flow2 742.81K  1.10G      0      0      0      0
flow3      0      0      0      0      0      0
flow1 67.73K 101.02M      0 21.04K  1.39M      0
flow2      0      0      0      0      0      0
flow3      0      0      0      0      0      0
      .      .      .      .      .      .
      .      .      .      .      .      .
      .      .      .      .      .      .
```

示例2 显示 RX 端统计信息

以下命令显示接收端统计信息。

```
# flowstat -r
FLOW  IPKTS  RBYTES  IDROPS
flow1 4.01M  5.98G      0
flow2 742.81K 1.10G      0
flow3      0      0      0
```

示例3 显示发送端统计信息

以下命令显示发送端统计信息，且时间间隔为 5 秒。

```
# flowstat -t
FLOW  OPKTS  OBYTES  ODROPS
flow1 24.37M  1.61G      0
flow2      0      0      0
flow3      4    216      0
```

示例4 显示一组特定的统计信息

以下命令显示一组指定的统计信息字段。

```
# flowstat -o FLOW,IPKTS
FLOW  IPKTS
flow1 68.58M
flow2 742.81K
```

示例 4 显示一组特定的统计信息 (续)

flow3 4

示例 5 显示网络历史使用情况

可以通过扩展记帐工具 **acctadm(1M)** 存储流使用统计信息。

```
# acctadm -e extended -f /var/log/net.log net
# acctadm net
      Network accounting: active
      Network accounting file: /var/log/net.log
      Tracked Network resources: extended
      Untracked Network resources: none
```

可以检索保存的历史数据，如下所示：

```
# flowstat -h -f /var/log/net.log
LINK      DURATION  IPACKETS  RBYTES      OPACKETS  OBYTES      BANDWIDTH
flowtcp   100       1031      546908      0          0           43.76Kbps
flowudp   0         0         0           0          0           0.00Mbps
```

显示 **flowtcp** 的日志信息，从 2008 年 2 月 19 日 10:38:46 开始并在同一天的 10:40:06 结束：

```
# flowstat -h -s 02/19/2008,10:39:06 -e 02/19/2008,10:40:06 \
-f /var/log/net.log flowtcp
FLOW      START      END          RBYTES  OBYTES      BANDWIDTH
flowtcp   10:39:06   10:39:26    1546    6539        3.23 Kbps
flowtcp   10:39:26   10:39:46    3586    9922        5.40 Kbps
flowtcp   10:39:46   10:40:06    240     216         182.40 bps
flowtcp   10:40:06   10:40:26    0        0           0.00 bps
```

将与上述信息相同的输出信息生成为 **plotfile**（绘图文件）：

```
# flowstat -h -s 02/19/2008,10:39:06 -e 02/19/2008,10:40:06 \
-F gnuplot -f /var/log/net.log flowtcp
# Time tcp-flow
10:39:06 3.23
10:39:26 5.40
10:39:46 0.18
10:40:06 0.00
```

退出状态

```
0
    所有操作都成功执行。
>0
    出现错误。
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

`/usr/sbin`

属性类型	属性值
可用性	system/core-os
接口稳定性	请参见下文

屏幕输出为 Uncommitted（未确定）。调用为 Committed（已确定）。

另请参见 [acctadm\(1M\)](#)、[dladm\(1M\)](#)、[dlstat\(1M\)](#)、[flowadm\(1M\)](#)、[ifconfig\(1M\)](#)、[prstat\(1M\)](#)、[route\(1M\)](#)

引用名 fmadm – 故障管理配置工具

用法概要 fmadm [-q] [*subcommand* [*arguments*]]

描述 管理员和服务人员可以使用 **fmadm** 实用程序来查看和修改由 Solaris Fault Manager **fmd(1M)** 维护的系统配置参数。fmd 接收与系统软件检测到的问题有关的遥测信息，诊断这些问题，启动主动式自我修复活动，如禁用故障组件等。

fmadm 可用于：

- 查看当前参与故障管理的诊断引擎和代理的集合，
- 查看被诊断为故障的系统组件列表，以及
- 执行与这些实体有关的管理任务。

Fault Manager 尝试将尽可能多的活动自动化，这样，通常就不需要使用 **fmadm** 了。当 Fault Manager 需要管理员、服务维修技术人员或 Oracle 的帮助时，会发出一条说明其需要的消息。它还会为您提供 Oracle 的 Web 站点上的知识文章。该 Web 站点可能会要求您使用 **fmadm** 或某个其他的故障管理实用程序来收集更多信息，或执行其他任务。**fmd(1M)**、**fmdump(1M)** 和 **fmstat(1M)** 的文档对于观察故障管理活动的工具进行了详细的描述。

Fault Manager 的一个职责是跟踪组件的位置。在机箱级别，**fmadm *-alias** 子命令管理机箱 *product-id.chassis-id* 到 *alias-id* 的映射。受管理的 *alias-id* 将以某种有意义的方式描述机箱的物理位置。

fmadm 实用程序要求用户拥有 **SYS_CONFIG** 特权。请参见《[Oracle Solaris 管理：安全服务](#)》了解有关如何配置 Solaris 特权的更多信息。**fmadm load** 子命令要求用户拥有所有特权。

子命令 **fmadm** 接受以下子命令：有些子命令接受或要求额外的选项和操作数。**load**、**unload**、**reset** 和 **rotate** 子命令是供经过培训的技术人员使用的。我们建议，在没有获得具体指导（如阅读知识库文章）之前，不要使用这些子命令。

fmadm acquit fmri | label [uuid]

通知 Fault Manager，不会将指定资源认定为 *uuid* 标识的故障事件中（如果未指定 UUID，则是在检测到的所有故障中）的可疑对象。**fmadm acquit** 子命令应仅在 Oracle 修复过程文档的指导下使用。管理员可能需要使用额外的命令来重新启用之前发生故障的资源。

fmadm acquit uuid

通知 Fault Manager，可以安全地忽略由 *uuid* 标识的故障事件。**fmadm acquit** 子命令应仅在 Oracle 修复过程文档的指导下使用。管理员可能需要使用额外的命令来重新启用之前发生故障的任何资源。

fmadm config

显示 Fault Manager 本身的配置，包括模块名称、版本以及每个组件模块的描述。Fault Manager 模块可提供包括自动诊断、自我修复以及向系统中存在的硬件和软件发送消息在内的各种服务。

fmadm faulty [-afgiprsv] [-n *max*] [-u *uid*]

显示 Fault Manager 当前认定为故障的资源的状态信息。

所支持的选项如下：

- a 显示所有故障。缺省情况下，**fmadm faulty** 命令只列出当前存在且发生故障的资源的输出。如果指定了 -a 选项，则会列出 Fault Manager 缓存的所有资源信息，包括已经自动更正或无需恢复操作的故障。列表中包括系统中可能已不存在的资源的信息。
- f 显示故障的 fru（Field replaceable units，现场可更换单元）。
- g 将具有相同 fru、类和故障消息的故障归为一组。
- i 显示 Fault Manager 中每个资源的持久性高速缓存标识符。
- n *max* 如果以 -a 或 -g 选项对故障或资源进行分组，会将输出限制为 *max* 项。
- p 以故障间的换页使输出通过页面调度程序。
- r 显示带有标识符的故障管理资源 (Fault Management Resource with their Identifier, FMRI) 及其故障管理状态。
- s 对每个故障事件显示一行的故障摘要。
- u *uid* 仅显示具有给定 *uid* 的故障。
- v 显示完整输出。

故障有多个可疑对象（具有不同的类或出现在不同的 fru 上）时会显示确定性百分比。如果同一 fru 上有多个资源，且无法 100% 确定故障与该 fru 有关，系统会显示该 fru 上可能的可疑对象的最大确定性百分比。

Fault Manager 将以下状态与收到其遥测信息的每个资源相关联：

ok（正常）

就 Fault Manager 而言，相应资源始终存在且正在使用中，并且没有已知的问题。

unknown（未知）

资源不存在或不可用，但是没有已知的问题。这可能表示，资源已被管理员禁用或取消配置。请通过相应的管理工具来了解更多信息。

faulted（错误）

资源存在但是不可用，因为已通过 Fault Manager 诊断出一个或多个问题。已禁用资源，以防止对系统造成进一步的损害。

degraded（降级）

资源存在且可以使用，但是 Fault Manager 诊断出资源中有一个或多个问题。

如果所有受影响的资源都处于相同状态，列表末端的消息中会说明这一点。否则，会在每个受影响的资源之后给出状态。

fmadm flush *fmri*

刷新由 Fault Manager 缓存的指定资源的信息（以其 FMRI 命名）。该子命令应仅在 Oracle 修复过程文档的指导下使用。通常，不需要使用该命令，因为 Fault Manager 会自动使其高速缓存保持最新。如果在高速缓存中刷新故障资源，管理员可能需要使用额外的命令来启用指定资源。

fmadm load *path*

装入指定的 Fault Manager 模块。*path* 必须是绝对路径，且必须指向某个为模块定义的目录中存在的模块。通常，不需要使用此命令，因为 Fault Manager 会在 Solaris 最初引导时或根据需要自动装入模块。

fmadm unload *module*

卸载指定的 Fault Manager 模块。使用 **fmadm config** 输出中所列的根基名称来指定 *module*。通常，不需要使用此命令，因为 Fault Manager 会根据系统配置自动装入和卸载模块。

fmadm repaired *fmri* | *label*

通知 Fault Manager，已经对指定的资源执行修复过程。**fmadm repaired** 子命令应仅在 Oracle 修复过程文档的指导下使用。管理员可能需要使用额外的命令来重新启用之前发生故障的资源。

fmadm replaced *fmri* | *label*

通知 Fault Manager，已经替换指定的资源。该命令应在 Fault Manager 无法自动检测替换的情况下使用。**fmadm replaced** 子命令应仅在 Oracle 修复过程文档的指导下使用。管理员可能需要使用额外的命令来重新启用之前发生故障的资源。

fmadm reset [-s *serd*] *module*

复位指定的 Fault Manager 模块或模块子组件。如果存在 **-s** 选项，将在模块内复位指定的 Soft Error Rate Discrimination (SERD) 引擎。如果不存在 **-s** 选项，会复位整个模块，并删除所有与该模块关联的持久性状态。**fmadm reset** 子命令应仅在 Oracle 修复过程文档的指导下使用。通常，不需要使用此命令，因为 Fault Manager 自动管理其模块。

fmadm rotate errlog | fltlog | infolog | infolog_hival

rotate 子命令是 **logadm(1M)** 的辅助命令，可用于确保 **logadm** 正确轮转活动的日志文件。不可直接调用（直接调用可能会丢失日志历史）。如果当前日志文件的大小不为零，可以使用以下命令之一轮转相应的日志文件：

```
# logadm -p now -s 1b /var/fm/fmd/errlog
# logadm -p now -s 1b /var/fm/fmd/fltlog
# logadm -p now -s 1b /var/fm/fmd/infolog
# logadm -p now -s 1b /var/fm/fmd/infolog_hival
```

fmadm add-alias *product-id.chassis-id alias-id* ['*comment*']

add-alias 子命令用于建立 *alias-id*，作为 *product-id.chassis-id* 机箱的受管理别名。定义受管理别名后，机箱的 **/dev/chassisdevchassis(7FS)** 名称空间表示将使用更有意义的 *alias-id* 来代替 *product-id.chassis-id*。

```
# fmadm add-alias SUN-Storage-J4410.1039QAQ007 RACK29.U25-28
```

上述命令将验证新映射是否会与现有映射产生冲突。如果产生冲突，则映射不会发生任何更改。当关联的名称空间更新完成时，该子命令即完成。如果更新后的名称空间没有使用新的 *alias-id*，则会显示一条警告，但是会更新映射。如果名称空间更新花费的时间过长，将显示一条警告。

如果提供可选注释，该注释会被保留下来，后续 `lookup-alias` 或 `list-alias` 命令将会显示该注释。

```
fmadm remove-alias alias-id | product-id.chassis-id
remove-alias 子命令用于删除 product-id.chassis-id 到 alias-id 的映射。
```

```
# fmadm remove-alias RACK29.U25-28
```

当关联的名称空间更新完成时，上述子命令即完成。

```
fmadm lookup-alias alias-id | product-id.chassis-id
lookup-alias 子命令可用于确定当前映射。下面是一个命令示例。
```

```
# fmadm lookup-alias SUN-Storage-J4410.1039QAQ007
```

```
fmadm list-alias
list-alias 子命令用于显示所有注释和映射。
```

```
fmadm sync-alias
sync-alias 子命令用于成批手动导入一组映射。将维护当前映射的两个副本：
```

- /etc/dev/chassis_aliases
- /etc/dev/.chassis_aliase

要成批导入一组映射，您可以更新 /etc/dev/chassis_aliases 文件，然后运行 `sync-alias`。

选项

支持以下选项：

```
-q    设置静默模式。fmadm 不会向标准输出发出表明成功操作结果的消息。
```

操作数

支持下列操作数：

- cmd* “子命令”中列出的一个子命令的名称。
- args* 一个或多个适用于选定 *subcommand* 的选项或参数，如“子命令”中所述。这些参数包括 `fmri`、`uuid` 和 `label`。这些参数会标识作为 `fmadm` 子命令的对象的资源。使用 `fmadm faulty` 获取用于目标资源的 `fmri`、`uuid` 和 `label`。请参见“示例”部分。通常，`label` 是这些操作数中用户友好程度最高的。

示例

示例1 调用 `faulty` 子命令

以下命令调用 `faulty` 子命令，它显示一个组件的 `uuid`、`label` 和 `fmri`。

```
# fmadm faulty
-----
TIME          EVENT-ID                                     MSG-ID      SEVERITY
-----
```

示例1 调用 faulty 子命令 (续)

```
Sep 09 16:15 96609fae-113c-e48c-b1cf-ebf4b0902d72 DISK-8000-3E Critical
                                     injected
```

```
Host      : x4170-brm-02
Platform  : SUN-FIRE-X4170-SERVER    Chassis_id : 0920XF508B
Product_sn:

Fault class: fault.io.scsi.cmd.disk.dev.rqs.derr
Affects : dev:///devid=id1,sd@n5000c5000940edbb//scsi_vhci/disk@g5000c\
5000940edbb out of service, but associated components no longer faulty
FRU      : "DISK 11"
(hc:///product-id=SUN-Storage-J4410:server-id=:chassis-id=:serial=000930\
G01CN4----3SJ01CN4:part=SEAGATE-ST330057SSUN300G:revision=0205/\
ses-enclosure=0/bay=11/disk=0)
replaced
...
...
```

在上面的输出中，uuid 是 EVENT-ID 列中的第一项
96609fae-113c-e48c-b1cf-ebf4b0902d72。label 是 FRU 行中的 DISK 11。fmri 是：

```
dev:///devid=id1,sd@n5000c5000940edbb//scsi_vhci/disk@g5000c\
5000940edbb

hc:///product-id=SUN-Storage-J4410:server-id=:chassis-id=:serial=000930\
G01CN4-----3SJ01CN4:part=SEAGATE-ST330057SSUN300G:revision=0205/\
ses-enclosure=0/bay=11/disk=0)
```

使用 fmdump -v 可以获得相同的值：

```
# fmdump -v
Sep 09 16:15:36.9252 96609fae-113c-e48c-b1cf-ebf4b0902d72 DISK-8000-3E \
Diagnosed 100% fault.io.scsi.cmd.disk.dev.rqs.derr

Problem in: hc:///scheme=:product-id=SUN-Storage-J4410:chassis-id=: \
server-id=/ses-enclosure=0/bay=11/disk=0

Affects: dev:///devid=id1,sd@n5000c5000940edbb//\
scsi_vhci/disk@g5000c5000940edbb
FRU: hc:///product-id=SUN-Storage-J4410:server-id=:chassis-id=: \
serial=000930G01CN4-----3SJ01CN4:part=SEAGATE-ST330057SSUN300G: \
revision=0205/ses-enclosure=0/bay=11/disk=0

Location: DISK 11
```

请注意，label 是使用最方便的标识符。

示例2 获取模块名称
以下命令显示每个组件的模块名称。模块名称被指定为 `fmadm unload` 命令的输入。

```
# fmadm config
MODULE          VERSION STATUS DESCRIPTION
cpumem-retire   1.1    active CPU/Memory Retire Agent
disk-transport  1.0    active Disk Transport Agent
eft             1.16   active eft diagnosis engine
..
```

退出状态 将返回下列退出值：

- 0 成功完成。
- 1 出现错误。错误包括与 `fmd` 的通信失败，或是执行请求操作的权限不足。
- 2 指定的命令行选项无效。

属性 有关以下属性的说明，请参见 `attributes(5)`：

属性类型	属性值
可用性	system/fault-management
接口稳定性	请参见下文。

命令行选项为 `Committed`（已确定）。用户可读输出为 `not-an-interface`（不是接口）。

另请参见 `fmd(1M)`、`fmdump(1M)`、`fmstat(1M)`、`logadm(1M)`、`syslogd(1M)`、`attributes(5)`、`devchassis(5)`
《Oracle Solaris 管理：常见任务》

引用名

fmd – Fault Manager 守护进程

用法概要

`/usr/lib/fm/fmd/fmd [-V] [-f file] [-o opt=val] [-R dir]`

描述

fmd 是运行在每个 Solaris 系统后台的守护进程。fmd 接收与系统软件检测到的问题有关的遥测信息，诊断这些问题，启动主动式自我修复活动，如禁用故障组件等。Fault Manager 还会在适当的情况下向 [syslogd\(1M\)](#) 服务发送消息，通知管理员已检测到故障。该消息可指示管理员查阅 Oracle 的 Web 站点 <https://support.oracle.com> 中的知识文章，了解有关故障影响的详细信息以及相应的应对措施。

Fault Manager 诊断的每个问题都分配了一个通用唯一标识符 (Universal Unique Identifier, UUID)。UUID 可在各种系统集之间唯一标识该特定问题。可以使用 [fmdump\(1M\)](#) 实用程序查看 Fault Manager 诊断的问题列表，以及它们的 UUID 和知识文章消息标识符。可以使用 [fmadm\(1M\)](#) 实用程序查看系统中确认存在故障的资源。可以使用 [fmstat\(1M\)](#) 实用程序报告 Fault Manager 保存的统计信息。当 Solaris 引导时会自动启动 Fault Manager，因此没有必要直接使用 fmd 命令。Oracle 的 Web 站点介绍了当前 Solaris 上可用于 Fault Manager 的功能的详细信息。

通知服务

syslog（软件包 `system/fault-management`）

新诊断的标准通知机制是通过 syslog 实现的，即，使用在与 fmd 本身相同的软件包中交付的 `syslog-msgs fmd` 模块。

缺省情况下，只有新的问题诊断是通过 `syslog-msgs` 发送消息，并使用如下表所列的 syslog 工具和严重性。管理员可使用 [svccfg\(1M\)](#) 请求通过 `syslog-msgs` 发送有关处于问题解决方案生命周期中的其他事件的消息。

svccfg setnotify event syslog:{active|inactive}

有关其他详细信息，请参见 [svccfg\(1M\)](#)。

Event	Disposition	Facility	Severity
problem-diagnosed	active	LOG_DAEMON	LOG_ERR
problem-updated	inactive	LOG_DAEMON	LOG_NOTICE
problem-repaired	inactive	LOG_DAEMON	LOG_NOTICE
problem-resolved	inactive	LOG_DAEMON	LOG_NOTICE

Email（软件包 `system/fault-management/smtp-notify`）

必须安装附加软件包时，可以通过电子邮件发送通知。SMF 服务

`svc:/system/fm/smtp-notify:default` 是通过

`system/fault-management/smtp-notify` 软件包以及通过 [svccfg\(1M\)](#) 配置的通知首选项进行交付的。有关其他详细信息，请参见 [sasinfo\(1M\)](#)。请注意，除了配置以上列出的问题生命周期事件（`problem-diagnosed` 等）的通知首选项，还可以通过 [svccfg\(1M\)](#) 配置该机制以提供 SMF 实例状态转换和其他事件的通知。

SNMP（软件包 `system/fault-management/snmp-notify`）

使用 SNMP 陷阱的新事件通知是一个通过 `system/fault-management/snmp-notify` 软件包交付的选项。服务 `svc:/system/fm/snmp-notify:default` 负责提出用于问题生命周期和其他指定事件（如果进行了配置，还包括 SMF 实例状态转换事件）的

SNMP陷阱。有关其他详细信息，请参见 [snmp-notify\(1M\)](#)。

全局和非全局 Solaris 区域 Fault Manager 服务 `svc:/system/fmd:default` 在全局 Solaris 区域和非全局 Solaris 区域中均进行了配置。在非全局区域中，各种面向硬件的 Fault Manager 模块未交付，因此它是在对应区域运行的精简 Fault Manager。在非全局区域中，Fault Manager 专注于软件事件。

选项 支持以下选项

- `-f file`
在搜索任何缺省的 Fault Manager 配置文件之前，读取指定的配置 *file*。
- `-o opt=value`
将指定的 Fault Manager 选项设置为指定的值。当前 Fault Manager 选项是专用接口；有关专用接口的信息，请参见 [attributes\(5\)](#)。
- `-R dir`
对 Fault Manager 评估的所有路径名称使用指定的根目录，而不是缺省的根目录 (`/`)。
- `-V`
将 Fault Manager 的版本输出到 `stdout` 中，然后退出。

退出状态 将返回下列退出值：

- 0 成功完成
- 1 出现一个阻止 Fault Manager 进行初始化的错误，例如无法打开遥测传输。
- 2 指定的命令行选项无效。

文件

<code>/etc/fm/fmd</code>	Fault Manager 配置目录
<code>/usr/lib/fm/fmd</code>	Fault Manager 库目录
<code>/var/fm/fmd</code>	Fault Manager 日志目录

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/fault-management
接口稳定性	Committed（已确定）

另请参见 [svcs\(1\)](#)、[fmadm\(1M\)](#)、[fmdump\(1M\)](#)、[fmstat\(1M\)](#)、[sasinfo\(1M\)](#)、[snmp-notify\(1M\)](#)、[svccfg\(1M\)](#)
<http://www.sun.com/msg/>

附注

Fault Manager 由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

```
svc:/system/fmd:default
```

可以使用 [svcs\(1\)](#) 命令来查询服务的状态。管理员不应该禁用 Fault Manager 服务。

引用名	fmdump – 故障管理日志查看器
用法概要	<pre>fmdump [[-e -i -I] -A] [-f] [-mvVp] [-c class] [-R root] [-t time] [-T time] [-u uid] [-n name[.name]*[=value]] [-E ENA] [file] ...</pre>
描述	fmdump 实用程序可用于显示与 Solaris Fault Manager fmd(1M) 有关的任何日志文件的内容。Fault Manager 在每个 Solaris 系统的后台运行。它接收与系统软件检测到的问题有关的遥测信息，诊断这些问题，启动主动式自我修复活动，如禁用故障组件等。 fmdump 实用程序不用作 Fault Manager 的主要管理接口。就此目的，请使用 fmadm（并参见 fmadm(1M)）。fmdump 实用程序只会转储 Fault Manager 历史日志（几乎没有更多解释），且可以包含没有解释的实现详细信息。请参见下面的“属性”部分。 Fault Manager 维护几组日志文件，供服务人员以及在少数情况下供管理员使用。 错误日志 记录错误遥测以及系统检测到的问题症状的日志。 信息日志 记录信息事件的日志。此类日志体现为两组日志文件：高价值信息事件和其他信息事件。 故障日志 记录故障诊断信息以及被认为解释了错误日志和信息日志中记录的症状之问题的日志。 日志文件组包括当前活动的日志文件，以及该组中数量不定的旧轮转日志文件。所有日志都由 logadm(1M) 管理，/etc/logadm.conf 下有相应的条目。 请注意，fmdump 实用程序会转储目标组的当前日志文件，以及所有的轮转日志文件。因此，它会显示完整的日志历史。特别是对于故障日志，必须认识到 fmdump 会显示曾经诊断出的所有问题而不是仅限制为当前的问题，这一点非常重要。（使用 fmadm 实用程序可了解此信息。） 缺省情况下，fmdump 会显示故障日志的内容，其中记录了 Fault Manager 或其某个组件模块做出的每个诊断的结果。错误日志可使用 -e 进行选择，信息日志使用 -i 进行选择，而高价值信息日志使用 -I 进行选择。或者，也可以在命令行中以 [file] 形式指定具体的日志文件路径（此时将仅转储此文件，不会查找日志的轮转版本）。用户也可以使用选项 -A 聚合所有日志，或是命令行上列出的一组日志文件路径。 下面是一个缺省 fmdump 显示示例： <pre># fmdump TIME UUID SUNW-MSG-ID EVENT Mar 23 14:06:35.2682 0a11a1a7-a8ce-c941-8527-8d7a9d320071 ZFS-8000-CS Diagnosed Mar 25 14:51:41.2261 0a11a1a7-a8ce-c941-8527-8d7a9d320071 FMD-8000-4M Repaired Mar 25 14:51:41.2523 0a11a1a7-a8ce-c941-8527-8d7a9d320071 FMD-8000-6U Resolved May 31 23:35:39.9146 c63ac52e-506b-c1cc-e965-ff3b8544490d SMF-8000-YX Diagnosed ...</pre>

（显示的输出在 80 个字符以内换行。）

这会转储故障日志，因为没有命令行选项或参数选择其他任何日志。故障日志记录了在 Fault Manager 或其组件模块诊断出的问题的整个生命周期，从最初问题诊断到问题解决的所有情况。

故障日志中记录的每个问题都具有以下标识：

- 事件时间戳—描述问题的生命周期状态更改。
- 通用唯一标识符 (UUID)—可在任何系统组之间唯一标识此特定问题。描述给定问题的生命周期状态更改的所有事件都将使用同一个 UUID（如上：我们看到最初诊断以及随后的修复和解决全部引用相同的问题 UUID）。
- 消息标示符—可用于访问 Oracle Web 站点上对应的知识文章：<http://www.oracle.com/us/sun/msg/>

如果问题需要管理员、服务技术人员采取操作，或是影响了系统行为，Fault Manager 也会向 `syslogd(1M)` 发布一条用户可读的消息。该消息提供一个问题总结，以及 Oracle Web 站点上的知识文章参考：<http://www.oracle.com/us/sun/msg/>。fmdump 实用程序可以使用 `-m` 选项，以类似格式将 `fltlog` 条目转储到提供的 `syslog`。

您可以使用 `-v` 和 `-V` 选项，将显示的单行总结扩展成每个日志记录事件的详细信息的多行显示。`-p` 选项可与 `-v` 一起使用，以请求“更漂亮”的输出。

可以使用 `-c`、`-t`、`-T`、`-n` 和 `-u` 选项来过滤输出，方法是仅选择那些与指定类、时间范围或 `uuid` 匹配的事件。如果命令行中存在多个过滤器选项，这些选项会组合在一起，从而仅显示由选项逻辑 AND 运算后选择的事件。如果命令行中存在同一过滤器选项的多个实例，同类选项会组合在一起以显示由选项逻辑 OR 运算后选择的事件。例如，命令：

```
# fmdump -u uuid1 -u uuid2 -t 02Dec09
```

...选择那些属性为 (`uuid1` OR `uuid2`) AND（时间在 2009 年 2 月 2 日当天或以后）的事件。

选项

所支持的选项如下：

-A

执行日志聚合。如果命令行中列出一个或多个日志文件路径，则会聚合这些文件，否则会聚合所有已知的日志类型，包括所有的 `logadm` 轮转文件。日志按照时间顺序合并，但是如果两个记录来自同一个日志文件，聚合时会按照它们在原始日志文件中的顺序（即它们被 Fault Manager 接收和处理的顺序，基本上就是时间顺序）。

不能将其他的日志组选择选项 `-e`、`-i` 或 `-I` 与 `-A` 一起使用。可以使用 `-c`、`-t`、`-T` 和 `-n` 等过滤器选项，但不能使用 `-u`。可以使用输出选项 `-v`、`-V` 和 `-p`，但不能使用 `-m`。选项 `-f` 将跟随所有选择的日志。

-c class

选择与指定类匹配的事件。class 参数可以使用 [sh\(1\)](#) 中描述的 glob 模式匹配语法。class 代表说明遥测事件类型的分层分类字符串。要了解关于 Oracle 遥测协议的更多信息，请访问 Oracle Web 站点 <http://www.oracle.com/us/sun/msg/>。

-e

显示故障管理错误日志（而不是故障日志）中的事件。

错误日志包含 Sun 自动诊断软件使用的“专用”遥测信息。记录此信息是为了便于问题的事后分析和事件重放，而不应对其进行解析或用于开发脚本和其他工具。请参见 [attributes\(5\)](#)，了解有关 Oracle Private（专用）接口规则的信息。

-E ENA

选择与指定 ENA 值匹配的任何一代事件。对于支持 ENA 的检测器，可以使用该选项显示与同一操作有关的多个事件。

-f

通过等待附加数据，跟踪日志文件（如果使用 -A，则有多文件）的增长。fmdump 进入无限循环，其间会休眠一秒钟，尝试读取并格式化日志文件中的新数据，然后返回休眠状态。可以通过发送中断 (Control-C) 随时终止此循环。

-m

输出与故障日志中的每个条目有关的本地化诊断消息。

-n name[.name]*[=value]

选择其属性具有匹配名称（也可以是匹配值）的日志事件（从命令行上选择的日志中进行选择）。对于字符串属性，值可以是一个正则表达式匹配项。正则表达式语法在 [regex\(5\)](#) 手册页的“扩展正则表达式”部分中进行了介绍。使用以下字符时要小心：

\$ * { ^ | () \

...或正则表达式，因为它们对 shell 是有意义的。最安全的方法是，将它们用单引号括起来。对于数字属性，其值可以是八进制、十六进制或十进制。

-p

与 -v（非常详细）选项组合使用，请求实际执行可用的漂亮打印选项。

-R dir

对 fmdump 访问的日志文件使用指定的根目录，而不是默认的根目录 (/)。

-t time

选择在指定时间或之后发生的事件。可以使用以下任一格式指定时间：

mm/dd/yy hh:mm:ss

月、日、年、小时（24 小时格式）、分钟和秒。以任意数量的空格分隔日期和时间。应该将参数用引号括起来，这样 Shell 会将两个字符串解释为一个参数。

mm/dd/yy hh:mm

月、日、年、小时（24 小时格式）和分钟。以任意数量的空格分隔日期和时间。应该将参数用引号括起来，这样 Shell 会将两个字符串解释为一个参数。

mm/dd/yy

指定月、日、年当天的凌晨 00:00:00。

ddMonyy hh:mm:ss

日、月名、年、小时（24 小时格式）、分钟和秒。以任意数量的空格分隔日期和时间。应该将参数用引号括起来，这样 Shell 会将两个字符串解释为一个参数。

ddMonyy hh:mm

日、月名、年、小时（24 小时格式）和分钟。以任意数量的空格分隔日期和时间。应该将参数用引号括起来，这样 Shell 会将两个字符串解释为一个参数。

Mon dd hh:mm:ss

当前年份的月、日、小时（24 小时格式）、分钟和秒。

yyyy-mm-dd [T hh:mm[:ss]]

年、月、日、小时（24 小时格式，可选）、分钟和秒。可以选择忽略秒或小时、分钟和秒。

ddMonyy

指定日、月名、年的凌晨 00:00:00。

hh:mm:ss

当天的小时（24 小时格式）、分钟和秒。

hh:mm

当天的小时（24 小时格式）和分钟。

Tns | Tnsec

T 纳秒以前，其中 *T* 是按照以 10 为底数的形式指定的整数值。

Tus | Tusec

T 微秒以前，其中 *T* 是按照以 10 为底数的形式指定的整数值。

Tms | Tmsec

T 毫秒以前，其中 *T* 是按照以 10 为底数的形式指定的整数值。

Ts | Tsec

T 秒以前，其中 *T* 是按照以 10 为底数的形式指定的整数值。

Tm | Tmin

T 分钟以前，其中 *T* 是按照以 10 为底数的形式指定的整数值。

Th | Thour

T 小时以前，其中 *T* 是按照以 10 为底数的形式指定的整数值。

Td | Tday

T 日以前，其中 *T* 是按照以 10 为底数的形式指定的整数值。

您可以向任意 *-t* 选项参数附加 *.n* 格式的小数，表示超出指定时间的秒数（小数形式）。

-T *time*

选择在指定时间或之前发生的事件。*time* 可以使用针对 **-t** 选项描述的任何时间格式进行指定。

-u *uuid*

在故障日志中选择与指定 *uuid* 完全匹配的问题诊断事件。为了便于识别，每个诊断都与一个通用唯一标识符 (Universal Unique Identifier, UUID) 相关联。**-u** 选项可以与其他选项（如 **-v**）组合使用，以显示与特定诊断相关的所有详细信息。请注意，多个故障日志事件可以与同一问题诊断 UUID 相关—描述一个问题生命周期的所有事件（从最初诊断到最终解决）引用同一个问题 UUID。

如果 **-e** 选项和 **-u** 选项都存在，会显示被指定诊断交叉引用的错误事件。

-v

显示详细的事件详细信息。事件显示会被放大，以显示选定事件的其他公共成员。

-V

显示非常详细的事件详细信息。事件显示会被放大，以显示与每个事件相关的名称-值对列表中的每个成员。此外，对于故障日志，事件显示还包括一个与诊断相关联的对应错误的交叉引用列表。

使用 **-p** 和 **-v** 选项可请求漂亮打印。

操作数

支持下列操作数：

file 指定要显示的备用日志文件（如果使用 **-A**，则为多个文件），而不是显示系统故障日志。**fmdump** 实用程序会自动确定指定日志的类型，并为选定日志生成相应的输出。

示例

示例 1 从 fmd 日志检索给定的类

使用以下任一命令从 fmd 日志中检索有关指定类的信息。完整的类名称是 `ereport.io.ddi.context`。

```
# fmdump -Ve -c 'ereport.io.ddi.context'
# fmdump -Ve -c 'ereport.*.context'
# fmdump -Ve -n 'class=ereport.io.ddi.context'
# fmdump -Ve -n 'class=ereport.*.context'
```

前面的任一命令都会生成以下输出：

```
Oct 06 2007 11:53:20.975021712 ereport.io.ddi.context
    nvlist version: 0
        class = ereport.io.ddi.context
        ena = 0x1b03a15ecf00001
        detector = (embedded nvlist)
    nvlist version: 0
        version = 0x0
        scheme = dev
        device-path = /
```

示例1 从fmd日志检索给定的类 (续)

```
(end detector)

__ttl = 0x1
__tod = 0x470706b0 0x3a1da690
```

示例2 从fmd日志检索特定的检测器设备路径
以下命令从fmd日志检索检测器设备路径。

```
# fmdump -Ve -n 'detector.device-path=.* /disk@1,0$'
Oct 06 2007 12:04:28.065660760 ereport.io.scsi.disk.rqs
nvlist version: 0
    class = ereport.io.scsi.disk.rqs
    ena = 0x453ff3732400401
    detector = (embedded nvlist)
        nvlist version: 0
            version = 0x0
            scheme = dev
            device-path = /pci@0,0/pci1000,3060@3/disk@1,0
        (end detector)

__ttl = 0x1
__tod = 0x4707094c 0x3e9e758
```

退出状态

将返回下列退出值：

- 0 成功完成。日志文件中的所有记录都已成功检查。
- 1 发生了致命错误。这会阻止对任何日志文件数据的检查，例如，无法打开指定文件。
- 2 指定的命令行选项无效。
- 3 日志文件成功打开，但是有一个或多个日志文件记录无法显示，原因可能是I/O错误，或是记录本身的格式不正确。fmdump对每个无法显示的记录发出一条警告消息，然后继续执行，并尝试显示其他记录。

文件

/var/fm/fmd	故障管理日志目录
/var/fm/fmd/errlog	故障管理错误日志
/var/fm/fmd/fltlog	故障管理故障日志
/var/fm/fmd/infolog_hival	高价值信息日志
/var/fm/fmd/infolog	信息日志

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/fault-management
接口稳定性	请参见下文。

命令行选项为 `Uncommitted`（未确定）。用户可读错误日志和信息日志输出为 `Private`（专用）。用户可读故障日志输出为 `Uncommitted`（未确定）。

另请参见 [sh\(1\)](#)、[fmadm\(1M\)](#)、[fmd\(1M\)](#)、[fmstat\(1M\)](#)、[logadm\(1M\)](#)、[syslogd\(1M\)](#)、[libexacct\(3LIB\)](#)、[at](#)
[《Oracle Solaris 管理：安全服务》](#)
<http://www.oracle.com/us/sun/msg/>

附注 故障日志包含对存储在错误日志中的记录的引用，可以通过 `fmdump -v` 显示这些记录引用，以了解特定故障诊断中使用的错误。如果在日志轮转中对某个错误日志进行了重命名，会保留这些链接。它们会在移除错误日志文件，或是移动到其他文件系统目录时被破坏。`fmdump` 不能显示此类断开链接的错误信息。它会继续显示故障日志中的信息。

引用名

fmstat – 报告故障管理模块统计信息

用法概要fmstat [-astZ] [-d u | d] [-m *module*] [*interval* [*count*]]**描述**

管理员和服务人员可以使用 **fmstat** 实用程序报告与 Solaris Fault Manager **fmd(1M)** 及其相关的模块集关联的统计信息。Fault Manager 在每个 Solaris 系统的后台运行。它接收与系统软件检测到的问题有关的遥测信息，诊断这些问题，启动主动式自我修复活动，如禁用故障组件等。

可使用 **fmstat** 查看当前参与故障管理的诊断引擎和代理的统计信息。**fmd(1M)**、**fmadm(1M)** 和 **fmdump(1M)** 的文档对于观察故障管理活动的工具进行了详细的描述。

如果存在 **-m** 选项或 **-t** 选项，**fmstat** 会报告指定故障管理模块保存的所有统计信息。可使用 **fmadm config** 获取模块列表。

如果 **-m** 选项不存在，**fmstat** 会报告它的每个客户机模块的以下统计信息：

module	故障管理模块的名称，如 fmadm config 所报告。
ev_recv	模块接收的遥测事件数。
ev_acpt	模块接受的与诊断相关的事件数。
wait	等待模块检查的遥测事件平均数。
svc_t	模块接收的遥测事件平均服务时间（以毫秒为单位）。
%w	存在等待模块检查的遥测事件的时间的百分比。
%b	模块忙于处理遥测事件的时间的百分比。
open	模块拥有的活动事例数（打开问题调查）。
solve	自从装入该模块后，它所解决的事例总数。
memsz	该模块当前所分配的动态内存量。
bufsz	该模块当前所分配的持久性缓冲区空间大小。

fmstat 实用程序要求用户拥有 **SYS_ADMIN** 特权。请参见《[Oracle Solaris 管理：安全服务](#)》了解有关如何配置 Solaris 特权的更多信息。

选项

所支持的选项如下：

-a	输出模块的所有统计信息，包括那些由 fmd 代表其保存的统计信息。如果 -a 选项不存在，则只报告由模块保存的统计信息。如果 -a 选项没有与 -m module 一起使用，将会显示一组与 fmd 关联的全局统计信息。
-d u d	显示时间戳。 指定 u 表示时间的内部表示形式的印刷表示形式。请参见 time(2) 。指定 d 表示标准日期格式。请参见 date(1) 。

- m *module* 输出与指定故障管理模块关联的统计信息报告，而不是缺省统计信息报告。模块可以发布任意一组统计信息以帮助 Oracle 自己维护故障管理软件。模块统计信息构成了一个 Private（专用）接口。请参见 [attributes\(5\)](#)，了解有关 Oracle Private（专用）接口规则的信息。不能根据故障管理模块统计信息的值来编写脚本，因为它们可能在不事先通知的情况下进行更改。
- s 输出与模块关联的 Soft Error Rate Discrimination (SERD) 引擎报告，而不是缺省模块统计信息报告。SERD 引擎是故障管理软件所使用的结构，用来确定是否超出了在特定时间测量为 *N* 个事件的统计信息阈值 *T*。-s 选项只能与 -m 选项结合使用。
- t 输出与每个故障管理事件传输关联的统计信息报告。每个故障管理模块可以实现一个或多个事件传输。
- T 输出与每个故障管理事件传输关联的授权信息表。如果存在 -m 选项，则只会显示与指定模块关联的传输。
- z 在与指定故障管理模块关联的报告中，省略具有零值的统计信息。-z 选项只能与 -m 选项结合使用。

操作数

- 支持下列操作数：
- count* 只输出计数报告，然后退出。
 - interval* 每隔 *interval* 秒输出一份新的报告。

如果未指定 *interval* 和 *count*，将会输出一份报告，然后 **fmstat** 退出。如果指定了 *interval* 但未指定 *count*，则 **fmstat** 会无限期地每隔 *interval* 秒输出报告，直至命令中断。

退出状态

- 将返回下列退出值：
- 0 成功完成。
 - 1 发生了致命错误。致命错误可能是无法与 [fmd\(1M\)](#) 进行通信。也可能是没有足够的特权来执行请求的操作。
 - 2 指定的命令行选项无效。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/fault-management
接口稳定性	请参见下文。

命令行选项为 **Committed**（已确定）。用户可读缺省报告为 **Uncommitted**（未确定）。用户可读模块报告为 **Private**（专用）。

另请参见

[fmadm\(1M\)](#)、[fmd\(1M\)](#)、[fmdump\(1M\)](#)、[attributes\(5\)](#)

《Oracle Solaris 管理：安全服务》

引用名 fmthard – 在硬盘上填充标签

用法概要

SPARC fmthard -d *data* | -n *volume_name* | -s *datafile* [-i] /dev/rdisk/c?
[t?] d?s2

x86 fmthard -d *data* | -n *volume_name* | -s *datafile* [-i] /dev/rdisk/c?
[t?] d?s2

描述

fmthard 命令更新硬盘上的 VTOC (Volume Table of Contents, 卷目录), 并且在 x86 系统上将向 Solaris fdisk 分区添加引导信息。必须使用选项 -s *datafile*、-d *data* 或 -n *volume_name* 中的一个或多个才能请求修改磁盘标签。要打印磁盘标签内容, 请参见 [prtvtoc\(1M\)](#)。/dev/rdisk/c?[t?]d?s2 文件必须是要安装新标签的设备的字符特殊文件。在 x86 系统上, 在运行 fmthard 之前必须在驱动器上运行 [fdisk\(1M\)](#)。

如果使用 x86 系统, 请注意, 此页中的术语“分区”是指 x86 计算机 x86 fdisk 分区内的分片。不要将 fmthard 创建的分片与 fdisk 创建的分片混淆。

选项

所支持的选项如下:

-d *data*

此选项的 *data* 参数是一个表示当前 VTOC 中特定分区信息的字符串。该字符串必须为 *part:tag:flag:start:size* 格式, 其中 *part* 是分区号, *tag* 是分区的 ID TAG, *flag* 是权限标志集, *start* 是分区的起始扇区号, *size* 是分区中的扇区数。有关这些字段的更多信息, 请参见下面 *datafile* 的介绍。

-i

此选项允许命令创建所需的 VTOC 表, 但将信息打印到标准输出中, 而不是修改磁盘上的 VTOC。

-n *volume_name*

此选项用于为磁盘提供 *volume_name*, 最多 8 个字符长度。

-s *datafile*

此选项用于根据用户创建的 *datafile* 填充 VTOC。如果 *datafile* 为 - (连字符), fmthard 将会从标准输入中读取。*datafile* 格式在下面进行说明。此选项会导致所有磁盘分区时间标记字段均设置为零。

按照惯例, fmthard 生成的每个 VTOC 也具有对应整个磁盘的分区 2。如果 *datafile* 的输入未指定分区 2 条目, 将会在 VTOC 中自动创建缺省分区 2 条目, 该缺省条目具有标记 V_BACKUP 且大小与整个磁盘的大小相等。

datafile 包含每个分区的一个规范行, 从分区 0 开始。每行以换行符 (\n) 来分隔。如果一行的第一个字符为星号 (*), 则将该行视为注释。每一行由与位置有关、以空格分隔并具有以下格式的条目组成:

partition tag flag starting_sector size_in_sectors

其中这些条目具有以下值:

分区

分区号。目前，对于 Solaris SPARC，一个磁盘最多可具有 8 个分区，0–7。即使 *partition* 字段具有 4 位，但目前只使用 3 位。对于 x86，会使用 4 位以允许分片 0–15。每个 Solaris fdisk 分区最多可具有 16 个分片。

tag (标记)

分区标记：十进制数。以下是保留代码：0 (V_UNASSIGNED)、1 (V_BOOT)、2 (V_ROOT)、3 (V_SWAP)、4 (V_USR)、5 (V_BACKUP)、6 (V_STAND)、7 (V_VAR)、8 (V_HOME)、12 (V_SYSTEM) 和 24 (V_BIOS_BOOT)。

flag

该标志允许将分区标记为无法挂载或只读，掩码为：V_UNMNT 0x01 和 V_RDONLY 0x10。对于可挂载分区，使用 0x00。

starting_sector

分区的起始扇区号（十进制）。

size_in_sectors

分区占用的扇区数（十进制）。

可将 prtvtoc 命令的输出保存到一个文件中，编辑该文件，然后将它作为 -s 选项的 *datafile* 参数使用。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见

[uname\(1\)](#)、[format\(1M\)](#)、[prtvtoc\(1M\)](#)、[attributes\(5\)](#)

仅限于 x86

[fdisk\(1M\)](#)、[installgrub\(1M\)](#)

附注

当覆盖现有 VTOC 时需要特别小心，因为不正确的条目可能导致当前数据无法访问。作为预防措施，应保存旧的 VTOC。

对于低于 2 兆字节的磁盘，fmthard 无法在无标签磁盘上编写 VTOC。要实现此操作，可使用 [format\(1M\)](#)。

引用名	format – disk partitioning and maintenance utility
用法概要	<pre>format [-f <i>command-file</i>] [-l <i>log-file</i>] [-x <i>data-file</i>] [-d <i>disk-name</i>] [-t <i>disk-type</i>] [-p <i>partition-name</i>] [-s] [-m] [-M] [-e] [<i>disk-list</i>] format -L <i>label-type</i> -d <i>disk-name</i></pre>
描述	format enables you to format, label, repair, and analyze disks on your system. Unlike previous disk maintenance programs, format runs under SunOS. Because there are limitations to what can be done to the system disk while the system is running, format is also supported within the memory-resident system environment. For most applications, however, running format under SunOS is the more convenient approach. format first uses the disk list defined in <i>data-file</i> if the -x option is used. format then checks for the FORMAT_PATH environment variable, a colon-separated list of filenames and/or directories. In the case of a directory, format searches for a file named format.dat in that directory; a filename should be an absolute pathname, and is used without change. format adds all disk and partition definitions in each specified file to the working set. Multiple identical definitions are silently ignored. If FORMAT_PATH is not set, the path defaults to /etc/format.dat. <i>disk-list</i> is a list of disks in the form c?t?d?, or /dev/rdisk/c?t?d?s?, /dev/chassis/?/disk. With the last two forms, shell wildcard specifications are supported. For example, specifying /dev/rdisk/c2* causes format to work on all drives connected to controller c2 only. If no <i>disk-list</i> is specified, format lists all the disks present in the system that can be administered by format. Removable media devices are listed only when users execute format in expert mode (option -e). This feature is provided for backward compatibility. Use rmformat(1) for rewritable removable media devices.
选项	The following options are supported: -d <i>disk-name</i> Specify which disk should be made current upon entry into the program. The disk is specified by its logical name (for instance, -d c0t1d0 or /dev/chassis/SYS/HD0/disk). This can also be accomplished by specifying a single disk in the disk list. -e Enable SCSI expert menu. Note this option is not recommended for casual use. -f <i>command-file</i> Take command input from <i>command-file</i> rather than the standard input. The file must contain commands that appear just as they would if they had been entered from the keyboard. With this option, format does not issue continue? prompts; there is no need to specify y(es) or n(o) answers in the <i>command-file</i>. In non-interactive mode, format does not initially expect the input of a disk selection number. The user

- must specify the current working disk with the `-d disk-name` option when `format` is invoked, or specify `disk` and the disk selection number in the *command-file*.
- `-l log-file` Log a transcript of the `format` session to the indicated *log-file*, including the standard input, the standard output and the standard error.
 - `-L label-type` Immediately, and non-interactively, write a default label of type *label-type*, to the disk specified with `-d`. *label-type* must be either `efi` or `vtoc`. Existing label, if any, will be overwritten with *label-type*. On an x86 machine, the whole disk will default to one Solaris partition labeled with *label-type*; all `fdisk` partitions will be lost.
 - `-m` Enable extended messages. Provides more detailed information in the event of an error.
 - `-M` Enable extended and diagnostic messages. Provides extensive information on the state of a SCSI device's mode pages, during formatting.
 - `-p partition-name` Specify the partition table for the disk which is current upon entry into the program. The table is specified by its name as defined in the data file. This option can be used only if a disk is being made current, and its type is either specified or available from the disk label.
 - `-s` Silent. Suppress all of the standard output. Error messages are still displayed. This is generally used in conjunction with the `-f` option.
 - `-t disk-type` Specify the type of disk which is current upon entry into the program. A disk's type is specified by name in the data file. This option can only be used if a disk is being made current as described above.
 - `-x data-file` Use the list of disks contained in *data-file*.

用法

When you invoke `format` with no options or with the `-e`, `-l`, `-m`, `-M`, or `-s` options, the program displays a numbered list of available disks and prompts you to specify a disk by list number. If the machine has more than a screenful of disks, press `SPACE` to see the next screenful of disks.

You can specify a disk by list number even if the disk is not displayed in the current screenful. For example, if the current screen shows disks 11-20, you can enter 25 to specify the twenty-fifth disk on the list. If you enter a number for a disk that is not currently displayed, `format` prompts you to verify your selection. If you enter a number from the displayed list, `format` silently accepts your selection.

After you specify a disk, `format` displays its main menu. This menu enables you to perform the following tasks:

analyze	Run read, write, compare tests, and data purge. The data purge function implements the National Computer Security Center Guide to Understanding Data Remnance (NCSC-TG-025 version 2) Overwriting Algorithm. See NOTES.
backup	Search for backup labels.
cache	Enable, disable, and query the state of the write cache and read cache. This menu item only appears when <code>format</code> is invoked with the <code>-e</code> option, and is only supported on SCSI devices..
current	Display the device name, the disk geometry, and the pathname to the disk device.
defect	Retrieve and print defect lists. This option is supported only on SCSI devices. IDE disks perform automatic defect management. Upon using the defect option on an IDE disk, you receive the message: Controller does not support defect management or disk supports automatic defect management.
disk	Choose the disk that will be used in subsequent operations (known as the current disk.)
fdisk	Run the <code>fdisk(1M)</code> program to create a fdisk partition for Solaris software (x86 based systems only).
format	Format and verify the current disk. This option is supported only on SCSI devices. IDE disks are pre-formatted by the manufacturer. Upon using the format option on an IDE disk, you receive the message: Cannot format this drive. Please use your manufacturer-supplied formatting utility.
inquiry	Display the vendor, product name, and revision level of the current drive.
label	Write a new label to the current disk.
partition	Create and modify slices.
quit	Exit the format menu.
repair	Repair a specific block on the disk.
save	Save new disk and slice information.
type	Select (define) a disk type.
verify	Read and display labels. Print information such as the number of cylinders, alternate cylinders, heads, sectors, and the partition table.
volname	Label the disk with a new eight character volume name.

- 环境变量

FORMAT_PATH

a colon-separated list of filenames and/or directories of disk and partition definitions. If a directory is specified, format searches for the file format.dat in that directory.
- 文件

/etc/format.dat

default data file
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[fmthard\(1M\)](#), [prtvtoc\(1M\)](#), [rmformat\(1\)](#), [format.dat\(4\)](#), [attributes\(5\)](#), [sd\(7D\)](#)
[《Oracle Solaris 管理：常见任务》](#)

x86 Only

[fdisk\(1M\)](#)

警告

When the format function is selected to format the Maxtor 207MB disk, the following message displays:

Mode sense page(4) reports rpm value as 0, adjusting it to 3600

This is a drive bug that may also occur with older third party drives. The above message is not an error; the drive will still function correctly.

Cylinder 0 contains the partition table (disk label), which can be overwritten if used in a raw disk partition by third party software.

format supports writing EFI-compliant disk labels in order to support disks or LUNs with capacities greater than one terabyte. However, care should be exercised since many software components, such as filesystems and volume managers, are still restricted to capacities of one terabyte or less. See the [《Oracle Solaris 管理：常见任务》](#) for additional information.

By default, on an unlabeled disk, EFI labels will be written on disks larger than 2 TB. When format is invoked with the -e option, on writing the label, the label type can be chosen. Booting is not currently supported on a disk with an EFI label.

附注

format provides a help facility you can use whenever format is expecting input. You can request help about what information is expected by simply entering a question mark (?) and format prints a brief description of what type of input is needed. If you enter a ? at the menu prompt, a list of available commands is displayed.

For SCSI disks, formatting is done with both Primary and Grown defects list by default. However, if only Primary list is extracted in defect menu before formatting, formatting will be done with Primary list only.

Changing the state of the caches is only supported on SCSI devices, and not all SCSI devices support changing or saving the state of the caches.

The NCSC-TG-025 algorithm for overwriting meets the DoD 5200.28-M (ADP Security Manual) Eraser Procedures specification. The NIST Guidelines for Media Sanitization (NIST SP 800-88) also reference this algorithm.

引用名 fpsd – Fp-scrubber daemon

用法概要 /usr/lib/fps/fpsd

描述 fpsd is a user-level daemon that periodically runs non-intrusive tests to validate proper functioning of FPU (Floating Point Unit) hardware in the system. A fault management action is initiated by means of [fmd\(1M\)](#), in case an error is detected by the test.

 fpsd is not supported on all Sun systems. On a system on which it is not supported, fpsd is automatically disabled.

退出状态 The following exit values are returned:

 0
 Successful completion.

 non-zero
 An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/floating-point-scrubber
Interface Stability	Uncommitted

另请参见 [svcs\(1\)](#), [fmd\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The fpsd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

 svc:/system/fpsd:default

 You can use [svccfg\(1M\)](#) to change the default fpsd behavior:

Property Name	Type	Description
-----	----	-----
config/exclude_cpus	astring	comma delimited list of CPU IDs to be excluded from proactive testing.

 Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

 Disabling the service can increase system's vulnerability to silent data corruption, if a fault were to develop in an FPU within the system.

引用名	fruadm – prints and updates customer data associated with FRUs
用法概要	<pre> /usr/platform/sun4u/sbin/fruadm /usr/platform/sun4u/sbin/fruadm -l /usr/platform/sun4u/sbin/fruadm [-r] path [text] </pre>
描述	fruadm prints or sets the customer data for Field-Replaceable Units (FRUs). Without arguments, fruadm prints the paths of all FRU ID-capable FRUs (containers) in the system, along with the contents of the customer data record, if present, for each such FRU; for FRUs without customer data, fruadm prints only the container's path. Only a privileged user can create or update data in containers. The privileges required to perform these write operations are hardware dependent. Typically, a default system configuration restricts write operations to the superuser or to the platform-administrator user.
选项	The following options are supported: -l List the system's frutree paths. -r Recursively display or update the data for all containers rooted at the argument <i>path</i>.
操作数	The following operands are supported: <i>path</i> A full or partial system frutree path for or under which to print or set the customer data. The first field of each line of output of fruadm -l gives the valid full frutree paths for the system. Paths can include shell meta-characters; such paths should be quoted appropriately for the user's shell. For partial paths, the first matching full path is selected for display or update. Without the -r option, the path must be that of a container; with the -r option, all containers (if any) under <i>path</i> will be selected. <i>text</i> Up to 80 characters of text set as the customer data. If the text contains white space or shell metacharacters, it should be quoted appropriately for the user's shell.
示例	示例 1 Displaying All Customer Data The following example prints all customer data available from FRUs on the system. For containers with no customer data, only the containers' paths will be listed. <pre>example% fruadm</pre> 示例 2 Displaying Customer Data For a Single FRU The following command prints the customer data, if present, for the specified FRU: <pre>example% fruadm /frutree/chassis/system-board</pre>

示例 3 Displaying Customer Data For a Single FRU

The following command prints the customer data, if present, for the first mem-module found:

```
example% fruadm mem-module
```

示例 4 Setting Customer Data

The following example sets the customer data for a FRU:

```
example# fruadm system-board 'Asset Tag 123456'
```

示例 5 Setting Customer Data

The following command sets the customer data for all FRUs under chassis:

```
example# fruadm -r /frutree/chassis "Property of XYZ, Inc."
```

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	system/fru-id
Interface Stability	Uncommitted

另请参见

[prtfru\(1M\)](#), [attributes\(5\)](#)

引用名	fsck – check and repair file systems
用法概要	<pre>fsck [-F FSType] [-m] [-V] [-v] [special]...</pre> <pre>fsck [-F FSType] [-n N y Y] [-V] [-v] [-o FSType-specific-options] [special]...</pre>
描述	fsck audits and interactively repairs inconsistent file system conditions. If the file system is inconsistent the default action for each correction is to wait for the user to respond yes or no. If the user does not have write permission fsck defaults to a no action. Some corrective actions will result in loss of data. The amount and severity of data loss can be determined from the diagnostic output. <i>FSType-specific-options</i> are options specified in a comma-separated (with no intervening spaces) list of options or keyword-attribute pairs for interpretation by the <i>FSType</i>-specific module of the command. <i>special</i> represents the character special device on which the file system resides, for example, /dev/rdisk/c1t0d0s7. Note: the character special device, not the block special device, should be used. fsck will not work if the block device is mounted. If no <i>special</i> device is specified fsck checks the file systems listed in /etc/vfstab. Those entries in /etc/vfstab which have a character special device entry in the fsckdev field and have a non-zero numeric entry in the fsckpass field will be checked. Specifying -F <i>FSType</i> limits the file systems to be checked to those of the type indicated. If <i>special</i> is specified, but -F is not, the file system type will be determined by looking for a matching entry in /etc/vfstab. If no entry is found, the default local file system type specified in /etc/default/fs will be used. If a file system type supports parallel checking, for example, ufs, some file systems eligible for checking may be checked in parallel. Consult the file system-specific man page (for example, fsck_ufs(1M)) for more information.
选项	The following generic options are supported: -F <i>FSType</i> Specify the file system type on which to operate. -m Check but do not repair. This option checks that the file system is suitable for mounting, returning the appropriate exit status. If the file system is ready for mounting, fsck displays a message such as: ufs fsck: sanity check: /dev/rdisk/c0t3d0s1 okay -n -N Assume a no response to all questions asked by fsck; do not open the file system for writing.

- V
Echo the expanded command line but do not execute the command. This option may be used to verify and to validate the command line.
 - v
Enables verbose output. Might not be supported by all filesystem-specific `fsck` implementations.
 - y | Y
Assume a yes response to all questions asked by `fsck`.
 - o *specific-options*
These *specific-options* can be any combination of the following separated by commas (with no intervening spaces).
 - b=*n*
Use block *n* as the super block for the file system. Block 32 is always one of the alternate super blocks. Determine the location of other super blocks by running `newfs(1M)` with the `-Nv` options specified.
 - c
If the file system is in the old (static table) format, convert it to the new (dynamic table) format. If the file system is in the new format, convert it to the old format provided the old format can support the file system configuration. In interactive mode, `fsck` will list the direction the conversion is to be made and ask whether the conversion should be done. If a negative answer is given, no further operations are done on the file system. In preen mode, the direction of the conversion is listed and done if possible without user interaction. Conversion in preen mode is best used when all the file systems are being converted at once. The format of a file system can be determined from the first line of output from `fstyp(1M)`. Note: the `c` option is seldom used and is included only for compatibility with pre-4.1 releases. There is no guarantee that this option will be included in future releases.
 - f
Force checking of file systems regardless of the state of their super block clean flag.
 - p
Check and fix the file system non-interactively (“preen”). Exit immediately if there is a problem requiring intervention. This option is required to enable parallel file system checking.
 - w
Check writable file systems only.
- 退出状态
- 0
file system is unmounted and OK
 - 1
erroneous parameters are specified

- 32 file system is unmounted and needs checking (`fscck -m` only)
- 33 file system is already mounted
- 34 cannot stat device
- 35 a filesystem that is mounted read/write was modified – reboot
- 36 uncorrectable errors detected - terminate normally
- 37 a signal was caught during processing
- 39 uncorrectable errors detected - terminate immediately
- 40 file system is mounted read-only and is OK

用法 The `fscck` command is *large file aware* for UFS file systems, per the [largefile\(5\)](#) man page.

文件 `/etc/default/fs`
default local file system type. Default values can be set for the following flags in `/etc/default/fs`. For example: `LOCAL=ufs`.

`LOCAL`
The default partition for a command if no `FSType` is specified.

`/etc/vfstab`
list of default parameters for each file system

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [clri\(1M\)](#), [fscck_ufs\(1M\)](#), [fsdb_ufs\(1M\)](#), [fsirand\(1M\)](#), [fstyp\(1M\)](#), [mkfs\(1M\)](#), [mkfs_ufs\(1M\)](#), [mountall\(1M\)](#), [newfs\(1M\)](#), [reboot\(1M\)](#), [vfstab\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#), [ufs\(7FS\)](#)

警告

The operating system buffers file system data. Running `fsck` on a mounted file system can cause the operating system's buffers to become out of date with respect to the disk. For this reason, the file system should be *unmounted* when `fsck` is used. If this is not possible, care should be taken that the system is quiescent and that it is rebooted immediately after `fsck` is run. Quite often, however, this will not be sufficient. A panic will probably occur if running `fsck` on a file system modifies the file system.

附注

This command may not be supported for all *FSTypes*.

Starting with Solaris 9, `fsck` manages extended attribute data on the disk. (See [fsattr\(5\)](#) for a description of extended file attributes.) A file system with extended attributes can be mounted on versions of Solaris that are not attribute-aware (versions prior to Solaris 9), but the attributes will not be accessible and `fsck` will strip them from the files and place them in `lost+found`. Once the attributes have been stripped, the file system is completely stable on versions of Solaris that are not attribute-aware, but would be considered corrupted on attribute-aware versions. In the latter circumstance, run the attribute-aware `fsck` to stabilize the file system before using it in an attribute-aware environment.

引用名

用法概要

描述

选项

fsck_pcfs – file system consistency check and interactive repair

```
fsck -F pcfs [generic_options] special
```

```
fsck -F pcfs [generic_options] [-o specific_options] special
```

The fsck utility audits and interactively repairs inconsistent conditions on file systems. *special* represents the character special device on which the file system resides, for example /dev/disk. The character special device, not the block special device, should be used.

In the case of correcting serious inconsistencies, by default, fsck asks for confirmation before making a repair and waits for the operator to respond either yes or no. If the operator does not have write permission on the file system, fsck defaults to a -n (no corrections) action. See [fsck\(1M\)](#).

Repairing some file system inconsistencies may result in loss of data. The amount and severity of data loss may be determined from the diagnostic output.

When executed with the verify option (-o v), fsck_pcfs automatically scans the entire file system to verify that all of its allocation units are accessible. If it finds any units inaccessible, it updates the file allocation table (FAT) appropriately. It also updates any effected directory entries to reflect the problem. This directory update includes truncating the file at the point in its allocation chain where the file data is no longer accessible. Any remaining accessible allocation units become orphaned.

Orphaned chains of accessible allocation units are, with the operator's concurrence, linked back into the file system as files in the root directory. These files are assigned names of the form fileNNNN.chk, where the Ns are digits in the integral range from 0 through 9.

After successfully scanning and correcting any errors in the file system, fsck displays a summary of information about the file system. This summary includes the size of the file system in bytes, the number of bytes used in directories and individual files, and the number of available allocation units remaining in the file system.

generic_options

The following generic options are supported:

-m

Check but do not repair. This option checks that the file system is suitable for mounting, returning the appropriate exit status. If the file system is ready for mounting, fsck displays a message such as:

```
pcfs fsck: sanity check: /dev/disk okay
```

-n | -N

Assume a no response to all questions asked by fsck; do not open the file system for writing.

-V

Echo the expanded command line, but do not execute the command. This option may be used to verify and to validate the command line.

- y | -Y

Assume a yes response to all questions asked by fsck.
- o *specific_options*

Specify pcfs file system specific options in a comma-separated list, in any combination, with no intervening spaces.
- v

Verify all allocation units are accessible prior to correcting inconsistencies in the metadata.
- p

Check and fix the file system non-interactively (preen). Exit immediately if there is a problem requiring intervention.
- w

Check writable file systems only.

文件

special

The device which contains the pcfs. A hard disk device or high-capacity removable device name much be qualified with a suffix to indicate the proper FDISK partition.

For example, in the names: /dev/rdisk/c0t0d0p0:c and /dev/rdisk/c0t4d0s2:c, the :c suffix indicates the first partition on the disk contains the pcfs.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/pcfs
Interface Stability	Committed

另请参见

[fsck\(1M\)](#), [fstyp\(1M\)](#), [fdisk\(1M\)](#), [mkfs\(1M\)](#), [mkfs_pcfs\(1M\)](#), [mountall\(1M\)](#), [attributes\(5\)](#), [pcfs\(7FS\)](#),

警告

The operating system buffers file system data. Running fsck on a mounted file system can cause the operating system's buffers to become out of date with respect to the disk. For this reason, the file system should be unmounted when fsck is used. If this is not possible, care should be taken that the system is quiescent and that it is rebooted immediately after fsck is run. Quite often, however, this is not sufficient. A panic will probably occur if running fsck on a file system modifies the file system.

引用名	fsck_udfs – file system consistency check and interactive repair
用法概要	<pre>fsck -F udfs [generic_options] [special ...]</pre> <pre>fsck -F udfs [generic_options] [-o specific_options] [special ...]</pre>
描述	fsck audits and interactively repairs inconsistent conditions on file systems. A file system to be checked can be specified by giving the name of the block or character special device or by giving the name of its mount point if a matching entry exists in <code>/etc/vfstab</code>. <i>special</i> represents the character special device, for example, <code>/dev/rdisk/c0t2d0s0</code>, on which the file system resides. The character special device, not the block special device should be used. fsck does not work on a mounted block device. If no special device is specified, all udfs file systems specified in the <code>vfstab</code> file with a <code>fsckdev</code> entry are checked. If the <code>-p</code> (preen) option is specified, udfs file systems with an <code>fsckpass</code> number greater than 1 are checked in parallel. See fsck(1M). In the case of correcting serious inconsistencies, by default, fsck asks for confirmation before making a repair and waits for the operator to respond with either yes or no. If the operator does not have write permission on the file system, fsck defaults to the <code>-n</code> (no corrections) option. See fsck(1M). Repairing some file system inconsistencies can result in loss of data. The amount and severity of data loss can be determined from the diagnostic output. fsck automatically corrects innocuous inconsistencies. It displays a message for each corrected inconsistency that identifies the nature of the correction which took place on the file system. After successfully correcting a file system, fsck prints the number of files on that file system and the number of used and free blocks. Inconsistencies checked are as follows: ■ Blocks claimed by more than one file or the free list ■ Blocks claimed by a file or the free list outside the range of the file system ■ Incorrect link counts in file entries ■ Incorrect directory sizes ■ Bad file entry format ■ Blocks not accounted for anywhere ■ Directory checks, file pointing to unallocated file entry and absence of a parent directory entry ■ Descriptor checks, more blocks for files than there are in the file system ■ Bad free block list format ■ Total free block count incorrect

诊断

If an unmount of the file system is not done before the system is shut down, the file system might become corrupted. In this case, a file system check needs to be completed before the next mount operation.

not writable

You cannot write to the device.

Currently Mounted on

The device is already mounted and cannot run `fsck`.

FILE SYSTEM WAS MODIFIED

File system has been modified to bring it to a consistent state.

Can't read allocation extent

Cannot read the block containing allocation extent.

Bad tag on alloc extent

Invalid tag detected when expecting an allocation extent.

Volume sequence tag error

Invalid tag detected in the volume sequence.

Space bitmap tag error

Invalid tag detected in the space bitmap.

UNEXPECTED INCONSISTENCY; RUN `fsck` MANUALLY

Use `fsck` in interactive mode.

引用名 fsck_ufs – file system consistency check and interactive repair

用法概要 fsck -F ufs [*generic-options*] [*special*] ...

fsck -F ufs [*generic-options*] [-o *specific-options*]
[*special*] ...

描述 The fsck utility audits and interactively repairs inconsistent conditions on file systems. A file system to be checked may be specified by giving the name of the block or character *special* device or by giving the name of its mount point if a matching entry exists in `/etc/vfstab`.

The *special* parameter represents the character special device, for example, `/dev/rdsk/c1t0d0s7`, on which the file system resides. The character special device, not the block special device should be used. The fsck utility will not work if the block device is mounted, unless the file system is error-locked.

If no *special* device is specified, all ufs file systems specified in the `vfstab` with a `fsckdev` entry will be checked. If the `-p` ("preen") option is specified, ufs file systems with an `fsckpass` number greater than 1 are checked in parallel. See [fsck\(1M\)](#).

In the case of correcting serious inconsistencies, by default, fsck asks for confirmation before making a repair and waits for the operator to respond either yes or no. If the operator does not have write permission on the file system, fsck will default to a `-n` (no corrections) action. See [fsck\(1M\)](#).

Repairing some file system inconsistencies can result in loss of data. The amount and severity of data loss can be determined from the diagnostic output.

The fsck utility automatically corrects innocuous inconsistencies such as unreferenced inodes, too-large link counts in inodes, missing blocks in the free list, blocks appearing in the free list and also in files, or incorrect counts in the super block. It displays a message for each inconsistency corrected that identifies the nature of the correction on the file system which took place. After successfully correcting a file system, fsck prints the number of files on that file system, the number of used and free blocks, and the percentage of fragmentation.

Inconsistencies checked include:

- Blocks claimed by more than one inode or the free list.
- Blocks claimed by an inode or the free list outside the range of the file system.
- Incorrect link counts.
- Incorrect directory sizes.
- Bad inode format.
- Blocks not accounted for anywhere.
- Directory checks, file pointing to unallocated inode, inode number out of range, and absence of `‘.’` and `‘. .’` as the first two entries in each directory.
- Super Block checks: more blocks for inodes than there are in the file system.

- Bad free block list format.
- Total free block and/or free inode count incorrect.

Orphaned files and directories (allocated but unreferenced) are, with the operator's concurrence, reconnected by placing them in the `lost+found` directory. The name assigned is the inode number. If the `lost+found` directory does not exist, it is created. If there is insufficient space in the `lost+found` directory, its size is increased.

An attempt to mount a `ufs` file system with the `-o nolargefiles` option will fail if the file system has ever contained a large file (a file whose size is greater than or equal to 2 Gbyte). Invoking `fsck` resets the file system state if no large files are present in the file system. A successful mount of the file system after invoking `fsck` indicates the absence of large files in the file system. An unsuccessful mount attempt indicates the presence of at least one large file. See [mount_ufs\(1M\)](#).

选项

The *generic-options* consist of the following options:

- m Check but do not repair. This option checks that the file system is suitable for mounting, returning the appropriate exit status. If the file system is ready for mounting, `fsck` displays a message such as:

ufs fsck: sanity check: /dev/rdisk/c0t3d0s1 okay
- n | N Assume a no response to all questions asked by `fsck`; do not open the file system for writing.
- V Echo the expanded command line, but do not execute the command. This option may be used to verify and to validate the command line.
- v Enables verbose output. Might not be supported by all filesystem-specific `fsck` implementations.
- y | Y Assume a yes response to all questions asked by `fsck`.

See generic [fsck\(1M\)](#) for the details for specifying *special*.

- o *specific-options* Specify `ufs` file system specific options. These options can be any combination of the following separated by commas (with no intervening spaces).
 - b=*n* Use block *n* as the super block for the file system. Block 32 is always one of the alternate super blocks. Determine the location of other super blocks by running [newfs\(1M\)](#) with the `-Nv` options specified.
 - f Force checking of file systems regardless of the state of their super block clean flag.

- p Check and fix the file system non-interactively (“preen”). Exit immediately if there is a problem requiring intervention. This option is required to enable parallel file system checking.
- w Check writable file systems only.

文件 /etc/vfstab list of default parameters for each file system

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [clri\(1M\)](#), [fsck\(1M\)](#), [fsdb_ufs\(1M\)](#), [fsirand\(1M\)](#), [fstyp\(1M\)](#), [mkfs\(1M\)](#), [mkfs_ufs\(1M\)](#), [mount_ufs\(1M\)](#), [mountall\(1M\)](#), [newfs\(1M\)](#), [reboot\(1M\)](#), [vfstab\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#), [ufs\(7FS\)](#)

警告 The operating system buffers file system data. Running `fsck` on a mounted file system can cause the operating system's buffers to become out of date with respect to the disk. For this reason, the file system should be *unmounted* when `fsck` is used. If this is not possible, care should be taken that the system is quiescent and that it is rebooted immediately after `fsck` is run. Quite often, however, this will not be sufficient. A panic will probably occur if running `fsck` on a file system modifies the file system.

附注 It is usually faster to check the character special device than the block special device.

Running `fsck` on file systems larger than 2 Gb fails if the user chooses to use the block interface to the device:

```
fsck /dev/dsk/c?t?d?s?
```

rather than the raw (character special) device:

```
fsck /dev/rdisk/c?t?d?s?
```

引用名	fsdb – file system debugger				
用法概要	fsdb [-F <i>FSType</i>] [-V] [-o <i>FSType-specific_options</i>] <i>special</i>				
描述	fsdb is a file system debugger that allows for the manual repair of a file system after a crash. <i>special</i> is a special device used to indicate the file system to be debugged. fsdb is intended for experienced users only. <i>FSType</i> is the file system type to be debugged. Since different <i>FSTypes</i> have different structures and hence different debugging capabilities, the manual pages for the <i>FSType</i>-specific fsdb should be consulted for a more detailed description of the debugging capabilities.				
选项	-F Specify the <i>FSType</i> on which to operate. The <i>FSType</i> should either be specified here or be determinable from /etc/vfstab by matching the <i>special</i> with an entry in the table, or by consulting /etc/default/fs. -V Echo the complete command line, but do not execute the command. The command line is generated by using the options and arguments provided by the user and adding to them information derived from /etc/vfstab. This option may be used to verify and validate the command line. -o Specify <i>FSType</i>-specific options.				
用法	See largefile(5) for the description of the behavior of fsdb when encountering files greater than or equal to 2 Gbyte (2 ³¹ bytes).				
文件	/etc/default/fs default local file system type. Default values can be set for the following flags in /etc/default/fs. For example: LOCAL=ufs LOCAL: The default partition for a command if no <i>FSType</i> is specified. /etc/vfstab list of default parameters for each file system				
属性	See attributes(5) for descriptions of the following attributes:				
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>system/core-os</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	system/core-os
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	system/core-os				
另请参见	vfstab(4) , attributes(5) , largefile(5) Manual pages for the <i>FSType</i> -specific modules of fsdb.				
附注	This command may not be supported for all <i>FSTypes</i> .				

引用名

fsdb_udfs – udfs file system debugger

用法概要

fsdb [-F] udfs [*generic_option*] [-o *specific_option*] *special*

描述

The `fsdb_udfs` command is an interactive tool that can be used to patch up a damaged `udfs` file system. `fsdb_udfs` has conversions to translate block and i-numbers into their corresponding disk addresses. Mnemonic offsets to access different parts of an inode are also included. Mnemonic offsets greatly simplify the process of correcting control block entries or descending the file system tree.

`fsdb` contains several error-checking routines to verify inode and block addresses. These can be disabled if necessary by invoking `fsdb` with the `-o` option or by using the `o` command.

`fsdb` reads one block at a time, and therefore works with raw as well as block I/O devices. A buffer management routine is used to retain commonly used blocks of data in order to reduce the number of read system calls. All assignment operations result in an immediate write-through of the corresponding block. In order to modify any portion of the disk, `fsdb` must be invoked with the `-w` option.

Wherever possible, `adb`-like syntax has been adopted to promote the use of `fsdb` through familiarity.

选项

The following options are supported:

<code>-o <i>specific_option</i></code>	Specify <code>udfs</code> file system specific options in a comma-separated list with no intervening spaces. The following specific options are supported:
<code>o</code>	Override some error conditions.
<code>p=<i>string</i></code>	Set prompt to <i>string</i> .
<code>w</code>	Open for write.
<code>?</code>	Display usage.

用法

Numbers are considered hexadecimal by default. The user has control over how data is to be displayed or accepted. The base command displays or sets the input and output base. Once set, all input defaults to this base and all output displays in this base. The base can be overridden temporarily for input by preceding hexadecimal numbers by `0x`, preceding decimal numbers with a `0t`, or octal numbers with a `0`. Hexadecimal numbers beginning with `a-f` or `A-F` must be preceded with a `0x` to distinguish them from commands.

Disk addressing by `fsdb` is at the byte level. However, `fsdb` offers many commands to convert a desired inode, directory entry, block, and so forth, to a byte address. After the address has been calculated, `fsdb` records the result in the current address (`dot`).

Several global values are maintained by `fsdb`:

- Current base (referred to as `base`)
- Current address (referred to as `dot`)

- Current inode (referred to as inode)
- Current count (referred to as count)
- Current type (referred to as type)

Most commands use the preset value of dot in their execution. For example,

```
> 2:inode
```

first sets the value of dot (.) to 2, colon (:), signifies the start of a command, and the inode command sets inode to 2. A count is specified after a comma (,). Once set, count remains at this value until a new command is encountered that resets the value back to 1 (the default).

So, if

```
> 2000,400/X
```

is entered, 400 hex longs are listed from 2000, and when completed, the value of dot is 2000 + 400 * sizeof (long). If a RETURN is then entered, the output routine uses the current values of dot, count, and type and displays 400 more hex longs. An asterisk (*) causes the entire block to be displayed. An example showing several commands and the use of RETURN would be:

```
> 2:ino; 0:dir?d
```

or

```
> 2:ino; 0:db:block?d
```

The two examples are synonymous for getting to the first directory entry of the root of the file system. Once there, subsequently entering a RETURN, plus (+), or minus (-) advances to subsequent entries. Notice that

```
> 2:inode; :ls
```

or

```
> :ls /
```

is again synonymous.

Expressions

The following symbols are recognized by fsdb:

- RETURN Update the value of dot by the current value of *type* and *display* using the current value of *count*.
- # Update the value of dot by specifying a numeric expression. Specify numeric expressions using addition, subtraction, multiplication, and division operators (+, -, *, and %). Numeric expressions are evaluated from left to right and can use parentheses. After evaluation, the value of dot is updated.

, <i>count</i>	Update the count indicator. The global value of <i>count</i> is updated to <i>count</i> . The value of <i>count</i> remains until a new command is run. A <i>count</i> specifier of * attempts to show a block's worth of information. The default for <i>count</i> is 1.
? <i>f</i>	Display in structured style with format specifier <i>f</i> . See Formatted Output.
/ <i>f</i>	Display in unstructured style with format specifier <i>f</i> . See Formatted Output.
.	Display the value of dot.
+ <i>e</i>	Increment the value of dot by the expression <i>e</i> . The amount actually incremented is dependent on the size of type: <code>dot = dot + e * sizeof (type)</code> The default for <i>e</i> is 1.
- <i>e</i>	Decrement the value of dot by the expression <i>e</i> . See +.
* <i>e</i>	Multiply the value of dot by the expression <i>e</i> . Multiplication and division don't use <i>type</i> . In the above calculation of dot, consider the <code>sizeof (type)</code> to be 1.
% <i>e</i>	Divide the value of dot by the expression <i>e</i> . See *.
< <i>name</i>	Restore an address saved in register <i>name</i> . <i>name</i> must be a single letter or digit.
> <i>name</i>	Save an address in register <i>name</i> . <i>name</i> must be a single letter or digit.
= <i>f</i>	Display indicator. If <i>f</i> is a legitimate format specifier (see Formatted Output), then the value of dot is displayed using format specifier <i>f</i> . Otherwise, assignment is assumed. See = [<i>s</i>] [<i>e</i>].
= [<i>s</i>] [<i>e</i>]	Change the value of dot using an assignment indicator. The address pointed to by dot has its contents changed to the value of the expression <i>e</i> or to the ASCII representation of the quoted (") string <i>s</i> . This can be useful for changing directory names or ASCII file information.
=+ <i>e</i>	Change the value of dot using an incremental assignment. The address pointed to by dot has its contents incremented by expression <i>e</i> .
=- <i>e</i>	Change the value of dot using a decremental assignment. Decrement the contents of the address pointed to by dot by expression <i>e</i> .

Commands

A command must be prefixed by a colon (:). Only enough letters of the command to uniquely distinguish it are needed. Multiple commands can be entered on one line by separating them by a SPACE, TAB, or semicolon (;).

To view a potentially unmounted disk in a reasonable manner, fsdb supports the `cd`, `pwd`, `ls`, and `find` commands. The functionality of each of these commands basically matches that of its UNIX counterpart. See `cd(1)`, `pwd(1)`, `ls(1)`, and `find(1)` for details. The *, ,, ?, and - wildcard characters are also supported.

The following commands are supported:

base[= <i>b</i>]	Display or set the base. All input and output is governed by the current base. Without the = <i>b</i> , displays the current base. Otherwise, sets the current base to <i>b</i> . Base is interpreted using the old value of base, so to ensure correctness use the 0, 0t, or 0x prefix when changing the base. The default for base is hexadecimal.
block	Convert the value of dot to a block address.
cd [<i>dir</i>]	Change the current directory to directory <i>dir</i> . The current values of inode and dot are also updated. If <i>dir</i> is not specified, changes directories to inode 2, root (/).
directory	If the current inode is a directory, converts the value of dot to a directory slot offset in that directory, and dot now points to this entry.
file	Set the value of dot as a relative block count from the beginning of the file. The value of dot is updated to the first byte of this block.
find <i>dir</i> [-name <i>n</i>] [-inum <i>i</i>]	Find files by name or i-number. Recursively searches directory <i>dir</i> and below for file names whose i-number matches <i>i</i> or whose name matches pattern <i>n</i> . Only one of the two options (-name or -inum) can be used at one time. The find -print is not necessary or accepted.
fill= <i>p</i>	Fill an area of disk with pattern <i>p</i> . The area of disk is delimited by dot and count.
inode	Convert the value of dot to an inode address. If successful, the current value of inode is updated as well as the value of dot. As a convenient shorthand, if : inode appears at the beginning of the line, the value of dot is set to the current inode and that inode is displayed in inode format.
ls [-R] [-l] <i>pat1 pat2...</i>	List directories or files. If no file is specified, the current directory is assumed. Either or both of the options can be used (but, if used, must be specified before the filename specifiers). Wild card characters are available and multiple arguments are acceptable. The long listing shows only the i-number and the name; use the inode command with ?i to get more information.
override	Toggle the value of override. Some error conditions might be overridden if override is toggled to on.
prompt " <i>p</i> "	Change the fsdb prompt to <i>p</i> . <i>p</i> must be enclosed in quotes.

pwd	Display the current working directory.
quit	Quit fsdb.
tag	Convert the value of dot and if this is a valid tag, print the volume structure according to the tag.
!	Escape to the shell.

Inode Commands

In addition to the above commands, several other commands deal with inode fields and operate directly on the current inode (they still require the colon (:)). They can be used to more easily display or change the particular fields. The value of dot is only used by the :db and :ib commands. Upon completion of the command, the value of dot is changed so that it points to that particular field. For example,

```
> :ln+=1
```

increments the link count of the current inode and sets the value of dot to the address of the link count field.

The following inode commands are supported:

at	Access time
bs	Block size
ct	Creation time
gid	Group id
ln	Link number
mt	Modification time
md	Mode
maj	Major device number
min	Minor device number
nm	This command actually operates on the directory name field. Once poised at the desired directory entry (using the <code>directory</code> command), this command allows you to change or display the directory name. For example, > 7:dir:nm="foo" gets the 7th directory entry of the current inode and changes its name to foo. Directory names cannot be made larger than the field allows. If an attempt is made to make a directory name larger than the field allows,, the string is truncated to fit and a warning message is displayed.
sz	File size

uid User ID

uniq Unique ID

Formatted Output

Formatted output comes in two styles and many format types. The two styles of formatted output are: structured and unstructured. Structured output is used to display inodes, directories, and so forth. Unstructured output displays raw data.

Format specifiers are preceded by the slash (/) or question mark (?) character. *type* is updated as necessary upon completion.

The following format specifiers are preceded by the ? character:

i Display as inodes in the current base.

d Display as directories in the current base.

The following format specifiers are preceded by the / character:

b Display as bytes in the current base.

c Display as characters.

o | O Display as octal shorts or longs.

d | D Display as decimal shorts or longs.

x | X Display as hexadecimal shorts or longs.

示例

示例 1 Using fsdb as a calculator for complex arithmetic

The following command displays 2010 in decimal format, and is an example of using fsdb as a calculator for complex arithmetic.

```
> 2000+400%(20+20)=D
```

示例 2 Using fsdb to display an i-number in inode format

The following command displays the i-number 386 in inode format. 386 becomes the current inode.

```
> 386:ino?i
```

示例 3 Using fsdb to change the link count

The following command changes the link count for the current inode to 4.

```
> :ln=4
```

示例 4 Using fsdb to increment the link count

The following command increments the link count by 1.

```
> :ln+=1
```

示例 5 Using fsdb to display the creation time as a hexadecimal long

The following command displays the creation time as a hexadecimal long.

```
> :ct=X
```

示例 6 Using fsdb to display the modification time in time format

The following command displays the modification time in time format.

```
> :mt=t
```

示例 7 Using fsdb to display in ASCII

The following command displays, in ASCII, block 0 of the file associated with the current inode.

```
> 0:file/c
```

示例 8 Using fsdb to display the directory entries for the root inode

The following command displays the first block's directory entries for the root inode of this file system. This command stops prematurely if the EOF is reached.

```
> 2:ino,*?d
```

示例 9 Using fsdb to change the current inode

The following command changes the current inode to that associated with the 5th directory entry (numbered from 0) of the current inode. The first logical block of the file is then displayed in ASCII.

```
> 5:dir:inode; 0:file,*/c
```

示例 10 Using fsdb to change the i-number

The following command changes the i-number for the 7th directory slot in the root directory to 3.

```
> 2:inode; 7:dir=3
```

示例 11 Using fsdb to change the name field

The following command changes the *name* field in the directory slot to name.

```
> 7:dir:nm="name"
```

示例 12 Using fsdb to display the a block

The following command displays the 3rd block of the current inode as directory entries.

示例 13 Using fsdb to set the contents of address

The following command sets the contents of address 2050 to 0xffffffff. 0xffffffff can be truncated, depending on the current type.

> 2050=0xffff

示例 14 Using fsdb to place an ASCII string at an address

The following command places the ASCII string this is some text at address 1c92434.

> 1c92434="this is some text"

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/udfs

另请参见

[clri\(1M\)](#), [fsck_udfs\(1M\)](#), [dir\(4\)](#), [attributes\(5\)](#)

引用名	fsdb_ufs – ufs file system debugger
用法概要	<code>fsdb -F ufs [generic_options] [specific_options] special</code>
描述	The <code>fsdb_ufs</code> command is an interactive tool that can be used to patch up a damaged UFS file system. It has conversions to translate block and i-numbers into their corresponding disk addresses. Also included are mnemonic offsets to access different parts of an inode. These greatly simplify the process of correcting control block entries or descending the file system tree. <code>fsdb</code> contains several error-checking routines to verify inode and block addresses. These can be disabled if necessary by invoking <code>fsdb</code> with the <code>-o</code> option or by the use of the <code>o</code> command. <code>fsdb</code> reads a block at a time and will therefore work with raw as well as block I/O devices. A buffer management routine is used to retain commonly used blocks of data in order to reduce the number of read system calls. All assignment operations result in an immediate write-through of the corresponding block. Note that in order to modify any portion of the disk, <code>fsdb</code> must be invoked with the <code>w</code> option. Wherever possible, <code>adb</code>-like syntax was adopted to promote the use of <code>fsdb</code> through familiarity.
选项	The following option is supported: -o Specify UFS file system specific options. These options can be any combination of the following separated by commas (with no intervening spaces). The options available are: ? Display usage - o Override some error conditions - p='string' set prompt to string - w open for write
用法	Numbers are considered hexadecimal by default. However, the user has control over how data is to be displayed or accepted. The base command will display or set the input/output base. Once set, all input will default to this base and all output will be shown in this base. The base can be overridden temporarily for input by preceding hexadecimal numbers with <code>'0x'</code>, preceding decimal numbers with <code>'0t'</code>, or octal numbers with <code>'0'</code>. Hexadecimal numbers beginning with <code>a-f</code> or <code>A-F</code> must be preceded with <code>'0x'</code> to distinguish them from commands. Disk addressing by <code>fsdb</code> is at the byte level. However, <code>fsdb</code> offers many commands to convert a desired inode, directory entry, block, superblock and so forth to a byte address. Once the address has been calculated, <code>fsdb</code> will record the result in dot (<code>.</code>). Several global values are maintained by <code>fsdb</code>: ■ the current base (referred to as <code>base</code>),

- the current address (referred to as dot),
- the current inode (referred to as inode),
- the current count (referred to as count),
- and the current type (referred to as type).

Most commands use the preset value of dot in their execution. For example,

```
> 2:inode
```

will first set the value of dot to 2, '.', will alert the start of a command, and the inode command will set inode to 2. A count is specified after a '.'. Once set, count will remain at this value until a new command is encountered which will then reset the value back to 1 (the default). So, if

```
> 2000,400/X
```

is typed, 400 hex longs are listed from 2000, and when completed, the value of dot will be $2000 + 400 * \text{sizeof}(\text{long})$. If a RETURN is then typed, the output routine will use the current values of dot, count, and type and display 400 more hex longs. A '*' will cause the entire block to be displayed.

End of fragment, block and file are maintained by fsdb. When displaying data as fragments or blocks, an error message will be displayed when the end of fragment or block is reached. When displaying data using the db, ib, directory, or file commands an error message is displayed if the end of file is reached. This is mainly needed to avoid passing the end of a directory or file and getting unknown and unwanted results.

An example showing several commands and the use of RETURN would be:

```
> 2:ino; 0:dir?d
      or
> 2:ino; 0:db:block?d
```

The two examples are synonymous for getting to the first directory entry of the root of the file system. Once there, any subsequent RETURN (or +, -) will advance to subsequent entries.

Note that

```
> 2:inode; :ls
      or
> :ls /
```

is again synonymous.

Expressions

The symbols recognized by fsdb are:

RETURN update the value of dot by the current value of type and display using the current value of count.

#	numeric expressions may be composed of +, -, *, and % operators (evaluated left to right) and may use parentheses. Once evaluated, the value of dot is updated.
, <i>count</i>	count indicator. The global value of count will be updated to count. The value of count will remain until a new command is run. A count specifier of '*' will attempt to show a <i>blocks's</i> worth of information. The default for count is 1.
? <i>f</i>	display in structured style with format specifier <i>f</i> . See FormattedOutput.
/ <i>f</i>	display in unstructured style with format specifier <i>f</i> . See FormattedOutput.
.	the value of dot.
+ <i>e</i>	increment the value of dot by the expression <i>e</i> . The amount actually incremented is dependent on the size of type: $\text{dot} = \text{dot} + e * \text{sizeof}(\text{type})$ The default for <i>e</i> is 1.
- <i>e</i>	decrement the value of dot by the expression <i>e</i> . See +.
* <i>e</i>	multiply the value of dot by the expression <i>e</i> . Multiplication and division don't use type. In the above calculation of dot, consider the <code>sizeof (type)</code> to be 1.
% <i>e</i>	divide the value of dot by the expression <i>e</i> . See *.
< <i>name</i>	restore an address saved in register <i>name</i> . <i>name</i> must be a single letter or digit.
> <i>name</i>	save an address in register <i>name</i> . <i>name</i> must be a single letter or digit.
= <i>f</i>	display indicator. If <i>f</i> is a legitimate format specifier, then the value of dot is displayed using the format specifier <i>f</i> . See FormattedOutput. Otherwise, assignment is assumed See =.
= [<i>s</i>] [<i>e</i>]	assignment indicator. The address pointed to by dot has its contents changed to the value of the expression <i>e</i> or to the ASCII representation of the quoted (") string <i>s</i> . This may be useful for changing directory names or ASCII file information.
=+ <i>e</i>	incremental assignment. The address pointed to by dot has its contents incremented by expression <i>e</i> .
=- <i>e</i>	decremental assignment. The address pointed to by dot has its contents decremented by expression <i>e</i> .

Commands

A command must be prefixed by a ':' character. Only enough letters of the command to uniquely distinguish it are needed. Multiple commands may be entered on one line by separating them by a SPACE, TAB or ';'.

In order to view a potentially unmounted disk in a reasonable manner, `fsdb` offers the `cd`, `pwd`, `ls` and `find` commands. The functionality of these commands substantially matches those of its UNIX counterparts. See individual commands for details. The '*', '?', and '[' wild card characters are available.

<code>base=b</code>	display or set base. As stated above, all input and output is governed by the current base. If the <code>=b</code> is omitted, the current base is displayed. Otherwise, the current base is set to <i>b</i> . Note that this is interpreted using the old value of <code>base</code> , so to ensure correctness use the '0', '0t', or '0x' prefix when changing the base. The default for <code>base</code> is hexadecimal.
<code>block</code>	convert the value of <code>dot</code> to a block address.
<code>cd dir</code>	change the current directory to directory <i>dir</i> . The current values of <code>inode</code> and <code>dot</code> are also updated. If no <i>dir</i> is specified, then change directories to inode 2 ("/").
<code>cg</code>	convert the value of <code>dot</code> to a cylinder group.
<code>directory</code>	If the current <code>inode</code> is a directory, then the value of <code>dot</code> is converted to a directory slot offset in that directory and <code>dot</code> now points to this entry.
<code>file</code>	the value of <code>dot</code> is taken as a relative block count from the beginning of the file. The value of <code>dot</code> is updated to the first byte of this block.
<code>find dir [-name n] [-inum i]</code>	find files by name or i-number. <code>find</code> recursively searches directory <i>dir</i> and below for filenames whose i-number matches <i>i</i> or whose name matches pattern <i>n</i> . Note that only one of the two options (<code>-name</code> or <code>-inum</code>) may be used at one time. Also, the <code>-print</code> is not needed or accepted.
<code>fill=p</code>	fill an area of disk with pattern <i>p</i> . The area of disk is delimited by <code>dot</code> and <code>count</code> .
<code>fragment</code>	convert the value of <i>dot</i> to a fragment address. The only difference between the <code>fragment</code> command and the <code>block</code> command is the amount that is able to be displayed.
<code>inode</code>	convert the value of <i>dot</i> to an inode address. If successful, the current value of <code>inode</code> will be updated as well as the value of <i>dot</i> . As a convenient shorthand, if <code>'inode'</code> appears at the beginning of the line, the value of <i>dot</i> is set to the current <code>inode</code> and that <code>inode</code> is displayed in inode format.
<code>log_chk</code>	run through the valid log entries without printing any information and verify the layout.

log_delta	count the number of deltas into the log, using the value of dot as an offset into the log. No checking is done to make sure that offset is within the head/tail offsets.
log_head	display the header information about the file system logging. This shows the block allocation for the log and the data structures on the disk.
log_otodb	return the physical disk block number, using the value of dot as an offset into the log.
log_show	display all deltas between the beginning of the log (BOL) and the end of the log (EOL).
ls	[-R] [-l] <i>pat1 pat2</i> . . . list directories or files. If no file is specified, the current directory is assumed. Either or both of the options may be used (but, if used, <i>must</i> be specified before the filename specifiers). Also, as stated above, wild card characters are available and multiple arguments may be given. The long listing shows only the i-number and the name; use the inode command with 'i' to get more information.
override	toggle the value of override. Some error conditions may be overridden if override is toggled on.
prompt <i>p</i>	change the fsdb prompt to <i>p</i> . <i>p</i> must be surrounded by (")s.
pwd	display the current working directory.
quit	quit fsdb.
sb	the value of <i>dot</i> is taken as a cylinder group number and then converted to the address of the superblock in that cylinder group. As a shorthand, 'sb' at the beginning of a line will set the value of <i>dot</i> to <i>the</i> superblock and display it in superblock format.
shadow	if the current inode is a shadow inode, then the value of <i>dot</i> is set to the beginning of the shadow inode data.
!	escape to shell

Inode Commands In addition to the above commands, there are several commands that deal with inode fields and operate directly on the current inode (they still require the ':'). They may be used to more easily display or change the particular fields. The value of *dot* is only used by the ':db' and ':ib' commands. Upon completion of the command, the value of *dot* is changed to point to that particular field. For example,

```
> :ln+=1
```

would increment the link count of the current inode and set the value of *dot* to the address of the link count field.

at access time.

bs block size.

ct creation time.

db use the current value of *dot* as a direct block index, where direct blocks number from 0 - 11. In order to display the block itself, you need to 'pipe' this result into the `block` or `fragment` command. For example,

```
> 1:db:block,20/X
```

would get the contents of data block field 1 from the inode and convert it to a block address. 20 longs are then displayed in hexadecimal. See `FormattedOutput`.

gid group id.

ib use the current value of *dot* as an indirect block index where indirect blocks number from 0 - 2. This will only get the indirect block itself (the block containing the pointers to the actual blocks). Use the `file` command and start at block 12 to get to the actual blocks.

ln link count.

mt modification time.

md mode.

maj major device number.

min minor device number.

nm although listed here, this command actually operates on the directory name field. Once poised at the desired directory entry (using the *directory* command), this command will allow you to change or display the directory name. For example,

```
> 7:dir:nm="foo"
```

will get the 7th directory entry of the current inode and change its name to foo. Note that names cannot be made larger than the field is set up for. If an attempt is made, the string is truncated to fit and a warning message to this effect is displayed.

si shadow inode.

sz file size.

uid user id.

Formatted Output

There are two styles and many format types. The two styles are structured and unstructured. Structured output is used to display inodes, directories, superblocks and the like. Unstructured displays raw data. The following shows the different ways of displaying:

?

- c display as cylinder groups
- i display as inodes
- d display as directories
- s display as superblocks
- S display as shadow inode data

/

- b display as bytes
- c display as characters
- o O display as octal shorts or longs
- d D display as decimal shorts or longs
- x X display as hexadecimal shorts or longs

The format specifier immediately follows the '/' or '?' character. The values displayed by '/b' and all '?' formats are displayed in the current base. Also, type is appropriately updated upon completion.

示例

示例 1 Displaying in Decimal

The following command displays 2010 in decimal (use of fsdb as a calculator for complex arithmetic):

```
> 2000+400%(20+20)=D
```

示例 2 Displaying an i-number in Inode Format

The following command displays i-number 386 in an inode format. This now becomes the current inode:

```
> 386:ino?i
```

示例 3 Changing the Link Count

The following command changes the link count for the current inode to 4:

```
> :ln=4
```

示例 4 Incrementing the Link Count

The following command increments the link count by 1:

```
> :ln+=1
```

示例 5 Displaying the Creation Time

The following command displays the creation time as a hexadecimal long:

```
> :ct=X
```

示例 6 Displaying the Modification Time

The following command displays the modification time in time format:

```
> :mt=t
```

示例 7 Displaying in ASCII

The following command displays in ASCII, block zero of the file associated with the current inode:

```
> 0:file/c
```

示例 8 Displaying the First Block's Worth of Directory Entries

The following command displays the first block's worth of directory entries for the root inode of this file system. It will stop prematurely if the EOF is reached:

```
> 2:ino,*?d
```

示例 9 Displaying Changes to the Current Inode

The following command displays changes the current inode to that associated with the 5th directory entry (numbered from zero) of the current inode. The first logical block of the file is then displayed in ASCII:

```
> 5:dir:inode; 0:file,*/c
```

示例 10 Displaying the Superblock

The following command displays the superblock of this file system:

```
> :sb
```

示例 11 Displaying the Cylinder Group

The following command displays cylinder group information and summary for cylinder group 1:

```
> 1:cg?c
```

示例 12 Changing the i-number

The following command changes the i-number for the seventh directory slot in the root directory to 3:

```
> 2:inode; 7:dir=3
```

示例 13 Displaying as Directory Entries

The following command displays the third block of the current inode as directory entries:

```
> 2:db:block,*?d
```

示例 14 Changing the Name Field

The following command changes the name field in the directory slot to *name*:

```
> 7:dir:nm="name"
```

示例 15 Getting and Filling Elements

The following command gets fragment 3c3 and fill 20 type elements with 0x20:

```
> 3c3:fragment,20:fill=0x20
```

示例 16 Setting the Contents of an Address

The following command sets the contents of address 2050 to 0xffffffff. 0xffffffff may be truncated depending on the current type:

```
> 2050=0xfffff
```

示例 17 Placing ASCII

The following command places the ASCII for the string at 1c92434:

```
> 1c92434="this is some text"
```

示例 18 Displaying Shadow Inode Data

The following command displays all of the shadow inode data in the shadow inode associated with the root inode of this file system:

```
> 2:ino:si:ino;0:shadow,*?S
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见	clri(1M) , fsck_ufs(1M) , dir_ufs(4) , attributes(5) , ufs(7FS)
警告	Since fsdb reads the disk raw, extreme caution is advised in determining its availability of fsdb on the system. Suggested permissions are 600 and owned by bin.
附注	The old command line syntax for clearing i-nodes using the ufs-specific ' - z i - number ' option is still supported by the new debugger, though it is obsolete and will be removed in a future release. Use of this flag will result in correct operation, but an error message will be printed warning of the impending obsolesence of this option to the command. The equivalent functionality is available using the more flexible clri(1M) command.

- 引用名

fsirand – install random inode generation numbers
- 用法概要

fsirand [-p] *special*
- 描述

fsirand installs random inode generation numbers on all the inodes on device *special*, and also installs a file system ID in the superblock. This helps increase the security of file systems exported by NFS.

fsirand must be used only on an unmounted file system that has been checked with [fsck\(1M\)](#) The only exception is that it can be used on the root file system in single-user mode, if the system is immediately re-booted afterwards.
- 选项

-p Print out the generation numbers for all the inodes, but do not change the generation numbers.
- 用法

See [largefile\(5\)](#) for the description of the behavior of fsirand when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[fsck\(1M\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

引用名	fssnap – create temporary snapshots of a file system										
用法概要	<pre>fssnap [-F FSType] [-V] -o special_options /mount/point fssnap -d [-F FSType] [-V] /mount/point dev fssnap -i [-F FSType] [-V] [-o special_options] [/mount/point dev]</pre>										
描述	The <code>fssnap</code> command creates a stable, read-only snapshot of a file system when given either an active mount point or a special device containing a mounted file system, as in the first form of the synopsis. A snapshot is a temporary image of a file system intended for backup operations. While the snapshot file system is stable and consistent, an application updating files when the snapshot is created might leave these files in an internally inconsistent, truncated, or otherwise unusable state. In such a case, the snapshot will contain these partially written or corrupted files. It is a good idea to ensure active applications are suspended or checkpointed and their associated files are also consistent during snapshot creation. File access times are not updated while the snapshot is being created. A path to the virtual device that contains this snapshot is printed to standard output when a snapshot is created.										
选项	The following options are supported: <table><tr><td>-d</td><td>Deletes the snapshot associated with the given file system.</td></tr><tr><td>-F <i>FSType</i></td><td>Specifies the file system type to be used. The <i>FSType</i> should either be specified here or be determined by matching the block special device with an entry in the <code>/etc/vfstab</code> table, or by consulting <code>/etc/default/fs</code>.</td></tr><tr><td>-i</td><td>Displays the state of any given <i>FSType</i> snapshot. If a mount-point or device is not given, a list of all snapshots on the system is displayed. When a mount-point or device is specified, detailed information is provided for the specified file system snapshot by default. The format and meaning of this information is file-system dependent. See the <i>FSType</i>-specific <code>fssnap</code> man page for details.</td></tr><tr><td>-o <i>special_options</i></td><td>See the <i>FSType</i>-specific man page for <code>fssnap</code>.</td></tr><tr><td>-V</td><td>Echoes the complete command line, but does not execute the command.</td></tr></table>	-d	Deletes the snapshot associated with the given file system.	-F <i>FSType</i>	Specifies the file system type to be used. The <i>FSType</i> should either be specified here or be determined by matching the block special device with an entry in the <code>/etc/vfstab</code> table, or by consulting <code>/etc/default/fs</code> .	-i	Displays the state of any given <i>FSType</i> snapshot. If a mount-point or device is not given, a list of all snapshots on the system is displayed. When a mount-point or device is specified, detailed information is provided for the specified file system snapshot by default. The format and meaning of this information is file-system dependent. See the <i>FSType</i> -specific <code>fssnap</code> man page for details.	-o <i>special_options</i>	See the <i>FSType</i> -specific man page for <code>fssnap</code> .	-V	Echoes the complete command line, but does not execute the command.
-d	Deletes the snapshot associated with the given file system.										
-F <i>FSType</i>	Specifies the file system type to be used. The <i>FSType</i> should either be specified here or be determined by matching the block special device with an entry in the <code>/etc/vfstab</code> table, or by consulting <code>/etc/default/fs</code> .										
-i	Displays the state of any given <i>FSType</i> snapshot. If a mount-point or device is not given, a list of all snapshots on the system is displayed. When a mount-point or device is specified, detailed information is provided for the specified file system snapshot by default. The format and meaning of this information is file-system dependent. See the <i>FSType</i> -specific <code>fssnap</code> man page for details.										
-o <i>special_options</i>	See the <i>FSType</i> -specific man page for <code>fssnap</code> .										
-V	Echoes the complete command line, but does not execute the command.										
操作数	The following operands are supported: <table><tr><td><i>/mount/point</i></td><td>The directory where the file system resides.</td></tr></table>	<i>/mount/point</i>	The directory where the file system resides.								
<i>/mount/point</i>	The directory where the file system resides.										

示例 See FSType-specific man pages for examples.

退出状态 The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

文件

- /etc/vfstab Specifies file system type.
- /etc/default/fs Specifies the default local file system type.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [fssnap_ufs\(1M\)](#), [attributes\(5\)](#)

附注 This command might not be supported for all FSTypes.

引用名	fssnap_ufs – create a temporary snapshot of a UFS file system
用法概要	<pre>fssnap [-F ufs] [-V] -o <i>backing-store=path</i>, [<i>specific-options</i>] /mount/point fssnap -d [-F ufs] [-V] /mount/point dev fssnap -i [-F ufs] [-V] [-o <i>specific-options</i>] /mount/point dev</pre>
描述	The <code>fssnap</code> command queries, creates, or deletes a temporary snapshot of a UFS file system. A snapshot is a point-in-time image of a file system that provides a stable and unchanging device interface for backups. When creating a file system snapshot, you must specify the file system to be captured and the backing-store file. The backing-store file(s) are where the snapshot subsystem saves old file system data before it is overwritten. Beyond the first backing-store file, <code>fssnap</code> automatically creates additional backing-store files on an as-needed basis. The number and size of the backing store files varies with the amount of activity in the file system. The destination path must have enough free space to hold the backing-store file(s). This location must be different from the file system that is being captured in a snapshot. The backing-store file(s) can reside on any type of file system, including another UFS file system or an NFS-mounted file system.
选项	The following options are supported: <code>-d</code> Deletes the snapshot associated with the given file system. <code>-i</code> Displays the state of one or all UFS snapshots. If a mount-point or device is not specified, a list of all snapshots on the system is displayed. When a mount-point or device is specified, detailed information is provided for the specified file system snapshot by default. Use the <code>-o</code> options with the <code>-i</code> option to specify what snapshot information is displayed. Since this feature is provided primarily for use in scripts and on the command line, no labels are displayed for the data. Sizes are all in bytes, and the output is not internationalized or localized. The information is displayed on one line per option. Unrecognized options display a single <code>?</code> on the line. One line per option guarantees that there are the same number of lines as options specified and there is a one-to-one correspondence between an output line and an option. The following <code>-o</code> options display specific information for a given snapshot. See the EXAMPLES section for examples of how to use these options. <code>snapnumber</code> Display the snapshot number. <code>blockdevname</code> Display the block device path.

`rawdevname`

Display the raw device path.

`mountpoint`

Display the mount point of the master file system.

`state`

Display the state of the snapshot device.

`backing-store`

Display the location of the first backing-store file for this snapshot. If there are multiple backing-store files, subsequent files have the same name as the first file, with the suffixes .2, .3, and so forth.

`backing-store-len`

Display the sum of the sizes of the backing-store files.

`maxsize`

Display the maxsize value specified for the backing-store file(s).

`createtime`

Display the time that the snapshot was created.

`chunksize`

Display the copy-on-write granularity.

`-o specific-options`

Without `-d` or `-i`, the default action is to create a snapshot. Specify the following options when creating a snapshot. All of these options are discretionary, except for the backing-store file, which is required.

`backing-store=path`

Uses *path* in the creation of the backing-store file(s). *path* must not reside on the file system that is being captured in a snapshot and must not be the name of an existing file. If *path* is a directory, then a backing-store file is created within it using a name that is generated automatically. If *path* is not a directory and does not already exist, then a backing-store file with that name is created. If more than one backing-store file is required, `fssnap` creates subsequent files automatically. The second and subsequent files have the same name as the first file, with suffixes of .2, .3, and so forth.

This option can be abbreviated as `bf=path` or `bs=path`.

`unlink`

Unlinks the backing-store file after the snapshot is created. This option specifies that the backing-store file does not need to be removed manually when the snapshot is deleted. This might make administration more difficult since the file is not visible in the file system. If this option is not specified, the backing-store files should be removed manually after the snapshot is deleted.

`chunksize=n [k,m,g]`

Uses *n* for the chunk size. Chunk size is the granularity of the data that is sent to the backing store.

Specify `chunksize` in the following units: k for kilobytes, m for megabytes, or g for gigabytes. By default, chunk size is four times the block size of the file system (typically 32k).

`maxsize=n[k,m,g]`

Does not allow the sum of the sizes of the backing-store file(s) to exceed *n*, where *n* is the unit specified. The snapshot is deleted automatically when the sum of the sizes of the backing-store file(s) exceeds `maxsize`.

Specify `maxsize` in the following units: k for kilobytes, m for megabytes, or g for gigabytes.

`raw`

Displays to standard output the name of the raw device instead of the block device when a snapshot is created. The block device is printed by default (when `raw` is not specified). This option makes it easier to embed `fssnap` commands in the command line for commands that require the raw device instead. Both devices are always created. This option affects only the output.

操作数

The following operands are supported:

mount-point

The directory where the file system resides.

special

The physical device for the file system, such as `/dev/dsk/c0t0d0s7`.

示例

示例 1 Creating a Snapshot of a File System

The following example creates a snapshot of a file system. The block special device created for the snapshot is `/dev/fssnap/0`.

```
# fssnap -F ufs -o backing-store=/var/tmp /export/home
/dev/fssnap/0
```

示例 2 Backing Up a File System Snapshot Without Having To Unmount the File System

The following example backs up a file system snapshot without having to unmount the file system. Since `ufsdump` requires the path to a raw device, the `raw` option is used. The `/export/home` file system snapshot is removed in the second command.

```
# ufsdump 0uf /dev/rmt/0 'fssnap -F ufs
-o raw,bs=/export/snap /export/home'
<output from ufsdump>
# fssnap -F ufs -d /export/home
```

示例 3 Backing Up a File System

When backing up a file system, do not let the backing-store file(s) exceed 400 Mbytes. The second command removes the /export/home file system snapshot.

```
# ufsdump 0uf /dev/rmt/0 'fssnap -F ufs
    -o maxsize=400m,backing-store=/export/snap,raw
    /export/home'
# fssnap -F ufs -d /export/home
```

示例 4 Performing an Incremental Dump of a Snapshot

The following example uses ufsdump to back up a snapshot of /var. Note the use of the N option to ufsdump, which writes the name of the device being dumped, rather than the name of the snapshot device, to /etc/dumpdates file. See [ufsdump\(1M\)](#) for details on the N flag.

```
# ufsdump lfNu /dev/rmt/0 /dev/rdisk/c0t3d0s2 'fssnap -F ufs
-o raw,bs=/export/scratch,unlink /var'
```

示例 5 Finding Out What Snapshots Currently Exist

The following command displays the currently existing snapshots.

```
# fssnap -i
0 /src
1 /export/home
<output continues>
```

示例 6 Mounting a File System Snapshot

The following example creates a file system snapshot. After you create a file system snapshot, mount it on /tmp/mount for temporary read-only access.

```
# fssnap -F ufs -o backing-store=/nfs/server/scratch /export/home
/dev/fssnap/1
# mkdir /tmp/mount
# mount -F ufs -o ro /dev/fssnap/1 /tmp/mount
```

示例 7 Creating a File System Snapshot and Unlinking the Backing-store File

The following example creates a file system snapshot and unlinks the backing-store file. After creating a file system snapshot and unlinking the backing-store file, check the state of the snapshot.

```
# fssnap -o bs=/scratch,unlink /src
/dev/fssnap/0
# fssnap -i /src
Snapshot number          : 0
Block Device             : /dev/fssnap/0
Raw Device               : /dev/rfssnap/0
Mount point              : /src
Device state             : active
```

示例 7 Creating a File System Snapshot and Unlinking the Backing-store File (续)

```
Backing store path      : /scratch/snapshot2 <UNLINKED>
Backing store size      : 192 KB
Maximum backing store size : Unlimited
Snapshot create time    : Sat May 06 10:55:11 2000
Copy-on-write granularity : 32 KB
```

示例 8 Displaying the Size and Location of the Backing-store File(s) and the Creation Time for the Snapshot

The following example displays the size of the backing-store file(s) in bytes, the location of the backing store, and the creation time for the snapshot of the /test file system.

```
# fssnap -i -o backing-store-len,backing-store,createtime /test
196608
/snapshot2
Sat May 6 10:55:11 2000
```

Note that if there are multiple backing-store files stored in /snapshot2, they will have names of the form *file* (for the first file), *file.1*, *file.2*, and so forth.

退出状态

The following exit values are returned:

- 0
Successful completion.
- >0
An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

The script-readable output mode is a stable interface that can be added to, but will not change. All other interfaces are subject to change.

另请参见

[mlock\(3C\)](#), [attributes\(5\)](#)

See the `ntpd` man page, delivered in the `SUNWntpu` package (not a SunOS man page).

附注

The `fssnap` device files should be treated like a regular disk block or character device.

The association between a file system and the snapshot is lost when the snapshot is deleted or the system reboots. Snapshot persistence across reboots is not currently supported.

To avoid unnecessary performance impacts, perform the snapshot and system backup when the system is least active.

It is not possible to perform a snapshot of a file system if any of the following conditions are true:

- The file system is in use by system accounting
- The file system contains a local swap file
- The file system is used as backing store by an application that uses `mlock(3C)` to lock its pages. Typically, these are real time applications, such as `ntpd` (delivered in the `SUNWntpu` package).

These conditions result in `fssnap` being unable to write lock the file system prior to performing the snapshot.

引用名

fsstat – report file system statistics

用法概要

```
fsstat [-a|f|i|n|v] [-T | u|d] {-F | {fstype|path}...}
      [interval [count]]
```

描述

fsstat reports kernel file operation activity by the file system type (*fstype*) or by the path name, which is converted to a mount point. The first set of lines of output reports all activity since:

- The file system module was loaded (in the case of *fstype*)
- The file system was mounted (in the case of mount point)

Statistics are gathered at the file system independent layer at both the *fstype* and the mount point levels. However, not all file system types are represented in the gathering of statistics. (See the NOTES section of this man page.)

The output of fsstat is dependent on the mode (option) requested. All statistic fields are displayed using “smart numbers” which automatically scale the units in a human readable form that fits in a maximum of 5 characters. For example:

```
100           is displayed as 100
2048          is displayed as 2K
3000000       is displayed as 2.86M
```

The unit modifiers are: K (Kbyte), M (Mbyte), G (Gbyte), T (terabyte), P (petabyte), and E (exabyte).

During the execution of fsstat, the state of the system can change. If relevant, a state change message is included in the fsstat output in one of the following forms:

```
<<mount point no longer available: {path}>>
<<file system module no longer loaded: {fstype}>>
```

After the state change messages are displayed, fsstat continues to display the statistics as directed. If all of the *fstypes* and mount points that fsstat was reporting on are no longer available, then fsstat exits.

The user is required to specify the -F option (all available file system types) or a list of one or more *fstypes* and/or mount points.

The default report shows general file system activity. This display combines similar operations into general categories as follows:

```
new file      Number of creation operations for file system objects (for example, files,
              directories, symlinks, etc.)
name remov    Number of name removal operations
name chng     Number of name change operations
```

attr get	Number of object attribute retrieval operations
attr set	Number of object attribute change operations
lookup ops	Number of object lookup operations
rddir ops	Number of read directory operations
read ops	Number of data read operations
read bytes	Bytes transferred by data read operations
write ops	Number of data write operations
write bytes	Bytes transferred by data write operations

The entity being reported on (*fstype* or mount point) is displayed in the last column.

选项

The following options are supported:

- a Report the activity for kernel attribute operations. The following statistics are reported:

getattr	Number of file attribute retrieval calls
setattr	Number of file attribute modification calls
getsec	Number of file security attribute retrieval calls
setsec	Number of file security attribute modification calls

The entity being reported on (*fstype* or mount point) is displayed in the last column.
- f Report the full activity for all kernel file operations. Each file operation is listed in the left column. The following statistics are reported for each operation:

#ops	Number of calls for this operation
bytes	Average transfer size in bytes (only applies to read, write, readdir)

The entity being reported on (*fstype* or mount point) is displayed in the first row.
- i Reports the activity for kernel I/O operations. The following statistics are reported:

read ops	Number of data read calls
read bytes	Number of bytes read
write ops	Number of data write calls
write bytes	Number of bytes written
rddir ops	Number of read directory calls

rddir bytes Number of bytes read by reading directories
 rwlock ops Number of internal file system lock operations
 rwunlock ops Number of internal file system unlock operations

The entity being reported on (*fstype* or mount point) is displayed in the last column.

-n Reports the activity for kernel naming operations. The following statistics are reported:

lookup Number of file name retrieval calls
 creat Number of file creation calls
 remov Number of file remove calls
 link Number of link calls
 renam Number of file renaming calls
 mkdir Number of directory creation calls
 rmdir Number of directory removal calls
 rddir Number of directory read calls
 symlink Number of symlink creation calls
 rdlink Number of symlink read calls

The entity being reported on (*fstype* or mount point) is displayed in the last column.

-v Reports the activity for calls to the virtual memory operations. The following statistics are reported.

map Number of calls mapping a file
 addmap Number of calls setting additional mapping to a mapped file
 delmap Number of calls deleting mapping to a file
 getpag Number of calls retrieving a page of data from a file
 putpag Number of calls writing a page of data to a file
 pagio Number of calls to transfer pages in file system swap files

The entity being reported on (*fstype* or mount point) is displayed in the last column.

-F Report on all available file system types.

-T *u|d* Display a time stamp.

Specify *u* for a printed representation of the internal representation of time (see [time\(2\)](#)) Specify *d* for the standard date format. (See [date\(1\)](#)). The time stamp is only used when an interval is set.

操作数

The following operands are supported:

count Display only *count* reports.

fstype Explicitly specify the file system type(s) to be reported. The file system module must be loaded.

interval Report once each *interval* seconds.

path Specify the path(s) of the mount point(s) to be reported. If path is not a mount point, the mount point containing path will be determined and displayed in the output.

If no *interval* and no *count* are specified, a single report is printed and fsstat exits. If an *interval* is specified but no *count* is specified, fsstat prints reports every *interval* seconds indefinitely until the command is interrupted.

示例

示例1 Displaying General Activity

The following example shows general activity for all file system types.

```
$ fsstat -F
new  name   name  attr  attr  lookup rddir  read  read  write write
file remov chng  get   set   ops    ops   ops  bytes ops  bytes
313K 214K 38.5K 2.16M 56.2K 8.36M 52.8K 19.7M 39.9G 18.8M 39.1G ufs
   0    0    0 2.95K   0  3.81K  282 2.52K 466K   0    0 proc
   0    0    0   0     0   0     0   0     0   0    0 nfs
  10    8    2   86    9   98    15  413  103M 8.43K 1.05G zfs
  13   14    4   98   16  125   10 1.01K 258M 15.9K 127M lofs
8.73K 3.29K 5.25K 55.3K   37 1.20M  44 37.9K 38.3M 47.2K 35.9M tmpfs
   0    0    0 4.93K   0   0     0 1.08K 913K   0    0 mntfs
   3    2    1   503   3  897   13  122 25.8K  128 272K nfs3
  10    8    0  615  10 10.1K   18   61 45.6K  292 2.26M nfs4
```

示例2 Displaying Naming Activity

The following example shows the naming activity for ufs, nfs, nfs3, nfs4, and tmpfs:

```
$ fsstat -n ufs nfs nfs3 nfs4 tmpfs
lookup creat remov  link renam mkdir rmdir rddir symlnk rdlnk
3.57M 3.10K 586    6   24  115  100 30.2K    5 330K ufs
   0    0    0    0   0   0     0   0     0   0 0 nfs
18.3K    3    5    0   0   0     0 1.03K    2  346 nfs3
 535    0    0    0   0   0     0   46    0    4 nfs4
```

示例2 Displaying Naming Activity (续)

```
146      24      15      0      0      4      0      4      0      0 tmpfs
```

示例3 Displaying Attribute Activity

The following example shows the attribute activity for the FS type ufs and the mounted file systems “/” and “/export/home” every three seconds for every third iteration:

```
# fsstat -a ufs / /export/home 3 3
getattr setattr getsec setsec
 378K   91.9K  11.8K      0 ufs
 367K   82.3K  11.6K      0 /
 11.3K    9.6K   198      0 /export/home
 4.97K   2.27K  163      0 ufs
 3.94K   1.36K  162      0 /
 1.03K    927    1      0 /export/home
 2.30K   1.06K   73      0 ufs
 1.95K    766    71      0 /
   361    317     2      0 /export/home
 2.33K   1.06K   78      0 ufs
 1.64K    451    77      0 /
   711    631     1      0 /export/home
```

示例4 Displaying File Operation Statistics

The following example shows the statistics for each file operation for “/” (using the -f option):

```
$ fsstat -f /
Mountpoint: /
operation #ops bytes
  open 8.54K
  close 9.8K
  read 43.6K 65.9M
  write 1.57K 2.99M
  ioctl 2.06K
  setfl 4
  getattr 40.3K
  setattr 38
  access 9.19K
  lookup 203K
  create 595
  remove 56
  link 0
  rename 9
  mkdir 19
  rmdir 0
  readdir 2.02K 2.27M
  symlink 4
```

示例 4 Displaying File Operation Statistics (续)

```
readlink 8.31K
fsync    199
inactive 2.96K
fid      0
rwlock   47.2K
rwunlock 47.2K
seek     29.1K
cmp       42.9K
frlock   4.45K
space     8
realvp   3.25K
getpage  104K
putpage  2.69K
map       13.2K
addmap   34.4K
delmap   33.4K
poll     287
dump      0
pathconf  54
pageio    0
dumpctl   0
dispose  23.8K
getsecattr 697
setsecattr  0
shrlock   0
vnevent   0
```

环境变量 See [environ\(5\)](#) for descriptions of the following environment variables that affect the execution of fsstat: LANG, LC_ALL, LC_CTYPE, LC_MESSAGES, LC_TIME, and NLSPATH.

退出状态 The following exit values are returned:

- 0 Successful completion.
- 1 A fatal error occurred. A fatal error could be a failed system call or another internal error.
- 2 Invalid command-line options were specified.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
CSI	Enabled

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	See below.

The command-line options are Uncommitted. The human-readable output is not considered an interface.

另请参见

[date\(1\)](#), [time\(2\)](#), [attributes\(5\)](#)

附注

All display options (`-a`, `-f`, `-i`, `-n`, `-v`) are mutually exclusive. Entering more than one of these options will result in an error.

The *fstype* and *path* operands must appear after the option, but before the *interval* or *count* on the command line. For example, “`fsstat -a fstype interval`”. Preference is given to *fstype* so that if a user wishes to see the statistics for a directory that has the same name as an *fstype* (for example, *ufs*), then the path must be specified unambiguously (for example, `./ufs`). Similarly, in order to define a file with a numeric name (for example, “10”) from an interval or count operand, the name should be prefixed accordingly (for example, `./10`).

When an interval is used, headers repeat after more than 12 lines of statistics have been displayed and the set of lines to be displayed in the current interval have completed.

Statistics are not displayed for all pseudo-file systems. The output displayed with the `-F` option shows which of the loaded filesystem types are supported.

Unbundled file systems may not be recognized by `fsstat`.

The command-line options are classified as Uncommitted and could change. The output is not considered to be an interface. The construction of higher level software tools depend on either the command-line options or the output of `fsstat` is not recommended.

引用名 fstyp – determine file system type

用法概要 fstyp [-a | -v] *special* [:*logical-drive*]

描述 fstyp allows the user to determine the file system type of unmounted file systems using heuristic programs.

An fstyp module for each file system type to be checked is executed; each of these modules applies an appropriate heuristic to determine whether the supplied *special* file is of the type for which it checks. If it is, the program prints on standard output the usual file system identifier for that type (for example, “ufs”) and exits with a return code of 0; if none of the modules succeed, the error message `unknown_fstyp (no matches)` is returned and the exit status is 1. If more than one module succeeds, the error message `unknown_fstyp (multiple matches)` is returned and the exit status is 2. Other errors are printed in the same format.

This command is unreliable and its results should not be used to make any decisions about subsequent use of a storage device or disk partition.

选项 -a Output all available file system attributes. If a file system has been successfully identified, and this option is specified, the *FSType* identifier is followed by one or more “name-value” pairs, one per line, in the format:

```
name: value
```

The following conventions are recognized for the file system attributes:

- String values are put in single quotes.
- Nested “name-value” list increases the indentation of its values by four whitespaces.
- For an array of “name-value” pairs, one array entry is printed per line, with the index following the name in square brackets.

For instance, in the following example, “*top_list*” is a “name-value” list, consisting of a string array “*string_array*” and a “name-value” list array “*list_array*”. The second “*list_array*” element is an integer array “*int_array*” containing three elements.

```
top_string: 'String'
top_list:
  string_array[0]: 'S0'
  string_array[1]: 'S1'
  list_array[0]:
    int_one: 1
    string_two: 'S2'
  list_array[1]:
    int_array[0]: 1
    int_array[1]: 2
    int_array[2]: 3
```

In addition to the *FSType*-specific attributes, the following generic attributes may be present:

- gen_clean

Allowable values are “true” or “false”. If the value is “false”, the file system is damaged or was not unmounted cleanly and the [fsck\(1M\)](#) command must be run before this file system can be mounted.
- gen_guid

Globally unique identifier. This string uniquely identifies the file system.
- gen_version

A string describing the file system version.
- gen_volume_label

Volume label, a human-readable string used to either describe or identify the file system.
- v

Produce verbose output. This is usually information about the file systems superblock and varies across different *FSTypes*. See [ufs\(7FS\)](#), [mkfs_ufs\(1M\)](#), and [tunefs\(1M\)](#) for details.

用法

See [largefile\(5\)](#) for the description of the behavior of `fstyp` when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[fsck\(1M\)](#), [mkfs_ufs\(1M\)](#), [tunefs\(1M\)](#), [attributes\(5\)](#), [libfstyp\(3LIB\)](#), [largefile\(5\)](#), [hsfs\(7FS\)](#), [ufs\(7FS\)](#), [pcfs\(7FS\)](#)

附注

The use of heuristics implies that the result of `fstyp` is not guaranteed to be accurate.

This command is unreliable and its results should not be used to make any decisions about subsequent use of a storage device or disk partition.

引用名	fuser – 识别文件和设备的用户
用法概要	<pre>/usr/sbin/fuser [-c -d -f] [-nu] [-k -s sig] files [[-] [-c -d -f] [-nu] [-k -s sig] files] ...</pre>
描述	fuser 实用程序显示正在使用指定为参数的 <i>files</i> 的进程的进程 ID。 每个进程 ID 后面跟有一个字母代码。字母代码的解释如下所述。如果进程将文件用作 - c 表示此进程正在使用该文件作为其当前目录。 - m 表示此进程正在使用由 mmap(2) 映射的文件。有关详细信息，请参见 mmap(2)。 - n 表示此进程正持有针对该文件的非阻塞强制性锁定。 - o 表示此进程正在将该文件作为打开的文件使用。 - r 表示此进程正在将该文件用作其根目录。 - t 表示此进程正在将该文件用作其文本文件。 - y 表示此进程正在将该文件用作其控制终端。 对于具有已挂载文件系统的块特殊设备，将列出使用该设备上的任何文件的所有进程。对于所有类型的文件（文本文件、可执行文件、目录、设备，等等），只会报告使用该文件的进程。 对于所有类型的设备，fuser 还会显示打开设备的任何已知内核使用者。内核使用者显示为下列格式之一： <pre>[module_name] [module_name,dev_path=path] [module_name,dev=(major,minor)] [module_name,dev=(major,minor),dev_path=path]</pre> 如果指定了多组文件，可能需要为其他每个文件组重新指定选项。单个短划线可取消当前施行的选项。 各个进程 ID 输出到标准输出中的单个行上，并以空格分隔。所有其他输出（包括单个终止换行符）均被写入到标准错误。 任何用户都可以运行 fuser，但只有超级用户可以终止其他用户的进程。
选项	所支持的选项如下： -c 报告用作文件系统的挂载点的文件，以及该已挂载文件系统内的所有文件。 -d 报告与指定的次要节点绑定到同一设备节点的所有次要节点的设备使用信息。此选项不报告已挂载文件系统内的文件的文件使用情况。 -f 为指定的文件输出报告，而不是已为挂载文件系统内的文件输出报告。

- k 向每个进程发送 SIGKILL 信号。因为此选项要针对每个进程生成终止信号，因此终止消息可能不会立即显示（请参见 [kill\(2\)](#)）。不会向内核文件使用者发送信号。
- n 只列出具有针对文件的非阻塞强制性锁定的进程。
- s sig 向每个进程发送信号。sig 选项参数指定在 <signal.h> 标头中定义的符号名称之一，或者指定一个十进制整数信号编号。如果 sig 是一个符号名称，识别它时将不区分大小写，且没有 SIG 前缀。-k 选项等效于 -s KILL 或 -s 9。不会向内核文件使用者发送信号。
- u 在进程 ID 后面的圆括号中显示用户登录名。

示例

示例1 报告挂载点和文件
以下示例将报告已挂载文件系统内的挂载点和文件。

```
example% fuser -c /export/foo
```

示例2 在报告挂载点和文件时对输出进行限制
以下示例将报告已挂载文件系统内的挂载点和文件，但将输出限制为持有非阻塞强制性锁定的那些进程。

```
example% fuser -cn /export/foo
```

示例3 向持有非阻塞强制性锁定的进程发送 SIGTERM
以下命令向持有针对文件 /export/foo/my_file 的非阻塞强制性锁定的所有进程发送 SIGTERM。

```
example% fuser -fn -s term /export/foo/my_file
```

环境变量

有关影响 fuser 执行的环境变量 LANG、LC_ALL LC_CTYPE、LC_MESSAGES 和 NLSPATH 的描述，请参见 [environ\(5\)](#)。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）
标准	请参见 standards(5) 。

另请参见

[ps\(1\)](#)、[mount\(1M\)](#)、[kill\(2\)](#)、[mmap\(2\)](#)、[signal\(3C\)](#)、[attributes\(5\)](#)、[environ\(5\)](#)、[standards\(5\)](#)

附注

因为 fuser 使用系统映像的快照，所以它可能会漏掉在 fuser 运行期间开始使用某个文件的进程。另外，被报告为正在使用某个文件的进程在 fuser 运行期间可能已停止使用该文件。考虑到这些因素，不鼓励使用 -k 选项。

引用名 fwflash – firmware query and update utility

用法概要

```
/usr/sbin/fwflash [-l [-c device_class | ALL ]]
                  | [-v] | [-h]

fwflash [-f file1,file2,file3,... | -r file]
         [-y] [-d device_path]
```

描述

The `fwflash` command writes a binary image file to supported flashable devices attached to a Solaris host. It also provides the ability to read firmware to a file if supported by the device. Because changing the firmware in a device can have significant impact on the stability of a system, only users with the privilege `ALL` are allowed to execute this command. Users authorized to run `fwflash` can be granted the “Firmware Flash Update” Rights Profile.

The first form of the command, above, provides information about devices. It lists all devices currently available on the system that are supported by `fwflash` for firmware upgrade. You can filter the list operation, to display only specified classes of devices. The second form of the command provides the operations to read or write the firmware images to specific devices.

选项 The following options are supported:

-c *device_class*

An optional parameter, valid only when used with the `-l` option. This option causes the command to list only devices of a specific class type. Currently supported classes are `IB`, `ses`, `sesgen`, or `ALL`. If `-c` is not specified for the `-l` option, the class defaults to `ALL`.

-d *dev_path*

The *dev_path* is absolute path name of the device that the user wants to modify with the `-f` or `-r` operation. If the device cannot be found, the command fails. If the `-d` option is specified, then either `-f` or `-r` must also be specified.

-f *file1,file2,file3,...*

Specify the path to one or more binary firmware image files you want to write to the device. `fwflash` will verify that each file is a valid firmware image for the specified device. If it is not, the command fails with an appropriate error message.

If multiple firmware image files are specified, each image is verified and flashed to the device in the order given on the command line. If any of the specified files cannot be successfully flashed, then an appropriate message is displayed.

After a new firmware image is flashed to a device, a reboot is required to correctly activate the new firmware.

-h

Display the command line usage message for `fwflash`.

-l

List the devices on a system available for firmware upgrade and display information specific to each device or device class.

For InfiniBand (IB) devices, the list operation displays the guids (Globally Unique Identifier) currently set for the HCA, as well as the current firmware revision installed. There are four separate guids on the HCA; two of them can be set with the same value.

For SCSI Enclosure Services (ses or sgen) devices, an identifying target-port worldwide name is displayed, if available.

-r file

Specify the path to a file to create when reading the firmware from the device. The -f and -r options are mutually exclusive.

Not all flashable devices support reading firmware images back from the device. At present, only InfiniBand (IB) devices are supported for this operation. A message will be displayed if the selected device does not support this operation.

-v

Display fwflash version information and exit.

-y

Valid only when a flash read (-r) or write (-f) operation is specified. This option causes fwflash not to prompt for confirmation during operation and operate non-interactively. Note that there is no option that allows you to forcibly flash an incompatible firmware image onto a device.

示例

示例 1 Entering Command Without Arguments

The following command shows fwflash when the command is entered without arguments.

```
example# fwflash
```

```
Usage:
```

```
Usage:
```

```
fwflash [-l [-c device_class | ALL]] | [-v] | [-h]
fwflash [-f file1,file2,file3,... | -r file] [-y] -d device_path
```

```
-l          list flashable devices in this system
-c device_class limit search to a specific class
              eg IB for InfiniBand, ses for SCSI Enclosures
-v          print version number of fwflash utility
-h          print this usage message

-f file1,file2,file3,...
              firmware image file list to flash
-r file      file to dump device firmware to
-y          answer Yes/Y/y to prompts
-d device_path pathname of device to be flashed
```

```
If -d device_path is specified, then one of -f <files>
or -r <file> must also be specified
```

示例 1 Entering Command Without Arguments (续)

If multiple firmware images are required to be flashed they must be listed together, separated by commas. The images will be flashed in the order specified.

示例 2 Listing Devices Available to Flash

The following command lists the devices available to be flashed.

```
example# fwflash -l
List of available devices:
Device[0], /devices/pci@0,0/pci8086,3595@2/pci8086,32a@0,2/\
pci15b3,5a46@c/pci15b3,5a44@0:devctl
Class [IB]
  GUID: System Image - 0002c901081e33b3
  Node      - 0000000000003446
  Port 1    - 0002c901081e33b1
  Port 2    - 0002c901081e33b2
  Firmware revision: 3.3.0002
  No hardware information available for this device

Device[1], /devices/pci@0,0/pci8086,3597@4/pci15b3,6278@0:devctl
Class [IB]
  GUID: System Image - 0002c9010a99e3b3
  Node      - 0002c9010a99e3b0
  Port 1    - 0002c9010a99e3b1
  Port 2    - 0002c9010a99e3b2
  Firmware revision: 4.8.00c8
  Product    : MTL P25208-CF256T (Lion cub)
  PSID      : MT_00B0000001
```

Alternatively, for a SAS Expander presented as a SCSI Enclosure Services device, we might see output such as this:

```
example# fwflash -l
List of available devices:
Device[0] /devices/pci@0/pci@0/pci@2/scsi@0/ses@3,0:ses
Class [sgen]
  Target port WWN : 500605b00002453d
  Vendor          : SUN
  Product         : 16Disk Backplane
  Firmware revision: 5021
```

示例 3 Flash Upgrading an IB HCA Device

The following command flash upgrades an IB HCA device.

```
example# fwflash -f ./version.3.2.0000 \
-d /devices/pci@0,0/pci8086,3597@4/pci15b3,6278@0:devctl
About to update firmware on:
```

示例 3 Flash Upgrading an IB HCA Device (续)

```
/devices/pci@0,0/pci8086,3597@4/pci15b3,6278@0:devctl
Continue (Y/N): Y
```

```
Updating . . . . .
Done. New image will be active after the system is rebooted.
```

Note that you are prompted before the upgrading proceeds and that it is mandatory that you reboot your host to activate the new firmware image.

The following command adds the `-y` option to the command.

```
example# fwflash -y -f ./version.3.2.0000 \
-d /devices/pci@0,0/pci8086,3597@4/pci15b3,6278@0:devctl
About to update firmware on:
/devices/pci@0,0/pci8086,3597@4/pci15b3,6278@0:devctl
```

```
Updating . . . . .
Done. New image will be active after the system is rebooted.
```

示例 4 Reading Device Firmware to File

The command shown below reads the device firmware to a file. The command uses the `-y` option so that read occurs without prompting.

```
example# fwflash -y -r /firmware.bin \
-d /devices/pci@1d,700000/pci@1/pci15b3,5a44@0:devctl
About to read firmware on:
/devices/pci@1d,700000/pci@1/pci15b3,5a44@0:devctl
to filename: /firmware.bin
```

```
Reading . . .
Done.
```

示例 5 When No Flashable Devices Are Found

The command output shown below informs the user that there are no supported flashable devices found in the system:

```
example# fwflash -l
fwflash: No flashable devices attached with the ses driver in this system
fwflash: No flashable devices attached with the sgen driver in this system
fwflash: No flashable devices attached with the tavor driver in this system
fwflash: No flashable devices in this system
```

Each plugin found in `/usr/lib/fwflash/identify` is loaded in turn, and walks the system device tree, determining whether any currently-attached devices can be flashed. For the list of device types and drivers that are currently supported, please see the NOTES section below.

返回值

The `fwflash` command returns the following values:

0

Success

1

Failure

属性

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/firmware-utilities
Interface Stability	Committed

另请参见

attributes(5), hermon(7D), ses(7D), tavor(7D)

The InfiniBand Trade Association website, <http://www.infinibandta.org>

The SCSI Storage Interfaces committee website, <http://www.t10.org>

SCSI Primary Commands-4, SPC4

SCSI Enclosure Services-2, SES2

Serial Attached SCSI-2, SAS2

附注

The fwflash command supports:

- InfiniBand Host Channel Adapters (IB HCAs) containing either the AMD or the Intel parallel flash parts.
- SCSI Enclosure Services devices such as SAS Expanders, attached with [ses\(7D\)](#) drivers.

引用名	fwtmp, wtmpfix – manipulate connect accounting records
用法概要	<pre>/usr/lib/acct/fwtmp [-ic] /usr/lib/acct/wtmpfix [file]...</pre>
描述	fwtmp reads from the standard input and writes to the standard output, converting binary records of the type found in /var/adm/wtmpx to formatted ASCII records. The ASCII version is useful when it is necessary to edit bad records. wtmpfix examines the standard input or named files in utmpx format, corrects the time/date stamps to make the entries consistent, and writes to the standard output. A hyphen (–) can be used in place of <i>file</i> to indicate the standard input. If time/date corrections are not performed, acctcon(1M) will fault when it encounters certain date-change records. Each time the date is set, a pair of date change records are written to /var/adm/wtmpx. The first record is the old date denoted by the string "old time" placed in the line field and the flag OLD_TIME placed in the type field of the utmpx structure. The second record specifies the new date and is denoted by the string new time placed in the line field and the flag NEW_TIME placed in the type field. wtmpfix uses these records to synchronize all time stamps in the file. In addition to correcting time/date stamps, wtmpfix will check the validity of the name field to ensure that it consists solely of alphanumeric characters or spaces. If it encounters a name that is considered invalid, it will change the login name to INVALID and write a diagnostic to the standard error. In this way, wtmpfix reduces the chance that acctcon will fail when processing connect accounting records.
选项	-ic Denotes that input is in ASCII form, and output is to be written in binary form.
文件	/var/adm/wtmpx history of user access and administration information
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/accounting/legacy-accounting

另请参见	acctcom(1) , ed(1) , acct(1M) , acctcms(1M) , acctcon(1M) , acctmerg(1M) , acctprc(1M) , acctsh(1M) , runacct(1M) , acct(2) , acct.h(3HEAD) , utmpx(4) , attributes(5) 《Oracle Solaris 管理：常见任务》

- 引用名

getdevpolicy – inspect the system's device policy
- 用法概要

/usr/sbin/getdevpolicy [*device...*]
- 描述

Without arguments, getdevpolicy outputs the device policy in effect to standard output.

With arguments, each argument is treated as a pathname to a device and the device policy in effect for that specific device is printed preceded by the supplied pathname.
- 用法

The device policy adds access restrictions over and above the file permissions.
- 退出状态

The following exit values are returned:

0	Successful completion.
non-zero	An error occurred.
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	See below

- 另请参见

The invocation is evolving. The output is unstable.

[add_drv\(1M\)](#), [rem_drv\(1M\)](#), [update_drv\(1M\)](#), [attributes\(5\)](#), [privileges\(5\)](#), [devfs\(7FS\)](#)

引用名

getent – get entries from administrative database

用法概要

getent *database* [*key*]...

描述

getent gets a list of entries from the administrative database specified by *database*. The information generally comes from one or more of the sources that are specified for the *database* in */etc/nsswitch.conf*.

database is the name of the database to be examined. This can be *passwd*, *group*, *hosts*, *ipnodes*, *services*, *protocols*, *auth_attr* or *exec_attr*. For each of these databases, getent uses the appropriate library routines described in [getpwnam\(3C\)](#), [getgrnam\(3C\)](#), [gethostbyaddr\(3NSL\)](#), [gethostbyname\(3NSL\)](#), [getipnodebyaddr\(3SOCKET\)](#), [getipnodebyname\(3SOCKET\)](#), [getservbyname\(3SOCKET\)](#), [getprotobyname\(3SOCKET\)](#), [getprofattr\(3C\)](#), [getauthattr\(3C\)](#), and [getexecattr\(3C\)](#), respectively.

Each *key* must be in a format appropriate for searching on the respective database. For example, it can be a *username* or *numeric-uid* for *passwd*; *hostname* or *IP address* for *hosts*; or *service*, *service/protocol*, *port*, or *port/proto* for *services*.

getent prints out the database entries that match each of the supplied keys, one per line, in the format of the matching administrative file: [passwd\(4\)](#), [group\(4\)](#), [project\(4\)](#), [networks\(4\)](#), [netmasks\(4\)](#), [user_attr\(4\)](#), [prof_attr\(4\)](#), [auth_attr\(4\)](#), or [exec_attr\(4\)](#). The key for [exec_attr\(4\)](#) is a profile name. If no key is given, all entries returned by the corresponding enumeration library routine, for example, `getpwent()` or `gethostent()`, are printed. Enumeration is not supported on *ipnodes*, *ethers*, and *netmasks*.

Key Interpretation for
passwd and group
Databases

When getent is invoked with database set to *passwd*, each key value is processed as follows:

- If the key value consists only of numeric characters, getent assumes that the key value is a numeric user ID and searches the user database for a matching user ID.
- If the user ID is not found in the user database or if the key value contains any non-numeric characters, getent assumes the key value is a user name and searches the user database for a matching user name.

Similarly, when getent is invoked with database set to *group*, each key value is processed as follows:

- If the key value consists only of numeric characters, getent assumes that the key value is a numeric group ID and searches the group database for a matching group ID.
- If the group ID is not found in the group database or if the key value contains any non-numeric characters, getent assumes the key value is a group name and searches the group database for a matching group name.

When getent is invoked with database set to [user_attr\(4\)](#), each key value is processed as follows:

- If the key value consists only of numeric characters, getent assumes that the key value is a numeric user ID and searches the user database for a matching user ID.

- If the key value contains any non-numeric characters, getent assumes the key value is a user name and searches the user database for a matching user name.

退出状态

The following exit values are returned:

- 0 Successful completion.
- 1 Command syntax was incorrect, an invalid option was used, or an internal error occurred.
- 2 At least one of the specified entry names was not found in the database.
- 3 There is no support for enumeration on this database.

文件

/etc/nsswitch.conf	name service switch configuration file
/etc/passwd	password file
/etc/group	group file
/etc/inet/hosts	IPv4 and IPv6 host name database
/etc/services	Internet services and aliases
/etc/project	project file
/etc/protocols	protocol name database
/etc/ethers	Ethernet address to hostname database or domain
/etc/networks	network name database
/etc/netmasks	network mask database
/etc/user_attr	Extended user attributes
/etc/security/prof_attr	Profile description database
/etc/security/auth_attr	Authorization description
/etc/security/exec_attr	Execution profiles database

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[ethers\(3SOCKET\)](#), [getgrnam\(3C\)](#), [gethostbyaddr\(3NSL\)](#), [gethostbyname\(3NSL\)](#), [gethostent\(3NSL\)](#), [getipnodebyaddr\(3SOCKET\)](#), [getipnodebyname\(3SOCKET\)](#), [getnetbyname\(3SOCKET\)](#), [getprojbyname\(3PROJECT\)](#), [getprotobyname\(3SOCKET\)](#), [getpwnam\(3C\)](#), [getservbyname\(3SOCKET\)](#), [getauthattr\(3C\)](#), [getexecattr\(3C\)](#),

```
getprofattr(3C), getuserattr(3C), ethers(4), group(4), hosts(4), netmasks(4),  
networks(4), nsswitch.conf(4), auth_attr(4), exec_attr(4), passwd(4), prof_attr(4),  
project(4), protocols(4), services(4), user_attr(4), attributes(5)
```

引用名 gettable – get DoD Internet format host table from a host

用法概要 /usr/sbin/gettable *host*

描述 gettable is a simple program used to obtain the DoD Internet host table from a “hostname” server. The specified *host* is queried for the table. The table is placed in the file `hosts.txt`.

gettable operates by opening a TCP connection to the port indicated in the service specification for “hostname”. A request is then made for all names and the resultant information is placed in the output file.

gettable is best used in conjunction with the [htable\(1M\)](#) program which converts the DoD Internet host table format to that used by the network library lookup routines.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/nis

另请参见 [htable\(1M\)](#), [attributes\(5\)](#) Harrenstien, Ken, Mary Stahl, and Elizabeth Feinler, HOSTNAME Server, RFC 953, Network Information Center, SRI International, Menlo Park, California, October 1985.

附注 Should allow requests for only part of the database.

引用名	getty – set terminal type, modes, speed, and line discipline
用法概要	<pre> /usr/lib/saf/ttymon [-h] [-t <i>timeout</i>] <i>line</i> [<i>speed</i> [<i>type</i> [<i>linedisc</i>]]] /usr/lib/saf/ttymon -c <i>file</i> </pre>
描述	getty sets terminal type, modes, speed, and line discipline. getty is a symbolic link to /usr/lib/saf/ttymon. It is included for compatibility with previous releases for the few applications that still call getty directly. getty can only be executed by the super-user, (a process with the user ID root). Initially getty prints the login prompt, waits for the user's login name, and then invokes the login command. getty attempts to adapt the system to the terminal speed by using the options and arguments specified on the command line. Without optional arguments, getty specifies the following: The <i>speed</i> of the interface is set to 300 baud, either parity is allowed, NEWLINE characters are converted to carriage return-line feed, and tab expansion is performed on the standard output. getty types the login prompt before reading the user's name a character at a time. If a null character (or framing error) is received, it is assumed to be the result of the user pressing the BREAK key. This will cause getty to attempt the next <i>speed</i> in the series. The series that getty tries is determined by what it finds in /etc/ttydefs .
选项	The following options are supported: -h If the -h flag is not set, a hangup will be forced by setting the speed to zero before setting the speed to the default or a specified speed. -t <i>timeout</i> Specifies that getty should exit if the open on the line succeeds and no one types anything in <i>timeout</i> seconds. -c <i>file</i> The -c option is no longer supported. Instead use /usr/sbin/sttydefs -l to list the contents of the /etc/ttydefs file and perform a validity check on the file.
操作数	The following operands are supported: <i>line</i> The name of a TTY line in /dev to which getty is to attach itself. getty uses this string as the name of a file in the /dev directory to open for reading and writing. <i>speed</i> The <i>speed</i> argument is a label to a speed and TTY definition in the file /etc/ttydefs. This definition tells getty at what speed to run initially, what the initial TTY settings are, and what speed to try next, (should the user press the BREAK key to indicate that the speed is inappropriate). The default <i>speed</i> is 300 baud. <i>type</i> and <i>linedisc</i> These options are obsolete and will be ignored.

文件 /etc/ttydefs

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [ct\(1C\)](#), [login\(1\)](#), [sttydefs\(1M\)](#), [ttypmon\(1M\)](#), [ioctl\(2\)](#), [attributes\(5\)](#), [tty\(7D\)](#)

引用名	gkadmin – Kerberos principals and policies administration GUI	
用法概要	<code>/usr/sbin/gkadmin</code>	
描述	<code>gkadmin</code> is an interactive graphical user interface (GUI) that enables you to maintain Kerberos principals and policies. <code>gkadmin</code> provides much the same functionality as the kadmin(1M) command. <code>gkadmin</code> does not support the management of keytabs. You must use <code>kadmin</code> for keytabs management. <code>gkadmin</code> uses Kerberos authentication and an encrypted RPC to operate securely from anywhere on the network. When <code>gkadmin</code> is invoked, the login window is populated with default values. For the principal name, <code>gkadmin</code> determines your user name from the <code>USER</code> environment variable. It appends <code>/admin</code> to the name (<code>username/admin</code>) to create a default user instance in the same manner as <code>kadmin</code>. It also selects appropriate defaults for realm and master KDC (<code>admin_server</code>) from the <code>/etc/krb5/krb5.conf</code> file. You can change these defaults on the login window. When you enter your password, a session is started with <code>kadmin</code>. Operations performed are subject to permissions that are granted or denied to the chosen user instance by the Kerberos ACL file. See kadm5.acl(4). After the session is started, a tabbed folder is displayed that contains a principal list and a policy list. The functionality is mainly the same as <code>kadmin</code>, with addition, deletion, and modification of principal and policy data available. <code>gkadmin</code> also includes an interface to specify principal key encryption types when modifying or creating principal records. The default set of encryption types is used if they are not selected through this interface. The default set of encryption types can be found in krb5.conf(4) under the <code>default_tkt_enctypes</code> section. In addition, <code>gkadmin</code> provides the following features: ■ New principal or policy records can be added either from default values or from the settings of an existing principal. ■ A comment field is available for principals. ■ Default values are saved in <code>\$HOME/.gkadmin</code>. ■ A logout option permits you to log back in as another user instance without exiting the tool. ■ Principal and policy lists and attributes can be printed or saved to a file. ■ Online context-sensitive help and general help is available in the <code>Help</code> menu.	
文件	<code>/etc/krb5/krb5.conf</code>	Kerberos configuration information on a Kerberos client. Used to search for default realm and master KDC (<code>admin_server</code>), including a port number for the master KDC.

\$HOME/.gkadmin	Default parameters used to initialize new principals created during the session.

属性

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/security/kerberos-5
Interface Stability	Committed

另请参见

```
kpasswd(1), kadmind(1M), kadmind(1M), kadmind.local(1M), kdb5_util(1M), kadmind.acl(4),  
kdc.conf(4), krb5.conf(4), attributes(5), kerberos(5)
```

诊断

The `gkadmin` interface is currently incompatible with the MIT `kadmin` daemon interface, so you cannot use this interface to administer an MIT-based Kerberos database. However, clients running the Solaris implementation of Kerberos can still use an MIT-based KDC.

引用名	groupadd – add (create) a new group definition on the system
用法概要	<pre>/usr/sbin/groupadd [-g <i>gid</i> [-o]] [-S <i>repository</i>] [-U <i>user1[,user2...]</i>] <i>group</i></pre>
描述	The groupadd command creates a new group definition on the system by adding the appropriate entry to the group database in the files and ldap repositories. An administrator must be granted the User Management rights profile or have <code>solaris.group.manage</code> authorization to be able to add a group. Once the group is successfully added, the administrator is granted the authorization to modify and delete the group. The format of the authorization is <code>solaris.group.assign/groupname</code>, and can be managed by means of usermod(1M).
选项	The following options are supported: -g <i>gid</i> Assigns the group id <i>gid</i> for the new group. This group id must be a non-negative decimal integer below MAXUID as defined in <code>/usr/include/sys/param.h</code>. The group ID defaults to the next available (unique) number above the highest number currently assigned. For example, if groups 100, 105, and 200 are assigned as groups, the next default group number is 201. (Group IDs from 0–99 are reserved by SunOS for future applications.) -o Allows the <i>gid</i> to be duplicated (non-unique). An administrator must have <code>solaris.group.assign</code> authorization to use this option. -S <i>repository</i> The valid repositories are files and ldap. The repository specifies which name service will be updated. When <i>repository</i> is not specified, the files repository is used. When the repository is files, the user name and other items can be present in other name service repositories and can be assigned to a group in the files repository. When the repository is ldap, all the assignable attributes must be present in the ldap repository. -U <i>user1[,user2]</i> Adds a list of users <i>user1</i>, <i>user2</i> to the group.
操作数	The following operands are supported: <i>group</i> A string consisting of characters from the set of lower case alphabetic characters and numeric characters. A warning message is written if the string exceeds MAXGLEN - 1, which is usually eight characters. The <i>group</i> field must contain at least one character; it accepts lower case or numeric characters or a combination of both, and must not contain a colon (:) or NEWLINE.
退出状态	The following exit values are returned: 0 Successful completion. 2 Invalid command syntax. A usage message for the groupadd command is displayed.

- 3 An invalid argument was provided to an option.
- 4 The *gid* is not unique (when -o option is not used).
- 9 The *group* is not unique.
- 10 The group database cannot be updated.

文件

- /etc/group
- /usr/include/userdefs.h

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[users\(1B\)](#), [groupdel\(1M\)](#), [groupmod\(1M\)](#), [grpck\(1M\)](#), [logins\(1M\)](#), [pwck\(1M\)](#), [useradd\(1M\)](#), [userdel\(1M\)](#), [usermod\(1M\)](#), [group\(4\)](#), [attributes\(5\)](#)

附注

groupadd adds a group definition to the system. If a network name service is being used to supplement the local /etc/group file with additional entries, groupadd verifies the uniqueness of a specified group name and group ID against the external name service and uses the entries in the files repository.

If the number of characters in a group entry exceeds 2047, group maintenance commands, such as [groupdel\(1M\)](#) and [groupmod\(1M\)](#), fail.

引用名	groupdel – delete a group definition from the system
用法概要	<code>/usr/sbin/groupdel [-S repository] group</code>
描述	The <code>groupdel</code> utility deletes a group definition from the system. It deletes the appropriate entry from the <code>/etc/group</code> file. In addition to <code>solaris.group.manage</code> authorization, an administrator must have either <code>solaris.group.assign</code> or a matching authorization of the form <code>solaris.group.assign/groupname</code> to delete a group. The authorization <code>solaris.group.assign/groupname</code> is automatically assigned to the administrator who created the group.
选项	<code>groupdel</code> supports the following option. <code>-S repository</code> The valid repositories are <code>files</code> and <code>ldap</code>. The repository specifies which name service will be updated. When <code>repository</code> is not specified, <code>groupdel</code> consults <code>nsswitch.conf(4)</code>.
操作数	<code>group</code> An existing group name to be deleted.
退出状态	The following exit values are returned: 0 Success. 2 Invalid command syntax. A usage message for the <code>groupdel</code> command is displayed. 6 <code>group</code> does not exist. 10 Cannot update the <code>/etc/group</code> file.
文件	<code>/etc/group</code> system file containing group definitions
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见	users(1B) , groupadd(1M) , groupmod(1M) , logins(1M) , useradd(1M) , userdel(1M) , usermod(1M) , nsswitch.conf(4) , attributes(5)
附注	The <code>groupdel</code> utility deletes a group definition that is in the group database in the repository. <code>groupdel</code> fails if a group entry (a single line in <code>/etc/group</code>) exceeds 2047 characters.

引用名	groupmod – modify a group definition on the system
用法概要	<pre>/usr/sbin/groupmod [-S repository] [-g gid [-o]] [-n name] [-U [+ -]user1[,user2]...] group</pre>
描述	The groupmod command modifies the definition of the specified group by modifying the appropriate entry in the group database in the repository. An administrator can modify any group for which it has a matching authorization of the form <code>solaris.group.assign/groupname</code>. This authorization is automatically assigned to the administrator who created the group. An administrator must have <code>solaris.group.assign</code> authorization to modify all other groups.
选项	The following options are supported: -g <i>gid</i> Specify the new group ID for the group. This group ID must be a non-negative decimal integer less than MAXUID, as defined in <code><param.h></code>. The group ID defaults to the next available (unique) number above 99. (Group IDs from 0-99 are reserved by SunOS for future applications.) -n <i>name</i> Specify the new name for the group. The <i>name</i> argument is a string of no more than eight bytes consisting of characters from the set of lower case alphabetic characters and numeric characters. A warning message will be written if these restrictions are not met. A future Solaris release may refuse to accept group fields that do not meet these requirements. The <i>name</i> argument must contain at least one character and must not include a colon (:) or NEWLINE (\n). -o Allow the <i>gid</i> to be duplicated (non-unique). An administrator must have <code>solaris.group.assign</code> authorization to use this option. -S <i>repository</i> The valid repositories are <code>files</code> and <code>ldap</code>. The repository specifies which name service will be updated. When <i>repository</i> is not specified, groupmod consults <code>nsswitch.conf(4)</code>. When the repository is <code>files</code>, the user name and other items can be present in other name service repositories and can be assigned to a group in the <code>files</code> repository. When the repository is <code>ldap</code>, all the assignable attributes must be present in the <code>ldap</code> repository. -U [+ -]user1[,user2] Updates the list of users for the group as follows: ▪ A prefix + before the list adds that list to existing users list. ▪ A prefix - before the list removes each user in the list from the existing users list. ▪ With no prefix before the list, replaces the existing users list with the new list of users specified.

- 操作数

The following operands are supported:
group An existing group name to be modified.
- 退出状态

The groupmod utility exits with one of the following values:

0 Success.
2 Invalid command syntax. A usage message for the groupmod command is displayed.
3 An invalid argument was provided to an option.
4 *gid* is not unique (when the -o option is not used).
6 *group* does not exist.
9 *name* already exists as a group name.
10 Cannot update the /etc/group file.
- 文件

/etc/group group file
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

- 另请参见

[users\(1B\)](#), [groupadd\(1M\)](#), [groupdel\(1M\)](#), [logins\(1M\)](#), [useradd\(1M\)](#), [userdel\(1M\)](#), [usermod\(1M\)](#), [group\(4\)](#), [nsswitch.conf\(4\)](#), [attributes\(5\)](#)
- 附注

The groupmod utility only modifies group definitions in the group database in the repository. If a network name service such as NIS is being used to supplement the local /etc/group file with additional entries, groupmod cannot change information supplied by the network name service. groupmod verifies the uniqueness of group name and group ID against the external name service and uses the entries in the files repository.

groupmod fails if a group entry (a single line in /etc/group) exceeds 2047 characters.

引用名 growfs – non-destructively expand a UFS file system

用法概要 /usr/sbin/growfs [-M *mount-point*] [*newfs-options*]
[*raw-device*]

描述 growfs non-destructively expands a mounted or unmounted UNIX file system (UFS) to the size of the file system's slice(s).

Typically, disk space is expanded by first adding a slice to a metadvice, then running the growfs command. When adding space to a mirror, you expand each submirror before expanding the file system.

growfs will “write-lock” (see [lockfs\(1M\)](#)) a mounted file system when expanding. The length of time the file system is write-locked can be shortened by expanding the file system in stages. For instance, to expand a 1 Gbyte file system to 2 Gbytes, the file system can be grown in 16 Mbyte stages using the -s option to specify the total size of the new file system at each stage. The argument for -s is the number of sectors, and must be a multiple of the cylinder size. Note: The file system cannot be grown if a cylinder size of less than 2 is specified. Refer to the [newfs\(1M\)](#) man page for information on the options available when growing a file system.

growfs displays the same information as mkfs during the expansion of the file system.

If growfs is aborted, recover any lost free space by unmounting the file system and running the fsck command, or run the growfs command again.

Note: If growfs is aborted and the file system is used before fsck is run on it, UFS metadata might be left in an incomplete state, with the result that the file system would be corrupted. In such a circumstance, you would have to restore the file system from backups.

选项 Root privileges are required for all of the following options.

-M <i>mount-point</i>	The file system to be expanded is mounted on <i>mount-point</i> . File system locking (lockfs) will be used.
<i>newfs-options</i>	The options are documented in the newfs man page.
<i>raw-device</i>	Specifies the name of a raw metadvice or raw special device, residing in /dev/md/rdisk, or /dev/rdisk, respectively, including the disk slice, where you want the file system to be grown.

示例 示例1 Expanding nonmetadvice slice for /export file system

The following example expands a nonmetadvice slice for the /export file system. In this example, the existing slice, /dev/dsk/clt0d0s3, is converted to a metadvice so additional slices can be concatenated.

```
# metainit -f d8 2 1 clt0d0s3 1 c2t0d0s3
# umount /export
```


示例 2 Associate /export with new metadvice

Edit the `/etc/vfstab` file to change the entry for `/export` to the newly defined metadvice, `d8`.

```
# mount /export
# growfs -M /export /dev/md/rdisk/d8
```

The first example starts by running the `metainit` command with the `-f` option to force the creation of a new concatenated metadvice `d8`, which consists of the existing slice `/dev/dsk/c1t0d0s3` and a new slice `/dev/dsk/c2t0d0s3`. Next, the file system on `/export` must be unmounted. The `/etc/vfstab` file is edited to change the entry for `/export` to the newly defined metadvice name, rather than the slice name. After the file system is remounted, the `growfs` command is run to expand the file system. The file system will span the entire metadvice when `growfs` completes. The `-M` option enables the `growfs` command to expand a mounted file system. During the expansion, write access for `/export` is suspended until `growfs` unlocks the file system. Read access is not affected, though access times are not kept when the lock is in effect.

示例 3 Dynamic Expansion of /export file system

The following example picks up from the previous one. Here, the `/export` file system mounted on metadvice `d8` is dynamically expanded.

```
# metattach d8 c0t1d0s2
# growfs -M /export /dev/md/rdisk/d8
```

This example begins by using the `metattach` command to dynamically concatenate a new slice, `/dev/dsk/c0t1d0s2`, to the end of an existing metadvice, `d8`. Next, the `growfs` command specifies that the mount-point is `/export` and that it is to be expanded onto the raw metadvice `/dev/md/rdisk/d8`. The file system will span the entire metadvice when `growfs` completes. During the expansion, write access for `/export` is suspended until `growfs` unlocks the file system. Read access is not affected, though access times are not kept when the lock is in effect.

示例 4 Expanding mounted file system to existing mirror

The following example expands a mounted file system `/files`, to an existing mirror, `d80`, which contains two submirrors, `d9` and `d10`.

```
# metattach d9 c0t2d0s5
# metattach d10 c0t3d0s5
# growfs -M /files /dev/md/rdisk/d80
```

In this example, the `metattach` command dynamically concatenates the new slices to each submirror. The `metattach` command must be run for each submirror. The mirror will automatically grow when the last submirror is dynamically concatenated. The mirror will grow to the size of the smallest submirror. The `growfs` command then expands the file system. The `growfs` command specifies that the mount-point is `/files` and that it is to be expanded

示例 4 Expanding mounted file system to existing mirror (续)

onto the raw metadevice /dev/md/rdisk/d80. The file system will span the entire mirror when the growfs command completes. During the expansion, write access for the file system is suspended until growfs unlocks the file system. Read access is not affected, though access times are not kept when the lock is in effect.

退出状态 The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm

另请参见 [fsck\(1M\)](#), [lockfs\(1M\)](#), [mkfs\(1M\)](#), [metattach\(1M\)](#), [newfs\(1M\)](#), [attributes\(5\)](#)

Solaris Volume Manager Administration Guide

Limitations Only UFS file systems (either mounted or unmounted) can be expanded using the growfs command. Once a file system is expanded, it cannot be decreased in size. The following conditions prevent you from expanding file systems: When acct is activated and the accounting file is on the target device. When C2 security is activated and the logging file is on the target file system. When there is a local swap file in the target file system. When the file system is root (/), /usr, or swap.

引用名	gsscred – add, remove, and list gsscred table entries																
用法概要	<pre>gsscred [-n user [-o oid] [-u uid]] [-c comment] -m mech -a gsscred [-n user [-o oid]] [-u uid] [-m mech] -r gsscred [-n user [-o oid]] [-u uid] [-m mech] -l</pre>																
描述	The gsscred utility is used to create and maintain a mapping between a security principal name and a local UNIX <i>uid</i>. The format of the user name is assumed to be GSS_C_NT_USER_NAME. You can use the -o option to specify the object identifier of the <i>name</i> type. The OID must be specified in dot-separated notation, for example: 1.2.3.45464.3.1 The gsscred table is used on server machines to lookup the <i>uid</i> of incoming clients connected using RPCSEC_GSS. When adding users, if no <i>user</i> name is specified, an entry is created in the table for each user from the passwd table. If no <i>comment</i> is specified, the gsscred utility inserts a comment that specifies the user name as an ASCII string and the GSS-API security mechanism that applies to it. The security mechanism will be in string representation as defined in the /etc/gss/mech file. The parameters are interpreted the same way by the gsscred utility to delete users as they are to create users. At least one of the following options must be specified: -n, -u, or -m. If no security mechanism is specified, then all entries will be deleted for the user identified by either the <i>uid</i> or <i>user</i> name. If only the security mechanism is specified, then all <i>user</i> entries for that security mechanism will be deleted. Again, the parameters are interpreted the same way by the gsscred utility to search for users as they are to create users. If no options are specified, then the entire table is returned. If the <i>user</i> name or <i>uid</i> is specified, then all entries for that <i>user</i> are returned. If a security mechanism is specified, then all <i>user</i> entries for that security mechanism are returned.																
选项	<table><tr><td>-a</td><td>Add a table entry.</td></tr><tr><td>-c <i>comment</i></td><td>Insert comment about this table entry.</td></tr><tr><td>-l</td><td>Search table for entry.</td></tr><tr><td>-m <i>mech</i></td><td>Specify the mechanism for which this name is to be translated.</td></tr><tr><td>-n <i>user</i></td><td>Specify the optional principal name.</td></tr><tr><td>-o <i>oid</i></td><td>Specify the OID indicating the name type of the user.</td></tr><tr><td>-r</td><td>Remove the entry from the table.</td></tr><tr><td>-u <i>uid</i></td><td>Specify the <i>uid</i> for the <i>user</i> if the <i>user</i> is not local.</td></tr></table>	-a	Add a table entry.	-c <i>comment</i>	Insert comment about this table entry.	-l	Search table for entry.	-m <i>mech</i>	Specify the mechanism for which this name is to be translated.	-n <i>user</i>	Specify the optional principal name.	-o <i>oid</i>	Specify the OID indicating the name type of the user.	-r	Remove the entry from the table.	-u <i>uid</i>	Specify the <i>uid</i> for the <i>user</i> if the <i>user</i> is not local.
-a	Add a table entry.																
-c <i>comment</i>	Insert comment about this table entry.																
-l	Search table for entry.																
-m <i>mech</i>	Specify the mechanism for which this name is to be translated.																
-n <i>user</i>	Specify the optional principal name.																
-o <i>oid</i>	Specify the OID indicating the name type of the user.																
-r	Remove the entry from the table.																
-u <i>uid</i>	Specify the <i>uid</i> for the <i>user</i> if the <i>user</i> is not local.																

示例

示例 1 Creating a gsscred Table for the Kerberos v5 Security Mechanism

The following shows how to create a gsscred table for the kerberos v5 security mechanism. gsscred obtains *user* names and *uid*'s from the passwd table to populate the table.

```
example% gsscred -m kerberos_v5 -a
```

示例 2 Adding an Entry for root/host1 for the Kerberos v5 Security Mechanism

The following shows how to add an entry for root/host1 with a specified *uid* of 0 for the kerberos v5 security mechanism.

```
example% gsscred -m kerberos_v5 -n root/host1 -u 0 -a
```

示例 3 Listing All User Mappings for the Kerberos v5 Security Mechanism

The following lists all user mappings for the kerberos v5 security mechanism.

```
example% gsscred -m kerberos_v5 -l
```

示例 4 Listing All Mappings for All Security Mechanism for a Specified User

The following lists all mappings for all security mechanisms for the user bsimpson.

```
example% gsscred -n bsimpson -l
```

退出状态

The following exit values are returned:

0 Successful completion.

>0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/library/security/gss
Interface Stability	Committed

另请参见

[gssd\(1M\)](#), [gsscred.conf\(4\)](#), [attributes\(5\)](#)

附注

Some GSS mechanisms, such as kerberos_v5, provide their own authenticated-name-to-local-name (uid) mapping and thus do not usually have to be mapped using gsscred. See [gsscred.conf\(4\)](#) for more information.

引用名 gssd – generates and validates GSS-API tokens for kernel RPC

用法概要 /usr/lib/gss/gssd

描述 gssd is the user mode daemon that operates between the kernel rpc and the Generic Security Service Application Program Interface (GSS-API) to generate and validate GSS-API security tokens. In addition, gssd maps the GSS-API principal names to the local user and group ids. By default, all groups that the requested user belongs to will be included in the grouplist credential. gssd is invoked by the Internet daemon [inetd\(1M\)](#) the first time that the kernel RPC requests GSS-API services.

退出状态 The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/kernel/security/gss
Interface Stability	Committed

另请参见 [kill\(1\)](#), [pkill\(1\)](#), [svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [gsscred\(1M\)](#), [svcadm\(1M\)](#), [gsscred.conf\(4\)](#), [resolv.conf\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

RFC 2078

附注 The following signal has the specified effect when sent to the server process using the [kill\(1\)](#) command:

SIGHUP gssd rereads the [gsscred.conf\(4\)](#) options.

When one of the mechanisms being used is Kerberos, then the gssd process must be restarted after adding or changing the [resolv.conf\(4\)](#) file.

The gssd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/rpc/gss:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名	hald – 用于支持硬件抽象层的守护进程										
用法概要	<pre>/usr/lib/hal/hald [--daemon={yes no}] [--help] [--use-syslog] [--verbose={yes no}] [--version]</pre>										
描述	hald 守护进程可以识别对符合硬件抽象层 (Hardware Abstraction Layer, HAL) 规范的设备所做的硬件更改。 可以通过服务管理工具 (service management facility, SMF) 执行对 hald 的禁用和启用操作（请参见 smf(5)）。使用故障管理资源标识符 (fault management resource identifier, FMRI) <code>svc:/system/hal</code> 管理 hald。										
选项	所支持的选项如下： <table><tr><td><code>--daemon=yes no</code></td><td>作为守护进程运行。</td></tr><tr><td><code>--help</code></td><td>显示用法信息并退出。</td></tr><tr><td><code>--use-syslog</code></td><td>将调试消息输出到 <code>syslog</code> 中，而不是输出到 <code>stderr</code> 中。如果 HAL 作为守护进程运行，请使用此选项记录调试消息。</td></tr><tr><td><code>--verbose=yes no</code></td><td>显示调试消息。</td></tr><tr><td><code>--version</code></td><td>显示版本信息并退出。</td></tr></table>	<code>--daemon=yes no</code>	作为守护进程运行。	<code>--help</code>	显示用法信息并退出。	<code>--use-syslog</code>	将调试消息输出到 <code>syslog</code> 中，而不是输出到 <code>stderr</code> 中。如果 HAL 作为守护进程运行，请使用此选项记录调试消息。	<code>--verbose=yes no</code>	显示调试消息。	<code>--version</code>	显示版本信息并退出。
<code>--daemon=yes no</code>	作为守护进程运行。										
<code>--help</code>	显示用法信息并退出。										
<code>--use-syslog</code>	将调试消息输出到 <code>syslog</code> 中，而不是输出到 <code>stderr</code> 中。如果 HAL 作为守护进程运行，请使用此选项记录调试消息。										
<code>--verbose=yes no</code>	显示调试消息。										
<code>--version</code>	显示版本信息并退出。										
文件	<table><tr><td><code>/usr/lib/hal</code></td><td>与 HAL 相关的文件</td></tr><tr><td><code>/etc/hal/fdi</code></td><td>设备信息文件</td></tr></table>	<code>/usr/lib/hal</code>	与 HAL 相关的文件	<code>/etc/hal/fdi</code>	设备信息文件						
<code>/usr/lib/hal</code>	与 HAL 相关的文件										
<code>/etc/hal/fdi</code>	设备信息文件										
属性	有关以下属性的说明，请参见 attributes(5) ：										

属性类型	属性值
可用性	system/hal
接口稳定性	Volatile（可变）

另请参见 [svccfg\(1M\)](#)、[attributes\(5\)](#)、[hal\(5\)](#)

- 引用名

hal-device – 管理 HAL 设备
- 用法概要

```
hal-device [-h] [--a udi | --r udi]
```
- 描述

硬件抽象层 (Hardware Abstraction Layer, HAL) 提供了连接到系统的各种硬件的视图。使用 `hal-device` 命令可以管理符合 HAL 标准的设备。具体而言，可以使用 `hal-device` 在 HAL 全局设备列表中添加或删除设备。设备属性是从遵循 `lshal(1M)` 语法的 `stdin` 中读取的。
- 选项

所支持的选项如下：

-a、--add *udi*

向 HAL 的全局设备列表中添加由通用设备标识符 (Universal Device Identifier, *udi*) 指定的设备。

-h

显示用法信息。

-r、--remove *udi*

从 HAL 的全局设备列表中删除由通用设备标识符 (Universal Device Identifier, *udi*) 指定的设备。
- 属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/hal
接口稳定性	Volatile（可变）

另请参见

[hald\(1M\)](#)、[attributes\(5\)](#)、[hal\(5\)](#)

- 引用名

hal-fdi-validate – 验证 HAL 设备信息文件
- 用法概要

hal-fdi-validate [-f dtd] file [file]...
- 描述

hal-fdi-validate 命令可以对一个或多个设备信息文件进行验证。请参见 [fdi\(4\)](#)。除非使用 -f 选项指定了其他文件，否则将使用标准 DTD 文件。
- 选项

支持以下选项：

-f dtd 指定 DTD 文件的路径。
- 操作数

hal-fdi-validate 命令接受以下操作数：

file [file...] 要验证的一个或多个 DTD 文件。
- 属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/hal
接口稳定性	Volatile（可变）

另请参见

[hald\(1M\)](#)、[fdi\(4\)](#)、[attributes\(5\)](#)、[hal\(5\)](#)

引用名

hal-find, hal-find-by-capability, hal-find-by-property – 搜索 HAL 全局设备列表

用法概要

```
hal-find-by-capability --capability capability [--help]
[--verbose] [--version]

hal-find-by-property --key key --string value [--help]
[--verbose] [--version]
```

描述

hal-find 命令 hal-find-by-capability 和 hal-find-by-property 按指定条件搜索硬件抽象层 (Hardware Abstraction Layer, HAL) 设备列表并将结果显示在标准输出中。hal-find-by-capability 按容量搜索，例如 volume 或 block。hal-find-by-property 按属性搜索，例如 block.is_volume 或 volume.disc.has_audio。

选项

所支持的选项如下：

--capability <i>capability</i>	要搜索的 HAL 设备容量。
--help	显示选项列表。
--key <i>key</i>	作为搜索基础的属性的键。
--string <i>value</i>	与作为搜索基础的属性关联的字符串值。
--冗长	详细模式。
--version	显示版本并退出。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/hal
接口稳定性	Volatile（可变）

另请参见

[hald\(1M\)](#)、[attributes\(5\)](#)、[hal\(5\)](#)

引用名 hal-get-property, hal-set-property – 获取和设置 HAL 设备属性

用法概要

```
hal-get-property --udi udi --key key [--help] [--verbose]
[--version]

hal-set-property --udi udi --key key {--int value | --uint64 value
| --string value | --bool value | --strlist-pre value
| --strlist-post value | --strlist-rem value | --double value
| --remove value} [--help] [--version]
```

描述

硬件抽象层 (Hardware Abstraction Layer, HAL) 提供了连接到系统的各种硬件的视图。当硬件配置通过热插拔或其他机制发生更改时，此视图也将随之进行动态更新。HAL 将硬件表示为设备对象。设备对象由唯一标识符进行标识并带有一组键/值对，这些键/值对称为设备属性。有些属性源自实际硬件，有些属性由设备信息文件（.fdi 文件）合并而来，有些属性与实际设备配置相关。

可以使用 `hal-get-property` 和 `hal-set-property` 命令获取和设置符合 HAL 规范的硬件的属性。

选项

所支持的选项如下：

```
--udi udi
    唯一设备 ID。

--key key
    要设置的属性的键。

--int
    将值设置为一个整数。接受前缀为 0x 或 x 的十进制或十六进制值。

--uint64
    将值设置为一个整数。接受前缀为 0x 或 x 的十进制或十六进制值。

--string value
    将值设置为一个字符串。

--double value
    将值设置为一个浮点数。

--boolean value
    将值设置为一个布尔值，即 true 或 false

--strlist-pre value
    将一个字符串附加到列表的开头。

--strlist-post value
    将一个字符串附加到列表的末尾。

--strlist-rem value
    从列表中删除一个字符串。
```

- `--remove value`
指示应删除的属性。
- `--version`
显示版本并退出。
- `--help`
显示选项列表并退出

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/hal
接口稳定性	Volatile（可变）

另请参见 [hald\(1M\)](#)、[attributes\(5\)](#)、[hal\(5\)](#)

引用名	halt, poweroff – 停止处理器
用法概要	<pre>/usr/sbin/halt [-dlnqy] /usr/sbin/poweroff [-dlnqy]</pre>
描述	halt 和 poweroff 实用程序将所有待处理信息写入磁盘，然后停止处理器。如果可能，poweroff 实用程序会关闭计算机电源。 halt 和 poweroff 实用程序通常将系统关机信息记录到系统日志守护进程 syslogd(1M) 中，并且会在登录记帐文件 <code>/var/adm/wtmpx</code> 中放置一条关机记录。如果提供了 <code>-n</code> 或 <code>-q</code> 选项，则会禁止这些操作。
选项	所支持的选项如下： -d 强制在重新引导前执行系统故障转储。有关配置系统故障转储的信息，请参见 dumpadm(1M)。 -l 禁止向系统日志守护进程 syslogd(1M) 发送有关谁执行了 halt 的消息。 -n 阻止在停止前执行 sync(1M)。 -q 快速停止。没有尝试正常关机。 -y 停止系统，即使是通过拨号终端停止。
文件	<code>/var/adm/wtmpx</code> 有关用户访问和管理信息的历史记录。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	system/core-os

另请参见	dumpadm(1M) 、 init(1M) 、 reboot(1M) 、 shutdown(1M) 、 sync(1M) 、 syslogd(1M) 、 inittab(4) 、 att
附注	halt 和 poweroff 实用程序不会完全关闭 smf(5) 服务，不执行 <code>/etc/rcnum.d</code> 中的脚本，也不执行 inittab(4) 中的关机操作。要确保完全关闭系统服务，请使用 shutdown(1M) 或 init(1M) 重新引导 Solaris 系统。

引用名	hextoalabel – convert an internal text label to its human readable equivalent
用法概要	<pre>/usr/sbin/hextoalabel [<i>internal-text-sensitivity-label</i>] /usr/sbin/hextoalabel -c [<i>internal-text-clearance</i>]</pre>
描述	hextoalabel converts an internal text label into its human readable equivalent and writes the result to the standard output file. This internal form is often hexadecimal. If no option is supplied, the label is assumed to be a sensitivity label. If no internal text label is specified, the label is read from the standard input file. The expected use of this command is emergency repair of labels that are stored in internal databases.
选项	-c Identifies the internal text label as a clearance.
退出状态	The following exit values are returned: 0 On success. 1 On failure, and writes diagnostics to the standard error file.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/trusted
Interface Stability	See below.

	The command output is Committed for systems with the same <code>label_encodings</code> file. The command invocation is Committed for systems that implement the DIA MAC policy.
文件	<code>/etc/security/tsol/label_encodings</code> The label encodings file contains the classification names, words, constraints, and values for the defined labels of this system.
另请参见	atohexlabel(1M), label_to_str(3TSOL), str_to_label(3TSOL), label_encodings(4), attributes(5) 《Trusted Extensions 配置和管理》
附注	The functionality described on this manual page is available only if the system is configured with Trusted Extensions. This file is part of the Defense Intelligence Agency (DIA) Mandatory Access Control (MAC) policy. This file might not be applicable to other MAC policies that might be developed for future releases of Solaris Trusted Extensions software.

引用名 host – DNS lookup utility

用法概要 host [-aCdilMrsTw] [-c *class*] [-N *ndots*] [-R *number*]
 [-t *type*] [-W *wait*] [-4 | -6] *name* [*server*]

描述 The `host` utility performs simple DNS lookups. It is normally used to convert names to IP addresses and IP addresses to names. When no arguments or options are given, `host` prints a short summary of its command line arguments and options.

The *name* argument is the domain name that is to be looked up. It can also be a dotted-decimal IPv4 address or a colon-delimited IPv6 address, in which case `host` by default performs a reverse lookup for that address. The optional *server* argument is either the name or IP address of the name server that `host` should query instead of the server or servers listed in `/etc/resolv.conf`.

选项 The following options are supported:

- 4
Use only IPv4 transport. By default, both IPv4 and IPv6 transports can be used. Options -4 and -6 are mutually exclusive.
- 6
Use only IPv6 transport. By default, both IPv4 and IPv6 transports can be used. Options -4 and -6 are mutually exclusive.
- a
Equivalent to setting the -v option and asking `host` to make a query of type ANY.
- c *class*
Make a DNS query of class *class*. This can be used to lookup Hesiod or Chaosnet class resource records. The default class is IN (Internet).
- C
Attempt to display the SOA records for zone *name* from all the listed authoritative name servers for that zone. The list of name servers is defined by the NS records that are found for the zone.
- d
Generate verbose output. This option is equivalent to -v. These two options are provided for backward compatibility. In previous versions, the -d option switched on debugging traces and -v enabled verbose output.
- i
Specifies that reverse lookups of IPv6 addresses should use the IP6.INT domain as defined in RFC 1886. The default is to use RFC 3152 domain IP6.ARPA.

- l
List mode. This option makes `host` perform a zone transfer for zone *name*, displaying the NS, PTR and address records (A/AAAA). If combined with `-a`, all records will be displayed. The argument is provided for compatibility with previous implementations. Options `-la` is equivalent to making a query of type AXFR.
- m
Sets the memory usage debugging flags: record, usage, and trace.
- N *ndots*
Set the number of dots that have to be in *name* for it to be considered absolute. The default value is that defined using the `ndots` statement in `/etc/resolv.conf`, or 1 if no `ndots` statement is present. Names with fewer dots are interpreted as relative names and will be searched for in the domains listed in the `search` or `domain` directive in `/etc/resolv.conf`.
- r
Make a non-recursive query. Setting this option clears the RD (recursion desired) bit in the query made by `host`. The name server receiving the query does not attempt to resolve *name*. The `-r` option enables `host` to mimic the behaviour of a name server by making non-recursive queries and expecting to receive answers to those queries that are usually referrals to other name servers.
- R *number*
Change the number of UDP retries for a lookup. The *number* argument indicates how many times `host` will repeat a query that does not get answered. The default number of retries is 1. If *number* is negative or zero, the number of retries will default to 1.
- s
Specifies that the host not send the query to the next name server if any server responds with a SERVFAIL response, which is the reverse of normal stub resolver behavior.
- t *type*
Select the query type. The *type* argument can be any recognised query type: CNAME, NS, SOA, SIG, KEY, and AXFR, among others. When no query type is specified, `host` automatically selects an appropriate query type. By default it looks for A, AAAA, and MX records, but if the `-C` option is specified, queries are made for SOA records. If *name* is a dotted-decimal IPv4 address or colon-delimited IPv6 address, `host` queries for PTR records.

If a query type of IXFR is chosen the starting serial number can be specified by appending an equal followed by the starting serial number (for example: `-t IXFR=12345678`).
- T
Use a TCP connection when querying the name server. TCP is automatically selected for queries that require it, such as zone transfer (AXFR) requests. By default `host` uses UDP when making queries.
- v
Generate verbose output. This option is equivalent to `-d`.

-w
Wait forever for a reply. The time to wait for a response will be set to the number of seconds given by the hardware's maximum value for an integer quantity.

-W *wait*
Wait for *wait* seconds for a reply. If *wait* is less than one, the wait interval is set to one second.

文件
/etc/resolv.conf
Resolver configuration file

属性
See for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	network/dns/bind
Interface Stability	Volatile

另请参见
[dig\(1M\)](#), [named\(1M\)](#), [attributes\(5\)](#)

RFC 1035, RFC 1886, RFC 3152

See the BIND 9 Administrator's Reference Manual. As of the date of publication of this man page, this document is available at <https://www.isc.org/software/bind/documentation>.

引用名	hostconfig – configure a system's host parameters
用法概要	<pre>/usr/sbin/hostconfig -p <i>protocol</i> [-d] [-h] [-n] [-v] [-i <i>interface</i>] [-f <i>hostname</i>]</pre>
描述	The hostconfig program uses a network protocol to acquire a machine's <i>host parameters</i> and set these parameters on the system. The program selects which protocol to use based on the argument to the required -p flag. Different protocols may set different host parameters. Currently, only one protocol (bootparams) is defined.
选项	The following options are supported: -d Enable debug output. -f <i>hostname</i> Run the protocol as if this machine were named <i>hostname</i>. -h Echo the received <i>hostname</i> to stdout, rather than setting <i>hostname</i> using the system name directly. -i <i>interface</i> Use only the named network interface to run the protocol. -n Run the network protocol, but do not set the acquired parameters into the system. -p <i>protocol</i> Run hostconfig using <i>protocol</i>. Currently, only one protocol (bootparams) is available. This option is required. Specifying the -p bootparams option uses the whoami call of the RPC bootparams protocol. This sets the system's hostname, domainname, and default IP router parameters. -v Enable verbose output.
示例	示例 1 Configuring Host Parameters with Verbose Output The following command configures a machine's host parameters using the whoami call of the RPC bootparams protocol with a verbose output. <pre>example# hostconfig -p bootparams -v</pre> 示例 2 Displaying Host Parameters The following command displays the parameters that would be set using the whoami call of the RPC bootparams protocol. <pre>example# hostconfig -p bootparams -n -v</pre>

示例 3 Configuring Host Parameters Less the System Name

The following command configures a machine's host parameters, less the system name, using the whoami call of the RPC bootparams protocol.

```
example# hostconfig='hostconfig -p bootparams -h'
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[hostname\(1\)](#), [domainname\(1M\)](#), [route\(1M\)](#), [attributes\(5\)](#)

引用名

hotplug – configure hotplug connectors and ports

用法概要

```
hotplug list [-c] [-d] [-l] [-v] [path [connection]]
hotplug online path port
hotplug offline [-f] [-q] path port
hotplug enable [path] connector
hotplug disable [-f] [-q] [path] connector
hotplug poweron [path] connector
hotplug poweroff [-f] [-q] [path] connector
hotplug set -o options [path] connector
hotplug get -ooptions [path] connector
hotplug install path port
hotplug uninstall [-f] [-q] path port
hotplug -?
```

描述

The `hotplug` command is used to manage hotplug connections. A connection can be a connector or port. A hotplug connector is a representation of a physical point in the system where components can be inserted or removed. A hotplug port is a representation of a logical point in the system device tree where the connection of a device to the system is managed.

The `hotplug` command only supports hotplug operations on hotplug connectors for PCI Express buses and PCI buses that implement the Standard PCI Hotplug feature. Hotplug ports on PCI Expresss and PCI buses in systems with PCI Express fabrics are also supported. Additional buses may be supported in the future.

The `hotplug` command operates on the following kinds of objects:

path

Hotplug connectors and ports are integrated into the Solaris device tree. The names of connectors and ports are unique relative only to their bus controller. A device path is required to uniquely reference a connector and the parent path of the device is required to uniquely reference a port.

The `hotplug` command can accept the case that a user specifies only a physical hotplug connector for the state-change operation subcommands. If no such connector exists, the command will fail. If multiple connectors exist with the same name in the system, the state change operations interact with the user to determine which connector needs to be operated upon. See the “Examples” section, below.

connector

If a hardware component supports being physically inserted or removed, then a hotplug connector represents the location where this action may occur. When a connector exists, it has a hierarchy of ports and device nodes that depend upon it.

port

All device nodes can be virtually hotplugged, even if their hardware does not support physical hotplugging. A hotplug port exists between a device node and its parent node in the system device tree. It represents the location where the device node and its dependents can be managed.

connection

A hotplug connection is a generic term to refer to either a hotplug connector or a hotplug port.

Hotplug connectors and ports are managed according to a state model. The `hotplug` command can list information about the hotplug connections in a system, or it can initiate change of state operations on specific hotplug connections.

Hotplug connectors can be in the following states:

empty

A component is not physically inserted in the connector.

present

A component is physically inserted in the connector, but the component is powered off. The component is not in use.

powered

A component is physically inserted in the connector, and the component is powered on. The component is disabled and is not in use.

enabled

A component is physically inserted in the connector. The component is powered on and has been probed and tested. The component is enabled and devices that represent its functions can be used.

Hotplug ports can be in the following states:

port-empty

No device exists for the hotplug port.

port-present

A device exists for the hotplug port, but the device has not been probed and it has no attached device driver. The device is not in use.

offline

A device exists for the hotplug port, and the device has been probed. A device driver is not attached, and the device is not in use.

online

A device exists for the hotplug port, and its device driver is fully attached. The device is in use.

maintenance

A device exists for the hotplug port, and its device driver is fully attached. The device is in use, but not fully operational. A maintenance or fault management operation is affecting the device. The reason that caused the device to enter maintenance state may vary. It is described by a sub-state under maintenance state. Currently there is only one sub-state defined:

maintenance-suspended

The device is live suspended.

The `hotplug` command can also access bus private properties for each hotplug connector. The current values of bus private properties can be displayed. New values for each bus private property can be set directly.

子命令

The following subcommands are supported:

list

Show information for hotplug connectors, ports, and their associated devices. Hotplug connectors and hotplug ports are integrated into the Solaris device tree hierarchy. The `list` subcommand therefore displays the hierarchy of device nodes with additional information included to show the locations of hotplug connectors and hotplug ports. The names of hotplug connectors are enclosed in square brackets, and the names of hotplug ports are enclosed in angled brackets. The current state of each hotplug connection is displayed next to its name.

online

Change the state of a hotplug port to the `online` state.

offline

Change the state of a hotplug port to the `offline` state.

enable

Change the state of a hotplug connector to the `enabled` state. All of the hotplug connector's dependent ports will be automatically probed and initialized into the `online` state.

disable

Change the state of a hotplug connector from the `enabled` state to the `powered` state. All dependent ports that are in the `online` state will first be transitioned to the `port-present` state.

poweron

Change the state of a hotplug connector from the `present` state to the `powered` state.

poweroff

Change the state of a hotplug connector from the `powered` or `enabled` state to the `present` state. All dependent ports that are in the `online` state will first be transitioned to the `port-present` state, and will then be removed.

set

Set bus-specific properties for a hotplug connector. The specified option string is a bus specific string of name and value pairs, as could be parsed by [getsubopt\(3C\)](#). The names and values will be passed directly to the bus controller that manages the specified hotplug connector to perform a bus-specific function.

get

Display the current values of bus specific properties for a hotplug connector. The specified option string is a bus specific string of named properties, as could be parsed by [getsubopt\(3C\)](#). The names will be passed directly to the bus controller to specify which properties should be returned. The current values of each named property will then be displayed.

The `install` and `uninstall` subcommands install and uninstall services the drivers of the ports can support.

install

The `install` subcommand install services the drivers of the hotplug port's device can support.

For example, this subcommand can be applied to the port of physical function of PCIe IO virtualization devices. It upgrades the port to `ONLINE` state (if it is not yet in that state) and then installs the virtual functions that the physical function (driver) supports. New hotplug ports will be created to represent each virtual function as a dependent of the specified physical function. The newly created ports will be initiated to `OFFLINE` state.

uninstall

The `uninstall` subcommand uninstall services the drivers of the hotplug port's device can support.

This subcommand can be applied to the ports of physical functions of PCIe IO virtualization devices. If the specified hotplug port has any dependent ports of virtual functions, the dependent ports and corresponding virtual function nodes will be removed.

选项

The following options are supported:

-l, --list-path

Show full paths to connections and device nodes. By default, the `list` subcommand shows hotplug connectors, ports, and devices in the format of a tree. This option enables the display of full paths to each connection and device node.

-c, --connectors

Display a table that summarizes the current status of all physical hotplug connectors. Device topologies and hierarchical information are not included. In general, the names of physical hotplug connectors should be unique. If multiple connectors of the same name exist in the system, `hotplug` displays a message specifying the same-named hotplug connectors and prompts for information to distinguish among those connectors. See “Examples,” below.

-d, --drivers

Show the binding driver names and the instance number of the device nodes. By default, the `list` subcommand shows only hotplug connectors, ports, and devices. This option enables the display of the binding driver names and the instance numbers of the device nodes.

-v, --verbose

Show verbose usage details. By default, the `list` subcommand shows only hotplug connectors, ports, and devices. This option enables the display of more detailed information about how the devices are currently consumed. Examples include mounted filesystems or plumbed network interfaces associated with individual devices.

Note that the `-v` option does not display information for disks under ZFS control.

-f, --force

Force the operation. Some change state operations that impact resources currently in use will fail with a warning. A forced operation will attempt to ignore these warnings and proceed.

This option should be used with extreme caution.

-n, --non-interactive

Disable the interactive feature. If the flag is specified and an ambiguous input is encountered, the command will exit immediately, returning a unique exit status to indicate the problem.

-q, --query

Query the operation. Instead of actually performing a change state operation, perform a test to predict if the operation would succeed or fail. If it would fail, show the error messages that would be expected if the operation had really been attempted.

It is not possible to predict every failure. An operation that succeeds during a query could still fail for another reason when actually attempted.

This option will not actually change the state of the system.

-o options, --options

Specify bus-specific properties for a `set` or `get` command. The options string conforms to the `getsubopt(3C)` syntax convention.

For the `get` subcommand, there are two special options that can be used. The special options value of `help` will display all supported properties and their possible values. The special options value of `all` will display the current value of all supported properties.

For the `set` subcommand, there is one special option that can be used. The special options value of `help` will display all supported properties which can be set and their possible values.

See “Notes” section for the properties supported by bus controllers.

-?, --help

Display a brief help message on proper use of the command.

示例

示例1 Showing All Hotplug Connections

The following command shows all hotplug connections:

```
# hotplug list -v
pci@0,0
pci108e,534a@2,1 <pci.2,1> ONLINE
    [pci30] EMPTY
pci10de,5d@e <pci.e,0> ONLINE
    display@b <pci.b,0> ONLINE
    [NEM0] ENABLED
    pci108e,534a@a,0 <pci.a,0> ONLINE
        { Network interface nge0 }
        { nge0: hosts IP addresses: 10.0.0.1 }
    pci108e,534a@a,1 <pci.a,1> (MAINTENANCE)
        [NEM1] (EMPTY)
pci108e,534a@c <pci.c,0> OFFLINE
pci108e,534a@d <pci.d,0> ONLINE
    pci1028,40d@0 <pci.0,0> (MAINTENANCE-SUSPENDED,
"activities=dma+pio+intr,reason=resource-rebalance")
        { Network interface bge0 }
        { bge0: hosts IP addresses: 10.0.1.1 }
```

To show the full paths of hotplug connections and devices, enter the following command:

```
# hotplug list -l
/pci@0,0
/pci@0,0/pci108e,4341 <pci.0,0> OFFLINE
/pci@0,0/pci8086,3408 <pci.1,0> OFFLINE
/pci@0,0/pci8086,340a@3 <pci.3,0> ONLINE
/pci@0,0/pci8086,340a@3 [Slot2] EMPTY
/pci@0,0/pci8086,340c@5 <pci.5,0> ONLINE
/pci@0,0/pci8086,340c@5 [Slot1] ENABLED
/pci@0,0/pci8086,340c@5/pci111d,8018@0 <pci.0,0> ONLINE
/pci@0,0/pci8086,340e@7 <pci.7,0> ONLINE
/pci@0,0/pci8086,340e@7 [pcie4] ENABLED
/pci@0,0/pci8086,3410@9 <pci.9,0> ONLINE
/pci@0,0/pci8086,3410@9 [pcie3] ENABLED
/pci@0,0/pci8086,342d <pci.13,0> OFFLINE
/pci@0,0/pci8086,342e <pci.14,0> OFFLINE
```

To show the binding driver names and instance numbers of the devices, enter the following command:

```
# hotplug list -d
pci@0,0 npe#0
    pci108e,4341 <pci.0,0> OFFLINE #0
```


示例1 Showing All Hotplug Connections (续)

```
pci8086,3408 <pci.1,0> OFFLINE pcieb#0
pci8086,340a@3 <pci.3,0> ONLINE pcieb#1
[Slot2] EMPTY
pci8086,340c@5 <pci.5,0> ONLINE pcieb#2
[Slot1] ENABLED
pci111d,8018@0 <pci.0,0> ONLINE pcieb#0
pci111d,8018@2 <pci.2,0> ONLINE pcieb#6
pci108e,f1bc@0 <pci.0,0> ONLINE e1000g#0
pci108e,f1bc@0,1 <pci.0,1> OFFLINE e1000g#1
pci8086,340e@7 <pci.7,0> ONLINE pcieb#3
[pcie4] ENABLED
pci8086,3410@9 <pci.9,0> ONLINE pcieb#4
[pcie3] ENABLED
pci8086,342d <pci.13,0> OFFLINE #0
pci8086,342e <pci.14,0> OFFLINE #0
```

To show the status of all physical hotplug connectors. enter the following command:

```
# hotplug list -c
```

Connection	State	Description

NEM0	ENABLED	PCIE-Native
PCI-EM0	ENABLED	PCIE-Native
NEM1	ENABLED	PCIE-Native
PCI-EM1	ENABLED	PCIE-Native

The table in the condition that multiple connectors with the same exist in the system:

```
# hotplug list -c
```

Connection	Note	State	Description

NEM0		ENABLED	PCIE-Native
PCI-EM0	*1	ENABLED	PCIE-Native
PCI-EM0	*2	ENABLED	PCIE-Native
NEM1		ENABLED	PCIE-Native

Note:

Multiple connectors with the same name exist:

```
[1] PCI-EM0 /pci@0,0/pci108e,4341@1a,2
[2] PCI-EM0 /pci@1,0/pci108e,4341@1a,2
```

示例2 Reporting Failure During State Change Operation

If a change of state operation fails, an explanation is displayed to describe the failure. An attempt to offline a hotplug port with dependent devices that are currently in use by the system might fail as follows:

示例 2 Reporting Failure During State Change Operation (续)

```
# hotplug offline /pci@0,0/pci10de,5d@e pci.a,0
ERROR: devices or resources are busy.
pci108e,534a@a,0:
  { Network interface nge0 }
  { nge0: hosts IP addresses: 10.0.0.1 }
  { Plumbed IP Address }
```

示例 3 Enabling User Interaction During State Change Operation

If there are multiple connectors with the same name in a system, the state change operations interact with the user to specify which connector needs to be operated upon.

```
# hotplug enable PCI-EM0
Multiple connectors with the same name exist:
```

```
[1] PCI-EM0 /pci@0,0/pci108e,4341@1a,2
[2] PCI-EM0 /pci@1,0/pci108e,4341@1a,2
```

Please select a connector, then press ENTER:

示例 4 Displaying Bus-Specific Properties and Values

The following command displays all supported bus-specific properties and their possible values:

```
# hotplug get -o help /pci@0,0 pci.2,1
power_led=<on|off|blink>
fault_led=<on|off|blink>
active_led=<on|off|blink>
attn_led=<on|off|blink>
card_type=<type description>
board_type=<type description>
```

示例 5 Displaying Bus-Specific Options

The following command displays the card type and the current state of the Power LED of a PCI hotplug connector:

```
# hotplug get -o card_type,power_led /pci@0,0 pci.2,1
card_type=fibre
power_led=on
```

示例 6 Setting a Bus-Specific Property

The following command turns on the attention LED of a PCI hotplug connector:

```
# hotplug set -o attn_led=on /pci@0,0 pci.2,1
```

示例 7 Installing Port Dependents

The following commands install dependent ports of an IOV physical function and then display the resulting IOV virtual functions that were probed.

```
# hotplug install /pci@400/pci@1/pci@0/pci@4 pci.0,1

# hotplug list -v /pci@400/pci@1/pci@0/pci@4 pci.0,1
<pci.0,1> (ONLINE)
    { IOV physical function }
    { IOV virtual function 'pci.0,81' }
    { IOV virtual function 'pci.0,83' }
    { IOV virtual function 'pci.0,85' }
    { IOV virtual function 'pci.0,87' }
<pci.0,81> (OFFLINE)
ethernet@0,81
<pci.0,83> (OFFLINE)
ethernet@0,83
<pci.0,85> (OFFLINE)
ethernet@0,85
<pci.0,87> (OFFLINE)
ethernet@0,87
```

示例 8 Uninstalling Port Dependents

The following command attempts to uninstall dependent ports of an IOV physical function, but fails when a dependent IOV virtual function is busy.

```
# hotplug uninstall /pci@400/pci@1/pci@0/pci@4 pci.0,0
ERROR: devices or resources are busy.
ethernet@0,81:
    { Network interface igvbf1 }
    { igvbf1: hosts IP addresses: 10.0.0.1 }
    { Plumbed IP Address }
```

退出状态

- 0 Successful completion.
- 1 Invalid command line options were specified.
- 2 The specified path or connection does not exist.
- 3 A fatal error occurred. One or more error messages are displayed on standard error.
- 4 The hotplug service is not available.
- 5 The connector name specified refers to multiple ports.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [cfgadm\(1M\)](#), [hotplugd\(1M\)](#), [getsubopt\(3C\)](#), [attributes\(5\)](#)

诊断 The following error message is displayed on systems that do not have any supported IO buses:

ERROR: there are no connections to display.
(See [hotplug\(1m\)](#) for more information.)

If this error message is seen, note that the system might still have other IO devices that support hotplugging, through the [cfgadm\(1M\)](#) command instead of [hotplug](#).

附注 The [hotplug](#) service (FMRI `svc:/system/hotplug`) must be enabled as a prerequisite for using the [hotplug](#) command. The service is disabled by default. See [hotplugd\(1M\)](#).

The authorization `solaris.hotplug.modify` must be granted in order to perform change-of-state operations, or to install and uninstall dependent ports. Alternatively, the rights profile “Hotplug Management” can be granted, which includes that authorization.

Verbose usage information is gathered from the RCM framework. Its format and content is subject to change.

The following bus specific properties are supported in PCI bus controllers:

`power_led | fault_led | attn_led | active_led`
States of a specific LED of a slot. The value could be on, off, or blink.

They can all be used with `get` subcommand, but only property `attn_led` can be used with `set` subcommand.

`card_type | board_type`
Type of a card or board of a slot.

They can all be used with `get` subcommand, but neither can be used with `set` subcommand.

引用名	hotplugd – hotplug daemon
用法概要	<code>/usr/lib/hotplugd [-d]</code>
描述	The hotplug daemon, <code>hotplugd</code>, provides user-level services for the management of hotplug connections. It is a system daemon started by the Service Management Facility (see smf(5)). Its fault management resource identifier (FMRI) is: <pre>svc:/system/hotplug:default</pre> Note that <code>hotplugd</code> is a Consolidation Private interface. See attributes(5). The hotplug(1M) command and any other client program that uses the private <code>libhotplug</code> library to query information about hotplug connections or initiate hotplug commands depends upon this daemon. The hotplug daemon is a door server which services requests from all <code>libhotplug</code> clients. The door interface is private. Client applications use the private <code>libhotplug</code> interface to administer hotplug connections. <code>libhotplug</code> uses the door interface to administer hotplug connections through the hotplug daemon service. The hotplug daemon acts as a central location to serialize all hotplug operations and coordinate activities with all other parts of the system.
选项	The following option is supported: <code>-d, --debug</code> Run the daemon in standalone debug mode. Messages will be displayed on the controlling terminal instead of to <code>syslog</code>. And increased verbosity will be enabled to display more details about the internal operations of the daemon.
示例	示例 1 Enabling the Hotplug Service The following command enables the hotplug service: <pre># svcadm enable svc:/system/hotplug:default</pre> 示例 2 Disabling the Hotplug Service The following command disables the hotplug service: <pre># svcadm disable svc:/system/hotplug:default</pre>
Errors	The hotplug daemon uses syslog(3C) to report status and error messages. All of the messages are logged with the <code>LOG_DAEMON</code> facility. Error messages are logged with the <code>LOG_ERR</code> and <code>LOG_NOTICE</code> priorities, and informational messages are logged with the <code>LOG_INFO</code> priority. The default entries in the <code>/etc/syslog.conf</code> file log all of the hotplug daemon error messages to the <code>/var/adm/messages</code> log.
文件	<pre>/var/run/hotplugd_door Hotplug daemon door /var/run/hotplugd_pid Hotplug daemon lock file</pre>

`/usr/lib/hotplugd`
Hotplug daemon binary

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Consolidation Private

另请参见 [svcs\(1\)](#), [hotplug\(1M\)](#), [svcadm\(1M\)](#), [syslog\(3C\)](#), [syslog.conf\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The `hotplugd` service is managed by the service management facility, [smf\(5\)](#), under the fault management resource identifier (FMRI):

`svc:/system/hotplug:default`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command. To administer the service, the authorization `solaris.smf.manage.hotplug` must be granted. Alternatively, the rights profile “Hotplug Management” can be granted.

The `hotplug` service must be enabled for the [hotplug\(1M\)](#) command and any other `libhotplug` client applications to function properly.

引用名	htable – convert DoD Internet format host table				
用法概要	<code>/usr/sbin/htable filename</code>				
描述	htable converts a host table in the format specified by RFC 952 to the format used by the network library routines. Three files are created as a result of running htable: hosts, networks, and gateways. The hosts file is used by the gethostbyname(3NSL) routines in mapping host names to addresses. The networks file is used by the getnetbyname(3SOCKET) routines in mapping network names to numbers. The gateways file is used by the routing daemon to identify “passive” Internet gateways. If any of the files <code>localhosts</code>, <code>localnetworks</code>, or <code>localgateways</code> are present in the current directory, the file's contents is prepended to the output file without interpretation. This allows sites to maintain local aliases and entries which are not normally present in the master database. htable is best used in conjunction with the gettable(1M) program which retrieves the DoD Internet host table from a host.				
文件	<code>localhosts</code> <code>localnetworks</code> <code>localgateways</code>				
属性	See attributes(5) for descriptions of the following attributes:				
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>system/network/nis</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	system/network/nis
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	system/network/nis				
另请参见	gettable(1M), gethostbyname(3NSL), getnetbyname(3SOCKET), attributes(5) Harrenstien, Ken, Mary Stahl, and Elizabeth Feinler, DoD Internet Host Table Specification, RFC 952, Network Information Center, SRI International, Menlo Park, California, October 1985.				
附注	htable does not properly calculate the gateways file.				

引用名	ickey – install a client key for WAN boot
用法概要	<pre>/usr/lib/inet/wanboot/ickey [-d] [-o type=3des] /usr/lib/inet/wanboot/ickey [-d] [-o type=aes] /usr/lib/inet/wanboot/ickey [-d] [-o type=sha1]</pre>
描述	The ickey command is used to install WAN boot keys on a running UNIX system so that they can be used the next time the system is installed. You can store three different types of keys: 3DES and AES for encryption and an HMAC SHA-1 key for hashed verification. ickey reads the key from standard input using getpass(3C) so that it does not appear on the command line. When installing keys on a remote system, you must take proper precautions to ensure that any keying materials are kept confidential. At a minimum, use ssh(1) to prevent interception of data in transit. Keys are expected to be presented as strings of hexadecimal digits; they can (but need not) be preceded by a 0x or 0X. The ickey command has a single option, described below. An argument of the type -o type=keytype is required.
选项	The ickey command the following option. -d Delete the key specified by the <i>keytype</i> argument.
退出状态	On success, ickey exits with status 0; if a problem occurs, a diagnostic message is printed and ickey exits with non-zero status.
文件	/dev/openprom WAN boot key storage driver
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/boot/wanboot
Interface Stability	Uncommitted

另请参见 [ssh\(1\)](#), [openprom\(7D\)](#), [attributes\(5\)](#)

引用名	id – return user identity
用法概要	<pre> /usr/bin/id [-p] [user] /usr/bin/id -a [-p] [user] /usr/bin/id -G [-n] [user] /usr/bin/id -g [-nr] [user] /usr/bin/id -u [-nr] [user] /usr/xpg4/bin/id [-p] [user] /usr/xpg4/bin/id -a [-p] [user] /usr/xpg4/bin/id -G [-n] [user] /usr/xpg4/bin/id -g [-nr] [user] /usr/xpg4/bin/id -u [-nr] [user] </pre>
描述	If no <i>user</i> operand is provided, the <i>id</i> utility writes the user and group IDs and the corresponding user and group names of the invoking process to standard output. If the effective and real IDs do not match, both are written. If multiple groups are supported by the underlying system, <i>/usr/xpg4/bin/id</i> also writes the supplementary group affiliations of the invoking process. If a <i>user</i> operand is provided and the process has the appropriate privileges, the user and group IDs of the selected user are written. In this case, effective IDs are assumed to be identical to real IDs. If the selected user has more than one allowable group membership listed in the group database, <i>/usr/xpg4/bin/id</i> writes them in the same manner as the supplementary groups described in the preceding paragraph.
Formats	The following formats are used when the LC_MESSAGES locale category specifies the "C" locale. In other locales, the strings <i>uid</i>, <i>gid</i>, <i>euid</i>, <i>egid</i>, and <i>groups</i> may be replaced with more appropriate strings corresponding to the locale. <pre> "uid=%u(%s) gid=%u(%s)\n" <real user ID>, <user-name>, <real group ID>, <group-name> </pre> If the effective and real user IDs do not match, the following are inserted immediately before the <i>\n</i> character in the previous format: <pre> " euid=%u(%s)" </pre> with the following arguments added at the end of the argument list: <pre> <effective user ID>, <effective user-name> </pre> If the effective and real group IDs do not match, the following is inserted directly before the <i>\n</i> character in the format string (and after any addition resulting from the effective and real user IDs not matching): <pre> " egid=%u(%s)" </pre>

with the following arguments added at the end of the argument list:

<effectivegroup-ID>, <effectivegroupname>

If the process has supplementary group affiliations or the selected user is allowed to belong to multiple groups, the first is added directly before the NEWLINE character in the format string:

" groups=%u(%s) "

with the following arguments added at the end of the argument list:

<supplementary group ID>, <supplementary group name>

and the necessary number of the following added after that for any remaining supplementary group IDs:

", %u(%s) "

and the necessary number of the following arguments added at the end of the argument list:

<supplementary group ID>, <supplementary group name>

If any of the user ID, group ID, effective user ID, effective group ID or supplementary/multiple group IDs cannot be mapped by the system into printable user or group names, the corresponding (%s) and name argument is omitted from the corresponding format string.

When any of the options are specified, the output format is as described under OPTIONS.

选项

The following options are supported by both `/usr/bin/id` and `/usr/xpg4/bin/id`. The `-p` and `-a` options are invalid if specified with any of the `-G`, `-g`, or `-u` options.

- p Reports additionally the current project membership of the invoking process. The project is reported using the format:

"projid=%u(%s) "

which is inserted prior to the `\n` character of the default format described in the `Formats` section. The arguments

<project ID>, <project name>

are appended to the end of the argument list. If the project ID cannot be mapped by the system into a printable project name, the corresponding (%s) and name argument is omitted from the corresponding format string.

- a Reports user name, user ID and all the groups to which the user belongs.
- G Outputs all different group IDs (effective, real and supplementary) only, using the format "%u\n". If there is more than one distinct group affiliation, output each such affiliation, using the format " %u", before the NEWLINE character is output.
- g Outputs only the effective group ID, using the format "%u\n".
- n Outputs the name in the format "%s" instead of the numeric ID using the format "%u".

- r Outputs the real ID instead of the effective ID.
- u Outputs only the effective user ID, using the format "%u\n".

操作数 The following operand is supported:

user The user (login) name for which information is to be written.

环境变量 See [environ\(5\)](#) for descriptions of the following environment variables that affect the execution of `id`: `LANG`, `LC_ALL`, `LC_CTYPE`, `LC_MESSAGES`, and `NLSPATH`.

退出状态 The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

/usr/bin/id

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os, system/library/processor
Interface Stability	Committed

/usr/xpg4/bin/id

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/xopen/xcu4
Interface Stability	Standard

另请参见 [fold\(1\)](#), [logname\(1\)](#), [who\(1\)](#), [getgid\(2\)](#), [getgroups\(2\)](#), [getprojid\(2\)](#), [getuid\(2\)](#), [attributes\(5\)](#), [environ\(5\)](#), [standards\(5\)](#)

附注 Output produced by the `-G` option and by the default case could potentially produce very long lines on systems that support large numbers of supplementary groups.

引用名 idmap – 配置和管理本机身份映射服务

用法概要

```
idmap
idmap -f command-file
idmap add [-d] name1 name2
idmap dump [-n] [-v]
idmap export [-f file] format
idmap flush [-a]
idmap get-namemap name
idmap help
idmap import [-F] [-f file] format
idmap list
idmap remove [-t|-f] name
idmap remove -a
idmap remove [-d] name1 name2
idmap set-namemap [-a authenticationMethod] [-D bindDN]
    [-j passwdfile] name1 name2
idmap show [-c] [-v] [-V] identity [target-type]
idmap unset-namemap [-a authenticationMethod] [-D bindDN]
    [-j passwdfile] name [target-type]
```

描述 idmap 实用程序用于配置和管理本机身份映射服务。

本机身份映射服务支持 Windows 安全标识符 (security identifier, SID) 与 POSIX 用户 ID 和组 ID (UID 和 GID) 之间的以下映射类型：

- **基于名称的映射。**管理员按名称映射 Windows 和 UNIX 的用户和组。
- **临时 ID 映射。**为每个尚未按名称映射的 SID 动态分配 UID 或 GID。
- **本地 SID 映射。**非临时 UID 或 GID 映射到按算法生成的本地 SID。

idmap 实用程序可以用于创建和管理基于名称的映射，以及监视生效的映射。

如果调用 idmap 实用程序时不带子命令或选项，则会从标准输入读取子命令。标准输入是 TTY 时，idmap 命令会输出用法消息，然后退出。

映射机制 [idmapd\(1M\)](#) 守护进程可将 Windows 用户和组 SID 映射到 UNIX UID 和 GID，如下所示：

1. SID 按名称映射。
此映射使用系统管理员手动设置的基于名称的映射。
2. 如果找不到基于名称的映射，SID 会映射到动态分配的临时 ID。

此分配使用 2^{31} 到 $2^{32} - 2$ 之间的下一个可用 UID 或 GID。

本地 SID 映射用于从 UNIX 到 Windows 的映射。

为了防止别名问题，所有文件系统、归档和备份格式以及协议都必须存储 SID，或将 2^{31} 到 $2^{32} - 2$ 范围内的所有 UID 和 GID 映射到 `nobody` 用户和组。

也可以创建对角映射。对角映射是指 Windows 组和 Solaris 用户之间以及 Solaris 组和 Windows 用户之间的映射。在 Windows 使用组身份作为文件所有者（或反之）的时候，需要这种映射。

基于名称的映射

基于名称的映射可以在 Windows 用户和组与 UNIX 名称服务中的用户和组之间建立等效名称映射关系。这些映射在重新引导后依然存在。例如，以下命令将 Windows 用户映射到具有相同名称的 UNIX 用户：

```
# idmap add "winuser:*@mywindomain.com" "unixuser:*
```

如果配置为使用目录服务，`idmapd(1M)` 将首先尝试使用存储在 Active Directory (AD) 和/或本机 LDAP 目录服务的用户或组对象中的映射信息。例如，可以对给定 Windows 用户或组的 AD 对象进行扩充，使之包括相应的 Solaris 用户或组名称或者数值型 ID。同样，也可以对给定 Solaris 用户或组的本机 LDAP 对象进行扩充，使之包括相应的 Windows 用户或组名称。

通过设置 `idmap` 服务的相应服务管理工具 (service management facility, SMF) 属性，可以将 `idmapd(1M)` 配置为使用 AD 和/或本机 LDAP 目录型名称映射。请参见下面的“服务属性”部分了解更多详细信息。

如果没有配置目录型名称映射或是配置了却找不到，则 `idmapd(1M)` 将处理本地存储的基于名称的映射规则。

`idmap` 支持对 Windows 已知名称的映射。下面列出其中几个：

```
Administrator
Guest
KRBtgt
Domain Admins
Domain Users
Domain Guest
Domain Computers
Domain Controllers
```

添加 `idmap` 规则时，这些已知名称会扩展成规范格式。也就是说，将会添加缺省域名（对于非已知名称），或者添加适当的内置域名。此域名可以为空或 `BUILTIN`，也可以是本地主机名，具体取决于特定的已知名称。

以下 `idmap` 命令序列说明了非已知名称 `fred`、已知名称 `administrator` 和 `guest` 的处理方法。

```
# idmap add winname:fred unixuser:fredf
add      winname:fred      unixuser:fredf
```

```
# idmap add winname:administrator unixuser:root
add      winname:administrator  unixuser:root

# idmap add winname:guest unixuser:nobody
add      winname:guest          unixuser:nobody

# idmap add wingroup:administrators sysadmin
add      wingroup:administrators unixgroup:sysadmin

# idmap list
add      winname:Administrator@examplehost unixuser:root
add      winname:Guest@examplehost        unixuser:nobody
add      wingroup:Administrators@BUILTIN   unixgroup:sysadmin
add      winname:fred@example.com          unixuser:fredf
```

临时映射

idmapd 守护进程会尝试在守护进程重新启动后保留临时 ID 映射。但是如果无法保留 ID，守护进程会将之前映射的每个 SID 映射到新的临时 UID 或 GID 值。守护进程决不会重复使用临时 UID 或 GID。如果 idmapd 守护进程用光了临时 UID 和 GID，它会返回错误，同时为无法按名称映射的 SID 返回缺省 UID 或 GID。

重新引导后不会保留动态 ID 映射。所以，重新引导系统后，任何动态映射到 UNIX UID 或 GID 的 SID 都很可能会映射到不同的 ID。

本地 SID 映射

如果找不到基于名称的映射，非临时 UID 或 GID 将映射到按算法生成的本地 SID。映射生成方式如下：

```
local SID for UID = <machine SID> - <1000 + UID>
local SID for GID = <machine SID> - <2^31 + GID>
```

<machine SID> 是由 idmap 服务为运行时所在的主机生成的唯一 SID。

规则查找顺序

将 Windows 名称映射到 UNIX 名称时，按以下顺序查找基于名称的映射规则：

1. *windows-name@domain* 到 ""
2. *windows-name@domain* 到 *unix-name*
3. *windows-name@** 到 ""
4. *windows-name@** 到 *unix-name*
5. **@domain* 到 *
6. **@domain* 到 ""
7. **@domain* 到 *unix-name*
8. **@** 到 *
9. **@** 到 ""
10. **@** 到 *unix-name*

将 UNIX 名称映射到 Windows 名称时，按以下顺序查找基于名称的映射规则：

1. *unix-name* 到 ""
2. *unix-name* 到 *windows-name@domain*
3. *** 到 **@domain*

4. * 到 ""
5. * 到 *windows-name@domain*

服务属性

服务属性决定 [idmapd\(1M\)](#) 守护进程的行为。这些属性存储在属性组 `config` 下的 SMF 系统信息库中（请参见 [smf\(5\)](#)）。您可以使用 [svccfg\(1M\)](#) 访问和修改它们，但这需要 `solaris.smf.value.idmap` 授权。idmap 服务的服务属性如下：

`config/ad_unixuser_attr`

指定包含 UNIX 用户名的 AD 属性名。无缺省值。

`config/ad_unixgroup_attr`

指定包含 UNIX 组名称的 AD 属性名。无缺省值。

`config/nldap_winname_attr`

指定包含 Windows 用户/组名称的本机 LDAP 属性名。无缺省值。

`config/directory_based_mapping`

控制对使用存储在目录服务中的数据进行身份映射的支持。

`none` 禁用基于目录的映射。

`name` 允许使用上述属性进行基于名称的映射。

`idmu` 启用使用 Microsoft Identity Management for UNIX (IDMU) 的映射。此 Windows 组件允许管理员为每个 Windows 用户指定一个 UNIX 用户 ID，以便将 Windows 身份映射到相应的 UNIX 身份。仅使用包含 Solaris 系统的域中的 IDMU 数据。

更改服务属性不会影响正在运行的 idmap 服务。更改在服务刷新（使用 [svcadm\(1M\)](#)）后才会生效。

操作数

idmap 命令使用以下操作数：

格式化

指定用户名称映射在 `export` 和 `import` 子命令中的描述格式。支持 `Netapp usermap.cfg` 和 `Samba smbusers` 外部格式。这些外部格式仅适用于用户而非组。

- `usermap.cfg` 规则-映射格式如下：

windows-username [*direction*] *unix-username*

windows-username 是 Windows 用户名，格式为 *domain\username* 或 *username@domain*。

unix-username 是 UNIX 用户名。

direction 为以下值之一：

- `==` 表示这是一个双向映射（此为缺省值）。
- `=>` 或 `<=` 表示这是一个单向映射。

不支持 IP 限定符。

- `smbusers` 规则-映射格式如下：

```
unixname = winname1 winname2 ...
```

如果 *winname* 包括空格，则通过用双引号将值括起来对空格进行转义。例如，以下文件显示如何以 *idmap* 命令的有效格式指定空格：

```
$ cat myusermap
terry="Terry Maddox"
pat="Pat Flynn"
cal=cbrown
```

映射作为从 Windows 名称到 UNIX 名称的单向映射导入。

格式依照 *smb.conf* 手册页的“用户名映射”条目的说明，可以在 samba.org Web 站点上找到这部分内容。支持在 *windows-name* 中使用星号(*)，但是不支持 *@group* 指令和映射链接。

缺省情况下，如果 *smbusers* 文件中没有映射项，Samba 会将 *windows-name* 映射到等效 *unix-name*（如果有）。如果要像 Samba 一样设置相同的映射，请使用以下 *idmap* 命令：

```
idmap add -d "winuser:*@" "unixuser:*
```

identity

指定用户名、用户 ID、组名称或组 ID。*identity* 以 *type:value* 格式指定。*type* 为以下值之一：

<i>usid</i>	Windows 用户 SID（文本格式）
<i>gsid</i>	Windows 组 SID（文本格式）
<i>sid</i>	既可以属于用户也可以属于组的 Windows 组 SID（文本格式）
<i>uid</i>	数值型 POSIX UID
<i>gid</i>	数值型 POSIX GID
<i>unixuser</i>	UNIX 用户名
<i>unixgroup</i>	UNIX 组名称
<i>winuser</i>	Windows 用户名
<i>wingroup</i>	Windows 组名称
<i>winname</i>	Windows 用户名或组名称

value 是一个对应于指定 *type* 的数字或字符串。例如，*unixgroup:staff* 指定 UNIX 组名称 *staff*。标识 *gid:10* 代表 GID 10，它与 UNIX 组 *staff* 相对应。

name

指定一个可以用于基于名称的映射规则的 UNIX 名称（*unixuser*、*unixgroup*）或 Windows 名称（*winuser*、*wingroup*）。

Windows 安全身份名称可以通过下列方式之一指定：

- *domain\name*
- *name@domain*
- *name*，使用缺省映射域

如果 *name* 是空字符串 ("")，则会禁止映射。请注意，"" 名称不能用于阻止未映射 Windows 用户的登录。

如果 *name* 使用了通配符 (*)，它将匹配未被其他映射匹配的所有名称。同样，如果 *name* 是通配符 Windows 名称 (*@*)，它将匹配未被其他映射匹配的所有域中的所有名称。

如果 *name* 在映射规则两端都使用通配符，那么 Windows 和 Solaris 用户的名称是相同的。例如，如果规则是 "*@domain" == "*", 那么 jp@domain Windows 用户名与该规则匹配，并会映射到 jp Solaris 用户名。

如果类型可以从命令行上指定的其他参数或类型中推断，您也可以选择是否指定 *name* 的类型。

target-type

与 *show* 和 *unset-namemap* 子命令结合使用。对于 *show*，指定应该显示的映射类型。例如，如果 *target-type* 是 *sid*，*idmap show* 会返回 SID，该 SID 映射到在命令行上指定的身份。对于 *unset-namemap*，确定由 *name* 操作数指定的对象中的属性。

选项

idmap 命令支持一个选项和一组子命令。子命令也有选项。

命令行选项

-f command-file

从 *command-file* 读取并执行 *idmap* 子命令。*idmap -f* - 命令从标准输入读取。此选项不能由任何子命令使用。

子命令

支持以下子命令：

add [-d] name1 name2

添加基于名称的映射规则。缺省情况下，名称映射是双向的。如果使用 *-d* 选项，会创建从 *name1* 到 *name2* 的单向映射。

name1 或 *name2* 的其中一个必须是 Windows 名称，而另外一个必须是 UNIX 名称。对于 Windows 名称，不能使用 *winname* 身份类型，而是要从 *winuser* 或 *wingroup* 类型中指定一个。有关 *name* 操作数的信息，请参见“操作数”部分。

请注意，相同的两个名称之间的两个方向相反的单向映射相当于一个双向映射。

该子命令需要 *solaris.admin.idmap.rules* 授权。

dump [-n] [-v]

转储自上次系统引导后缓存的所有映射。*-n* 选项也会显示名称。缺省情况下，只会显示 *sid*、*uid* 和 *gid*。*-v* 选项显示映射是如何生成的。

export [-f *file*] *format*

将基于名称的映射规则以指定的 *format* 导出至标准输出。-f *file* 选项将规则写入指定的输出文件。

flush [-a]

刷新身份映射高速缓存，以便将来的映射请求会完全按照当前的规则和目录信息进行处理。这是一个非中断操作。更改规则时会自动刷新高速缓存；该手动操作可用于强制使新更改的目录信息生效。

使用 -a 选项可将高速缓存清除干净。该操作可能会中断进行中的操作，因此应仅在静默系统上使用。可能需要使用 -a 选项来设置 "clean slate" 测试用例，虽然一般情况下并不需要。

get-namemap *name*

从指定名称的 AD 或本机 LDAP 用户或组对象获取目录型名称映射信息。

help

显示用法消息。

import [-F] [-f *file*] *format*

使用指定的 *format* 从标准输入导入基于名称的映射规则。-f *file* 选项从指定文件读取规则。-F 选项在添加新规则前刷新现有的基于名称的映射规则。

不管使用的外部格式为何，导入的规则将使用前面“规则查找顺序”部分描述的语义和顺序进行处理。

该子命令需要 `solaris.admin.idmap.rules` 授权。

list

列出所有基于名称的映射规则。所有规则都以 `idmap add` 格式显示。

remove [-t|-f] *name*

删除涉及指定名称的任何基于名称的映射规则。*name* 可以是 UNIX 或 Windows 用户名或组名称。

-f 选项删除使用 *name* 作为源的规则。-t 选项删除使用 *name* 作为目标的规则。这些选项是互斥的。

该子命令需要 `solaris.admin.idmap.rules` 授权。

remove -a

删除所有基于名称的映射规则。

该子命令需要 `solaris.admin.idmap.rules` 授权。

remove [-d] *name1 name2*

删除 *name1* 和 *name2* 之间基于名称的映射规则。如果指定 -d 选项，则会删除 *name1* 到 *name2* 的规则。

name1 或 *name2* 的其中一个必须是 Windows 名称，而另外一个必须是 UNIX 名称。

该子命令需要 `solaris.admin.idmap.rules` 授权。

`set-namemap [-a authenticationMethod] [-D bindDN] [-j passwdfile] name1 name2`

设置 AD 或本机 LDAP 用户或组对象中的名称映射信息。*name1* 或 *name2* 的其中一个必须是 Windows 名称，而另外一个必须是 UNIX 名称。

如果 *name1* 是一个 Windows 名称，那么 UNIX 名称 *name2* 会添加到 *name1* 所代表的 AD 对象。同样，如果 *name1* 是一个 UNIX 名称，那么 Windows 名称 *name2* 会添加到 *name1* 所代表的本机 LDAP 项。

所支持的选项如下：

`-a authenticationMethod`

指定修改本机 LDAP 项时的验证方法。有关详细信息，请参见 [ldapaddent\(1M\)](#)。缺省值是 `sasl/GSSAPI`。

`-D bindDN`

使用标识名 *bindDN* 绑定到目录。

`-j passwdfile`

指定一个文件，该文件包含对目录进行身份验证所需的口令。

`show [-c] [-v] [-V] name [target-type]`

显示类型为 *target-type* 的身份，它是指定 *name* 映射的对象。如果省略可选的 *target-type*，则会显示非对角映射。

缺省情况下，该子命令只显示已经建立的映射。`-c` 选项强制评估基于名称的映射配置或 ID 的动态分配。

`-v` 选项显示映射是如何生成的，以及映射是刚刚生成的还是从高速缓存获取的。`-v` 选项详细说明了要确定映射所需采取的具体步骤，包括被拒绝的临时步骤和方法。

`unset-namemap [-a authenticationMethod] [-D bindDN] [-j passwdfile] name [target-type]`

从指定名称和可选目标类型代表的 AD 或本机 LDAP 用户或组对象中取消设置目录名称映射信息。

请参见 `set-namemap` 子命令以了解各个选项。

示例

示例1 在基于名称的映射规则两端使用通配符

如果有一项存在且没有其他映射，以下命令可将 `xyz.com` 域中的所有 Windows 用户名映射到具有相同名称的 UNIX 用户。如果匹配此类规则，但是 UNIX 用户名不存在，则会使用临时 ID 映射。

```
# idmap add "winuser:*@xyz.com" "unixuser:*
```

示例2 在基于名称的映射规则一端使用通配符

以下命令将 `xyz.com` 域中的所有未映射 Windows 用户映射到 `guest` UNIX 用户。`-d` 选项指定一个从 `*@xyz.com` 用户到 `guest` 用户的单向映射。

示例2 在基于名称的映射规则一端使用通配符 (续)

```
# idmap add -d "winuser:*@xyz.com" unixuser:guest
```

示例3 添加双向的基于名称的映射规则

以下命令将 Windows 用户 foobar@example.com 映射到 UNIX 用户 foo (反之亦可) :

```
# idmap add winuser:foobar@example.com unixuser:foo
```

此命令显示如何删除由先前命令添加的映射 :

```
# idmap remove winuser:foobar@example.com unixuser:foo
```

示例4 显示 UID 到 SID 的映射

- 以下命令显示指定 UID uid:50000 映射到的 SID :

```
# idmap show uid:50000 sid
uid:50000 -> usid:S-1-5-21-3223191800-2000
```

- 以下命令显示指定 Windows 用户名 joe@example.com 映射到的 UNIX 用户名 :

```
# idmap show joe@example.com unixuser
winuser:joe@example.com -> unixuser:joes
```

示例5 列出缓存的 SID 到 UID 的映射

以下命令显示高速缓存中所有的 SID 到 UID 的映射 :

```
# idmap dump | grep "uid:"
usid:S-1-5-21-3223191800-2000 == uid:50000
usid:S-1-5-21-3223191800-2001 == uid:50001
usid:S-1-5-21-3223191800-2006 == uid:50010
usid:S-1-5-21-3223191900-3000 == uid:2147491840
usid:S-1-5-21-3223191700-4000 => uid:60001
```

示例6 批处理 idmap 请求

以下命令显示如何批处理 idmap 请求。此特定命令序列执行如下操作 :

- 删除 foobar@example.com 此前的所有规则。
- 将 Windows 用户 foobar@example.com 映射到 UNIX 用户 bar, 反之亦可。
- 将 Windows 组 members 映射到 UNIX 组 staff, 反之亦可。

```
# idmap <<EOF
    remove winuser:foobar@example.com
    add winuser:foobar@example.com unixuser:bar
    add wingroup:members unixgroup:staff
EOF
```

示例7 列出基于名称的映射规则

以下命令显示如何列出基于名称的映射规则 :

示例7 列出基于名称的映射规则 (续)

```
# idmap list
add winuser:foobar@example.com unixuser:bar
add wingroup:members unixgroup:staff
```

示例8 从 usermap.cfg 文件导入基于名称的映射规则

usermap.cfg 文件可以用于配置基于名称的映射规则。以下 usermap.cfg 文件显示了从 Windows 用户 foo@example.com 到 UNIX 用户 foo，以及从 foobar@example.com 到 UNIX 用户 foo 的映射规则。

```
# cat usermap.cfg
foo@example.com == foo
foobar@example.com => foo
```

以下 idmap 命令将 usermap.cfg 信息导入到 idmapd 数据库：

```
# cat usermap.cfg | idmap import usermap.cfg
```

此命令与先前的命令功能相同：

```
# idmap import -f usermap.cfg usermap.cfg
```

以下命令相当于先前的 idmap import 命令：

```
# idmap <<EOF
    add winuser:foo@example.com unixuser:foo
    add -d winuser:foobar@example.com unixuser:foo
EOF
```

示例9 结合使用基于名称的和临时的 ID 映射与身份功能映射及例外情况

以下命令将 example.com Windows 域中的所有用户映射到具有相同名称的 UNIX 用户帐户。命令还会指定以下 Windows 用户的映射：

joe@example.com、jane.doe@example.com 和 administrator@example.com。所有域中的 administrator 都映射到 nobody。没有对应 UNIX 帐户的 Windows 用户会动态映射到可用的临时 UID。

```
# idmap import usermap.cfg <<EOF
    joe@example.com == joes
    jane.doe@example.com == janed
    administrator@* => nobody
    *@example.com == *
    *@example.com => nobody
EOF
```

示例10 将目录型名称映射添加到 AD 用户对象

以下命令可将 UNIX 名称添加到 joe@example.com 的 AD 对象，从而将 Windows 用户 joe@example.com 映射到 UNIX 用户 joe。

示例 10 将目录型名称映射添加到 AD 用户对象 (续)

```
# idmap set-namemap winuser:joe@example.com joes
```

示例 11 将目录型名称映射添加到本机 LDAP 用户对象
以下命令可将 Windows 名称添加到 foo 的本机 LDAP 对象，从而将 UNIX 用户 foo 映射到 Windows 用户 foobar@example.com。

```
# idmap set-namemap unixuser:foo foobar@example.com
```

示例 12 从 AD 用户对象中删除目录型名称映射
以下命令从代表 joe@example.com 的 AD 对象中删除 UNIX 用户名 unixuser。

```
# idmap unset-namemap winuser:joe@example.com unixuser
```

退出状态

0 成功完成。
>0 出现错误。向标准错误写入一条诊断消息。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Uncommitted（未确定）

另请参见

[svcs\(1\)](#)、[idmapd\(1M\)](#)、[ldapaddent\(1M\)](#)、[svcadm\(1M\)](#)、[svccfg\(1M\)](#)、[ad\(5\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

附注

[idmapd](#) 服务由服务管理工具 [smf\(5\)](#) 管理。[idmapd](#) 服务的服务标识符是 `svc:/system/idmap`。

使用 `svcadm` 命令执行此服务的管理操作，如启用、禁用或重新启动服务。这些操作需要 `solaris.smf.manage.idmap` 授权。使用 `svcs` 命令查询服务状态。

Windows 用户名不区分大小写，而 UNIX 用户名则区分大小写。[idmap](#) 名称-规则和 `idmap show` 命令行中显示的 Windows 名称的大小写是不相关的。

由于在 UNIX 环境下的通常做法是使用全部小写的用户名，因此通配符名称-规则会按照如下方式将 Windows 名称映射到 UNIX 用户/组名称：首先，使用规范的 Windows 名称（即大小写情况与其在目录中的显示相同）作为 UNIX 用户名或组名称。如果没有这样的 UNIX 身份，则 Windows 名称转换为小写，这样就可以用作 UNIX 用户名或组名称。

这样区分对待大小写的结果是，看起来相同的用户名可能不会被识别为相互匹配。您必须创建规则来处理这种只有大小写不同的字符串的配对情况。例如，要将 Windows 用户 `sam@example` 映射到 Solaris 用户 `Sam`，必须创建以下规则：

```
# idmap add "winuser:*@example" "unixuser:*"  
# idmap add winuser:sam@example unixuser:Sam
```

有关修改 Active Directory 模式的指导信息，请参考 Microsoft 文档《Step-by-Step Guide to Using Active Directory Schema and Display Specifiers》，该文档可以在 Microsoft technet Web 站点 <http://technet.microsoft.com/> 上找到。

引用名	idmapd – 本机身份映射服务守护进程
用法概要	<code>/usr/lib/idmapd</code>
描述	<code>idmapd</code> 守护进程可以将 Windows 安全标识符 (Security Identifier, SID) 映射到 POSIX 标识符 (UID/GID)，或者进行反向映射。 idmap(1M) 实用程序提供了 <code>idmapd</code> 守护进程的前端。
文件	<code>/var/idmap/idmap.db</code> 用于存储本地的基于名称的 ID 映射规则的数据库。该数据库的内容是专用的。不应该直接访问或修改该数据库。 <code>/var/run/idmap/idmap.db</code> 用于缓存由临时 ID 映射和基于名称的映射生成的 ID 映射的数据库。该数据库的内容是专用的。不应该直接访问或修改该数据库。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	system/core-os
接口稳定性	请参见下文。

这些组件的接口稳定性如下所述：

`svc:/system/idmap`
Committed（已确定）

`/var/idmap/idmap.db`
Project Private（项目专用）

`/var/run/idmap/idmap.db`
Project Private（项目专用）

另请参见 [svcs\(1\)](#)、[idmap\(1M\)](#)、[svcadm\(1M\)](#)、[svccfg\(1M\)](#)、[defaultdomain\(4\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

附注

`idmapd` 服务由服务管理工具 (service management facility, SMF) 管理。`idmapd` 服务的服务标识符是 `svc:/system/idmap`。

使用 `svcadm` 命令执行此服务的管理操作，如启用、禁用或重新启动服务。这些操作需要 `solaris.smf.manage.idmap` 授权。使用 `svcs` 命令查询服务状态。

在将来的 Solaris 操作系统发行版中，该守护进程的功能可能会更改。

引用名	idsconfig – prepare a Directory Server Enterprise Edition (DSEE) to be populated with data and serve LDAP clients
用法概要	<pre>/usr/lib/ldap/idsconfig [-v] [-i <i>input_configfile</i>] [-o <i>output_configfile</i>]</pre>
描述	Use the <code>idsconfig</code> tool to set up a Directory Server Enterprise Edition (DSEE). You can specify the input configuration file with the <code>-i</code> option on the command line. Alternatively, the tool will prompt the user for configuration information. The input configuration file is created by <code>idsconfig</code> with the <code>-o</code> option on a previous run. The first time a server is set up, the user is prompted for all the required information. Future installations on that machine can use the configuration file previously generated by <code>idsconfig</code> using the <code>-o</code> option. The output configuration file contains the directory administrator's password in clear text. Thus, if you are creating an output configuration file, take appropriate security precautions. You should back up the directory server's configuration and data prior to running this command.
选项	The following options are supported: <code>-i <i>input_configfile</i></code> Specify the file name for <code>idsconfig</code> to use as a configuration file. This file will be read by <code>idsconfig</code>, and the values in the file will be used to configure the server. Do not manually edit <i>input_configfile</i>. The <i>input_configfile</i> is only partially validated, as <code>idsconfig</code> assumes that the file was created by a previous invocation of the command. <code>-o <i>output_configfile</i></code> Create a configuration file. <code>-v</code> Verbose output.
操作数	The following operands are supported: <i>input_configfile</i> Name of configuration file for <code>idsconfig</code> to use. <i>output_configfile</i> Configuration file created by <code>idsconfig</code>.
示例	示例 1 Prompting the User for Input In the following example, the user is prompted for information to set up DSEE. <pre>example# idsconfig</pre>

示例 2 Creating an Output Configuration File

In the following example, the user is prompted for information to set up DSEE, and an output configuration file, `config.1`, is created when completed.

```
example# idsconfig -o config.1
```

示例 3 Setting up DSEE Using the Specified Configuration File

In the following example, DSEE is set up by using the values specified in the configuration file, `config.1`. The verbose mode is specified, so detailed information will print to the screen.

```
example# idsconfig -v -i config.1
```

退出状态

The following exit values are returned:

- 0
Successful completion.
- >0
An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/nis
Interface Stability	Committed

另请参见

[ldapadd\(1\)](#), [ldapdelete\(1\)](#), [ldaplist\(1\)](#), [ldapmodify\(1\)](#), [ldapmodrdn\(1\)](#), [ldapsearch\(1\)](#), [ldap_cachemgr\(1M\)](#), [ldapaddent\(1M\)](#), [ldapclient\(1M\)](#), [resolv.conf\(4\)](#), [attributes\(5\)](#), [ldap\(5\)](#)

引用名	ifconfig – 配置网络接口参数
用法概要	<pre> ifconfig interface [address_family] [address [/prefix_length] [dest_address]] [addif address [/prefix_length]] [removeif address [/prefix_length]] [arp -arp] [auth_algs authentication_algorithm] [encr_algs encryption_algorithm] [encr_auth_algs authentication_algorithm] [auto-revarp] [broadcast_address] [deprecated -deprecated] [preferred -preferred] [destination dest_address] [ether [address]] [failover -failover] [group [name ""]] [index if_index] [ipmp] [metric n] [modlist] [modinsert mod_name@pos] [modremove mod_name@pos] [mtu n] [netmask mask] [plumb] [unplumb] [private -private] [nud -nud] [set [address] [/netmask]] [standby -standby] [subnet subnet_address] [tdst tunnel_dest_address] [token address/prefix_length] [tsrc tunnel_src_address] [trailers -trailers] [up] [down] [usesrc [name none]] [xmit -xmit] [encaplimit n -encaplimit] [thoplimit n] [router -router] [zone zonename -zone -all-zones] ifconfig [address_family] interface {auto-dhcp dhcp} [primary] [wait seconds] drop extend inform ping release start status </pre>
描述	ifconfig 命令用于向网络接口分配地址以及配置网络接口参数。由 ifconfig 命令配置的网络接口在重新引导后即不再存在。必须使用 ipadm(1M) 命令来配置持久性的网络接口。如果未指定任何选项，ifconfig 会显示网络接口的当前配置。如果指定了地址族，ifconfig 只报告特定于该地址族的详细信息。只有特权用户才能修改网络接口的配置。以大括号 { } 括起的选项表明必须指定这些选项之一。
DHCP 配置	使用 auto-dhcp 或 dhcp 参数的 ifconfig 用于控制接口的动态主机配置协议 (Dynamic Host Configuration Protocol, DHCP) 配置。在此模式中，ifconfig 用于控制 DHCP 客户机守护进程 dhcpgent(1M) 的运行。一旦使用 start 操作数将接口置于 DHCP 控制下，就不能在正常运行时使用 ifconfig 修改该接口的地址或特征。如果受 DHCP 控制的接口地址发生更改，dhcpgent 将会取消 DHCP 对该接口的控制。
选项	所支持的选项如下： addif address 在指定的物理接口中创建下一个未使用的逻辑接口。 all-zones 使接口可用于系统上的每个共享 IP 区域。要将数据传送到哪个适当区域由 tnzonecfg 数据库确定。仅当系统配置有 Solaris Trusted Extensions 功能时，此选项才可用。 tnzonecfg 数据库在 tnzonecfg(4) 手册页中进行了介绍，该手册页包含在 Solaris Trusted Extensions Reference Manual 中。

anycast

通过设置 **ANYCAST** 标志，将逻辑接口标记为任播地址。有关任播的更多信息，请参见下文的“接口标志”。

-anycast

通过清除 **ANYCAST** 标志，将逻辑接口标记为非任播地址。

arp

允许在网络级别地址与链路级别地址之间的映射中使用地址解析协议 (Address Resolution Protocol, ARP) (缺省值)。当前已对 IPv4 地址与 MAC 地址之间的映射实现此选项。

-arp

禁止对物理接口使用 ARP。不能对 IPMP IP 接口禁用 ARP。

auth_algs authentication algorithm

对于隧道，使用指定的验证算法启用 IPsec AH。该算法可以是编号或算法名称，也可以是 *any*，表示没有首选算法。必须在同一命令行上指定所有 IPsec 隧道属性。要禁用隧道安全性，请将 **auth_alg** 指定为 **none**。

现在，配置隧道的安全属性时最好使用 **ipseccnf(1M)** 命令。如果使用 **ipseccnf** 设置隧道的安全属性，则此关键字不会对隧道有影响。

auto-dhcp

使用 DHCP 自动获取此接口的地址。此选项有一个完全等效的别名，称为 **dhcp**。

对于 IPv6，指定的接口必须是第零个逻辑接口（物理接口名称），该接口具有链路本地地址。

primary

将接口定义为 **primary**。该接口定义为用于传送客户机范围配置数据的首选接口。在任何给定时间只能有一个接口是主接口。如果随后选择另一个接口作为主接口，它会取代先前的主接口。如果在客户机工作站引导后将接口指定为主接口，则没有多大意义，因为此时许多应用程序都已启动并使用从前一个主接口读取的数据进行了配置。

wait seconds

ifconfig 命令在操作完成或达到指定时间间隔（以这两个时间中较短的为准）之前将一直处于等待状态。如果未指定等待时间间隔，并且操作无法立即完成，则 **ifconfig** 将等待 30 秒直到请求的操作完成。也可以使用符号值 **forever**，表示永远。

drop

从 DHCP 控制中删除指定接口而不通知 DHCP 服务器，并记录当前租用供以后使用。此外，对于 IPv4，将 IP 地址设置为零。对于 IPv6，取消激活由 **dhcpgent** 激活的所有逻辑接口。

extend

尝试延长对接口 IP 地址的租用。此选项不是必需的，因为代理会在租期届满很久之前就自动延长租用。

inform

从 DHCP 获取网络配置参数，而无需获取对 IP 地址的租用。在通过 DHCP 之外的其他机制获取 IP 地址时，此选项十分有用。

ping

检查指定的接口是否受 DHCP 控制，这意味着该接口由 DHCP 代理管理并且工作正常。退出状态为 0 时表示成功。

release

通过通知服务器并放弃当前租用，放弃接口的 IP 地址。对于 IPv4，会将 IP 地址设置为零。对于 IPv6，会取消激活由 `dhcpageant` 激活的所有逻辑接口。

start

对接口启动 DHCP。

status

显示接口的 DHCP 配置状态。

auto-revarp

使用反向地址解析协议 (Reverse Address Resolution Protocol, RARP) 自动获取此接口的地址。如果该接口不支持 RARP，此操作将会失败；例如，IPoIB (IP over InfiniBand) 和 IPv6 接口。

broadcast address

仅适用于 IPv4。指定用于表示向网络进行广播的地址。缺省广播地址是主机部分全部为 1 的地址。如果为 `broadcast` 值指定“+”（加号），则会导致将广播地址重置为（可能是新的）地址和网络掩码的相应缺省值。`ifconfig` 的参数按从左至右的方向解释。因此，

```
example% ifconfig -a netmask + broadcast +
```

和

```
example% ifconfig -a broadcast + netmask +
```

可能会导致向接口的广播地址分配不同的值。

deprecated

将逻辑接口标记为已过时。与已过时接口关联的地址不会用作传出包的源地址，除非接口上没有其他可用地址或应用程序已显式绑定到此地址。状态显示会将 **DEPRECATED** 显示为标志的一部分。有关 `ifconfig` 支持的标志的信息，请参见[接口标志](#)。

-deprecated

将逻辑接口标记为未过时。与此类接口关联的地址可用作传出包的源地址。

preferred

将逻辑接口标记为首选接口。此选项仅适用于 IPv6 地址。与系统上配置的所有其他地址相比，分配给首选逻辑接口的地址将作为首选源地址，除非该地址相对于目

标地址处于不适当的范围。无论分配给哪一个物理接口，首选地址都将用作源地址。例如，可以配置回送接口的首选源地址，然后使用路由协议通告此地址的可访问性。

-preferred

将逻辑接口标记为非首选接口。

destination *dest_address*

为点对点接口设置目标地址。

dhcp

此选项是 **auto-dhcp** 选项的别名

down

将逻辑接口标记为“关闭”。（即，禁用 **IFF_UP** 位。）当逻辑接口标记为“关闭”时，系统不会尝试将分配给该接口的地址用作传出包的源地址，并且不会将发往该地址的传入包识别为传送到此主机。此外，当给定物理接口上的所有逻辑接口均“关闭”时，该物理接口本身也将处于禁用状态。

当逻辑接口关闭时，将该接口指定为输出（使用 **route(1M)** 命令中的 **-ifp** 选项或 **route(7P)** 套接字中的 **RTA_IFP**）的所有路由都将从转发表中删除。如果该接口恢复启用，则标记为 **RTF_STATIC** 的路由将返回到该表中，而未标记为 **RTF_STATIC** 的路由会直接被删除。

如果可能用于访问特定网关地址的所有逻辑接口均已关闭（在不带 **interface** 选项的情况下指定，如上一段落所述），则受影响的网关路由将被视为已设置 **RTF_BLACKHOLE** 标志。所有匹配的包都会因该网关不可访问而被放弃。

encaplimit *n*

将接口的隧道封装限制设置为 **n**。此选项仅适用于 IPv4-in-IPv6 和 IPv6-in-IPv6 隧道，它只修改底层 IPv6 隧道链路的 **encaplimit** 链路属性（请参见 **dladm(1M)**）。隧道封装限制用于控制包在离开任何隧道之前可进入的隧道的数量（即隧道嵌套级别）。

此选项已过时，已被 **dladm(1M)** **encaplimit** 链路属性所取代。

-encaplimit

禁止生成隧道封装限制。此选项仅适用于 IPv4-in-IPv6 和 IPv6-in-IPv6 隧道。此选项只是将底层 IPv6 隧道链路的 **encaplimit** 链路属性设置为 0（请参见 **dladm(1M)** **encaplimit**）。

此选项已过时，已被 **dladm(1M)** **encaplimit** 链路属性所取代。

encr_auth_algs *authentication algorithm*

对于隧道，使用指定的验证算法启用 IPsec ESP。此选项可以是编号或算法名称，也可以是 **any** 或 **none**，表示没有首选算法。如果指定了 ESP 加密算法但未指定验证算法，则 ESP 验证算法的缺省值将为 **any**。

现在，配置隧道的安全属性时最好使用 **ipseconf(1M)** 命令。如果使用 **ipseconf** 设置隧道的安全属性，则此关键字不会对隧道有影响。

encr_algs *encryption algorithm*

对于隧道，使用指定的加密算法启用 IPsec ESP。该选项可以是编号或算法名称。请注意，必须在同一命令行上指定所有 IPsec 隧道属性。要禁用隧道安全性，请将 **encr_alg** 的值指定为 **none**。如果指定了 ESP 验证算法但未指定加密算法，则 ESP 加密的缺省值将为 **null**。

现在，配置隧道的安全属性时最好使用 **ipseconf(1M)** 命令。如果使用 **ipseconf** 设置隧道的安全属性，则此关键字不会对隧道有影响。

ether [*address*]

如果未提供地址并且用户是超级用户或具有打开底层数据链路的足够特权，则会显示当前以太网地址信息。

在其他情况下，如果用户是超级用户或具有足够特权，会将接口的以太网地址设置为 *address*。该地址是表示为 *x:x:x:x:x:x* 格式的以太网地址，其中 *x* 是一个介于 0 与 FF 之间的十六进制数字。类似地，对于 IPoIB (IP over InfiniBand) 接口，该地址是 20 字节的以冒号分隔的十六进制数字，这些数字也介于 0 和 FF 之间。

某些（不是全部）以太网接口卡具有其自己的地址。要使用没有自己地址的卡，请参见 IEEE 802.3 规范第 3.2.3(4) 节中有关本地管理的地址空间的定义。请注意，IPMP 组中的所有 IP 接口都必须具有唯一的硬件地址；请参见 **in.mpathd(1M)**。

-failover

对逻辑接口设置 **NOFAILOVER**。这使关联的地址可供 **in.mpathd** 用来对关联的物理 IP 接口执行基于探测器的故障检测。不利的一面是，这也会对逻辑接口设置 **DEPRECATED**。不允许对 IPMP IP 接口执行此操作。

failover

清除逻辑接口的 **NOFAILOVER**。这是缺省值。在启用时这些逻辑接口会被迁移（请参见 **IP MULTIPATHING GROUPS**）。

group [*name* | ""]

当应用于物理接口时，此选项会将该接口放在指定的组中。如果组不存在，将会创建该组以及一个或多个 IPMP IP 接口（对于 IPv4 和/或 IPv6）。未标记为 **NOFAILOVER** 的所有 UP 地址都会被迁移到 IPMP IP 接口（请参见 **IP MULTIPATHING GROUPS**）。如果将组名指定为 ""，则会从该组中删除物理 IP 接口。

当应用于物理 IPMP IP 接口时，此选项会将 IPMP 组重命名为新名称。如果新名称已存在或将名称指定为 ""，此操作会失败。建议不要重命名 IPMP 组。应在通过 **ipmp** 子命令创建 IPMP IP 接口时，为该接口指定一个有意义的名称，系统也会将该名称用作 IPMP 组名。

index *n*

更改接口的接口索引。*n* 的值必须是未由其他接口使用的接口索引 (*if_index*)。 *if_index* 是用于在系统上唯一标识网络接口的非零正数。

ipmp

使用指定名称创建 IPMP IP 接口。必须分别创建供 IPv4 和 IPv6 使用的接口。*address_family* 参数控制该命令是应用于 IPv4 还是 IPv6（如果未指定，则应用于 IPv4）。所有 IPMP IP 接口都设置了 IPMP 标志。

metric *n*

将接口的路由度量设置为 *n*；如果未指定值，则缺省为 0。路由度量由路由协议使用。较高的度量会降低相应路由的适用性。度量为目标网络或主机的附加跳数。

modinsert *mod_name@pos*

将名为 *mod_name* 的模块插入设备流的 *pos* 位置处。该位置是相对于流头的位置。位置 0 表示紧接在流头下的位置。

根据 *modlist* 选项中的示例，使用以下命令将名为 *ipqos* 的模块插入到 *ip* 模块之下、防火墙模块之上：

```
example% ifconfig eri0 modinsert ipqos@2
```

随后会列出设备流中的所有模块，如下所示：

```
example% ifconfig eri0 modlist
0 arp
1 ip
2 ipqos
3 firewall
4 eri
```

modlist

列出设备流中的所有模块。

以下示例列出设备流中的所有模块：

```
example% ifconfig eri0 modlist
0 arp
1 ip
2 firewall
4 eri
```

modremove *mod_name@pos*

从设备流中删除 *pos* 位置的名为 *mod_name* 的模块。该位置是相对于流头的位置。

根据 *modinsert* 选项中的示例，在插入 *ipqos* 模块后使用以下命令从设备流中删除防火墙模块：

```
example% ifconfig eri0 modremove firewall@3
```

随后会列出设备流中的所有模块，如下所示：

```
example% ifconfig eri0 modlist
0 arp
1 ip
2 ipqos
3 eri
```


请注意，核心 IP 栈模块（例如，`ip` 和 `tun` 模块）无法被删除。

`mtu n`

将接口的最大传输单元设置为 `n`。对于许多类型的网络，`mtu` 具有上限；例如，对于以太网，上限为 `1500`。此选项会对受影响的接口设置 `FIXEDMTU` 标志。

`netmask mask`

仅适用于 IPv4。指定将网络划分为子网络时要保留的地址部分。掩码包含本地地址的网络部分和子网部分，子网部分从地址的主机字段获取。掩码由 1 和 0 组成，其中 1 表示 32 位地址中用于网络和子网部分的位位置，0 表示主机部分。掩码应至少包含标准网络部分，并且子网字段应与网络部分相邻。可通过以下四种方法之一指定掩码：

1. 使用带有前导 `0x` 的单个十六进制数字
2. 使用点记法地址
3. 使用“+”（加号）地址
4. 使用在网络数据库 `networks(4)` 中找到的伪主机名/伪网络名。

如果为网络掩码值指定“+”（加号），将在 `netmasks(4)` 数据库中查找掩码。此查找过程会在该数据库中查找匹配的最长网络掩码，方法是首先从作为键的接口 IPv4 地址开始，然后逐个将地址的各个低序位设置到掩码中。这种迭代式查找可确保在网络号中使用可变长度子网掩码时可以使用 `netmasks(4)` 数据库指定网络掩码。

如果将伪主机名/伪网络名提供为网络掩码值，则网络掩码数据可能位于 `hosts` 或 `networks` 数据库中。首先使用 `gethostbyname(3NSL)` 查找名称。如果在此处未找到，则在 `getnetbyname(3SOCKET)` 中查找。这些接口可能又会使用 `nsswitch.conf(4)` 确定要用来获取实际值的数据存储。

对于 `inet` 和 `inet6`，可将通过 `mask` 传达的相同信息指定为附加到 `address` 参数的 `prefix_length`。

`nud`

点对点物理接口启用相邻节点不可访问性检测机制。

`-nud`

点对点物理接口禁用相邻节点不可访问性检测机制。

`plumb`

对于物理 IP 接口，打开与物理接口名称关联的数据链路并设置 IP 使用该数据链路所需的激活。当与逻辑接口名称一起使用时，此命令用于在现有物理 IP 接口上创建特定的命名逻辑接口。

必须根据 `address_family` 参数分别为 IPv4 和 IPv6 激活接口（如果未指定该参数，则为 IPv4 激活接口）。在激活接口之前，`ifconfig -a` 不会显示该接口。

请注意，IPMP IP 接口不是绑定到特定的数据链路，而是使用 `ipmp` 子命令创建的。

`private`

向 `in.routed` 路由守护进程指示不能通告指定的逻辑接口。

-private

指定不可通告的接口。

removeif *address*

删除指定物理接口上与指定 *address* 匹配的逻辑接口。

router

对接口启用 IP 转发。如果启用，则会将接口标记为 **ROUTER** 并可通过该接口转发 IP 包。如果对 IPMP 组中的任何 IP 接口启用 **ROUTER**，也会对该 IPMP 组中的所有 IP 接口启用此选项。

-router

对接口禁用 IP 转发。IP 包不通过该接口转发。如果对 IPMP 组中的任何 IP 接口禁用 **ROUTER**，也会对该 IPMP 组中的所有 IP 接口禁用此选项。

set

为逻辑接口设置 *address* 和/或 *prefix_length*。

standby

将物理 IP 接口标记为 **STANDBY** 接口。如果接口标记为 **STANDBY** 并且属于 IPMP 组，将不会使用该接口进行数据通信，除非 IPMP 组中的另一个接口无法使用。当 **STANDBY** 接口已运行但未用于数据通信时，它还会标记为 **INACTIVE**。不允许对 IPMP IP 接口执行此操作。

-standby

清除接口的 **STANDBY**。这是缺省值。

subnet

为接口设置子网地址。

tdst *tunnel_dest_address*

设置隧道的目标地址。该地址不能与隧道的 *dest_address* 相同，因为包不通过此类隧道传出系统。

此选项已过时，已被 **dladm(1M)** `create-iptun` 和 `modify-iptun` 子命令所取代。

thoplimit *n*

为隧道接口设置跳数限制。对于 IPv6-in-IPv4 和 IPv4-in-IPv4 隧道，该跳数限制值用作 IPv4 头中的 TTL。对于 IPv6-in-IPv6 和 IPv4-in-IPv6 隧道，该跳数限制值用作 IPv6 头中的跳数限制。此选项只修改底层 IP 隧道链路的 `hoplimit` 链路属性（请参见 **dladm(1M)**）。

此选项已过时，已被 **dladm(1M)** `hoplimit` 链路属性所取代。

token *address/prefix_length*

设置接口的 IPv6 令牌以用于地址自动配置。

```
example% ifconfig eri0 inet6 token ::1/64
```

trailers

此标志以前会在某些链路级别导致 IPv4 包的非标准封装。此发行版附带的驱动程序将不再使用此标志。该标志仅用于兼容目的，将被忽略。

-trailers

禁止使用“尾部”链路级封装。

tsrc tunnel_src_address

设置隧道的源地址。这是外部封装 IP 头中的源地址。该地址必须是已使用 **ifconfig** 进行配置的另一个接口的地址。

此选项已过时，已被 **dladm(1M)** **create-iptun** 和 **modify-iptun** 子命令所取代。

unplumb

对于物理或 IPMP 接口，删除所有关联的逻辑 IP 接口并取消 IP 使用该接口所需的所有激活。对于 IPMP IP 接口，如果组不为空，则此命令将会失败。对于逻辑接口，将会删除该逻辑接口。

必须根据 *address_family* 参数分别为 IPv4 和 IPv6 取消激活接口（如果未指定该参数，则为 IPv4 取消激活接口）。此命令成功执行后，接口名称将不再显示在 **ifconfig -a** 的输出中。

up

将逻辑接口标记为 UP。因此，IP 模块将接受发往关联地址的包（除非该地址为零）以及任何关联的多播和广播 IP 地址。同样，IP 模块也允许以关联地址作为源地址发送包。至少必须有一个逻辑接口为 UP，这样关联的物理接口才能收发包。

usesrc [name | none]

指定要用于选择源地址的物理接口。如果使用关键字 **none**，则将清除以前的所有选择。

如果应用程序未使用 **bind(3SOCKET)** 选择非零源地址，系统将会根据传出接口和地址选择规则来选择适当的源地址（请参见 **ipaddrsel(1M)**）。

如果指定了 **usesrc** 并在输出的转发表中选择了指定的接口，则在选择源地址时系统会首先考虑指定的物理接口及其关联的逻辑接口。如果转发表中未列出可用地址，将应用普通的选择规则。例如，如果输入以下内容：

```
# ifconfig eri0 usesrc vni0
```

并且，**vni0** 分配有地址 10.0.0.1，系统会将 10.0.0.1 首选为源自本地连接并通过 **eri0** 发出的任何包的源地址。示例部分中提供了详细示例。

虽然可指定任何物理接口（甚至回送），但请注意您也可以指定虚拟 IP 接口（请参见 **vni(7d)**）。虚拟 IP 接口不与任何物理硬件关联，因此不易受硬件故障的影响。可以指定任意数量的物理接口来使用在单个虚拟接口上承载的源地址。这可以简化基于路由的多路径配置。如果有一个物理接口失败，通信将会通过其余某个正常运行的物理接口继续进行。此方案假定以某一方式（例如，路由协议）通告在虚拟接口上承载的地址的可访问性。

由于 `ifconfig preferred` 选项适用于所有接口，因此它的粒度比 `usesrc` 选项粗。`usesrc` 和 `setsrc`（路由命令）将按先后顺序覆盖此选项。

`xmit`

允许逻辑接口传送包。这是逻辑接口处于启用状态时的缺省行为。

`-xmit`

禁止通过接口发送包。但该接口仍将会继续接收包。

`zone zonename`

将逻辑接口放入区域 `zonename` 中。命名区域在处于就绪或运行中状态的内核中必须处于活动状态。当区域停止或重新引导时，接口将被取消激活。必须将区域配置为共享 IP 区域。`zonecfg(1M)` 用于向专用 IP 区域分配网络接口名称。

`-zone`

将 IP 接口放在全局区域中。这是缺省值。

操作数

下面介绍了 `interface` 操作数以及影响它的地址参数。

`interface`

采用以下格式之一的字符串：

- 名称后跟物理单元，例如 `eri0` 或 `ce1`
- 名称后跟物理单元:逻辑单元，例如 `eri0:1`
- `ip.tunN`、`ip6.tunN` 或 `ip6to4.tunN`（表示隐式 IP 隧道链路）

如果接口名称以短划线 (-) 开头，它将解释为指定一组接口的一组选项。在这种情况下，必须将 `-a` 包含在这些选项中，并可按任意顺序添加以下其他选项。如果给出其中某一接口名称，则其后跟的命令将应用于所有匹配的接口。

`-a`

将命令应用于指定地址族的所有接口。如果未在命令行或未通过 `/etc/default/inet_type` 提供地址族，将选择所有地址族。

`-d`

将命令应用于系统中的所有“关闭”接口。

`-D`

将命令应用于不受 DHCP（Dynamic Host Configuration Protocol，动态主机配置协议）控制的所有接口。

`-u`

将命令应用于系统中的所有“启动”接口。

`-Z`

将命令应用于用户区域中的所有接口。

`-4`

将命令应用于所有 IPv4 接口。

`-6`

将命令应用于所有 IPv6 接口。

address_family

地址族由 *address_family* 参数指定。ifconfig 命令当前支持以下族：inet 和 inet6。如果未指定地址族，则缺省值为 inet。

ifconfig 在显示接口信息时遵从 /etc/default/inet_type 文件中的 DEFAULT_IP 设置。如果 DEFAULT_IP 设置为 IP_VERSION4，ifconfig 将会省略与 IPv6 接口相关的信息。但是，如果在 ifconfig 命令行显式指定地址族（inet 或 inet6），该命令行会覆盖 DEFAULT_IP 设置。

address

对于 IPv4 族 (inet)，*address* 是位于主机名数据库（请参见 [hosts\(4\)](#)）或网络信息服务 (Network Information Service, NIS) 映射 hosts 中的主机名，或者用 Internet 标准“点记法”表示的 IPv4 地址。

对于 IPv6 族 (inet6)，*address* 是位于主机名数据库（请参见 [hosts\(4\)](#)）或网络信息服务 (Network Information Service, NIS) 映射 ipnode 中的主机名，或者是采用 Internet 标准以冒号分隔的十六进制格式的 IPv6 地址，该格式表示为 *x:x:x:x:x:x:x*，其中 *x* 是一个介于 0 和 FFFF 之间的十六进制数字。

prefix_length

对于 IPv4 和 IPv6 族 (inet 和 inet6)，*prefix_length* 是一个介于 0 和地址位数之间的数字。对于 inet，地址位数为 32；对于 inet6，地址位数为 128。*prefix_length* 指示网络掩码中的前导设置位数。

dest_address

如果提供了 *dest_address* 参数以及 *address* 参数，此选项指定点对点链路另一端通信方的地址。

tunnel_dest_address

可通过所配置隧道以外的其他接口访问的地址。此选项向隧道指示将通过隧道传送的包发送到的位置。此地址不能与所配置的接口目标地址相同。

tunnel_src_address

一个附加到已配置接口（已使用 ifconfig 配置为启动）的地址。

接口标志

ifconfig 命令支持以下接口标志。此上下文中的“地址”一词指逻辑接口（例如 eri0:0），而“接口”指物理接口（例如 eri0）。

ADDRCONF

地址来自无状态 **addrconf**。利用无状态机制，主机可以使用路由器所通告的信息与本地可用信息的组合生成其自己的地址。路由器会通告用于标识与链路关联的子网的前缀，而主机会生成用于在子网中唯一标识某个接口的“接口标识符”。如果缺少路由器提供的信息，主机可以生成链路本地地址。此标志特定于 IPv6。

ANYCAST

指示 **anycast** 地址。**anycast** 地址标识提供特定类型服务的系统组的最近成员。**anycast** 地址将分配给系统组。包将传送给 **anycast** 地址所标识的最近组成员，而不是传送给该组的所有成员。

BROADCAST

此 broadcast 地址是有效地址。此标志与 POINTTOPOINT 互斥

CoS

此接口支持某种形式的服务类 (Class of Service, CoS) 标记。此标志的一个示例是 VLAN 接口支持的 802.1D 用户优先级标记。对于 IPMP IP 接口，仅当组中的所有接口都设置了 CoS 时才会设置此标志。

请注意，仅对位于 VLAN 链路和以太网链路且其 [dladm\(1M\)](#) tagmode 链路属性设置为 normal 的接口设置此标志。

DEPRECATED

此地址已过时。此地址不会用作传出包的源地址，除非此接口上没有任何其他地址或应用程序已显式绑定到该地址。已过时的 IPv6 地址包含在 IPv6 重新编号标准机制中，在不用时会被彻底删除。对于 IPv4 和 IPv6，DEPRECATED 也会在所有 NOFAILOVER 地址上设置，但这一点在以后的发行版中可能会更改。

DHCRUNNING

逻辑接口的地址由 [dhcagent\(1M\)](#) 管理。对于 IPv6，如果已在第零个逻辑接口上启动 DHCPv6，也会在该接口上设置此标志；请参见 [in.ndpd\(1M\)](#)。

DUPLICATE

逻辑接口因该接口上配置的 IP 地址是重复地址而被禁用。网络上的其他某一节点正在使用此地址。如果此地址是由 DHCP 配置的或者是临时地址，系统可能会自动选择另一个地址。否则，系统会定期尝试恢复此地址，并在网络中的地址冲突消除后恢复接口。如果更改此地址或网络掩码，或将逻辑接口设置为 up，将会重新启动重复地址检测。如果将接口设置为 down，则会终止恢复并删除 DUPLICATE 标志。

FAILED

[in.mpathd](#) 守护进程已确定接口故障。FAILED 接口无法用于收发 IP 数据通信。如果对 IPMP 组中的某一物理 IP 接口设置了此标志，IP 数据通信将继续经由 IPMP 组中的其他可用物理 IP 接口传送。如果对 IPMP IP 接口设置了此标志，则整个组将失败，无法通过该组中的任何接口收发数据通信。

FIXEDMTU

已使用 -mtu 选项设置 MTU。此标志为只读标志。设置有此标志的接口具有固定 MTU 值，当驱动程序向 IP 通知链路 MTU 更改时可能发生的动态 MTU 更改不会影响该值。

INACTIVE

物理接口正常但无法用于根据管理策略收发数据通信。此标志最初由 standby 子命令设置，随后由 [in.mpathd](#) 进行控制。启用 FAILBACK=no 模式时也会设置该标志（请参见 [in.mpathd\(1M\)](#)），以指示 IP 接口已修复但未在使用。

IPMP

指示这是 IPMP IP 接口。

LOOPBACK

指示这是回送接口。

MULTI_BCAST

指示使用广播地址在此接口上进行多播。

MULTICAST

接口支持多播。IP 假定支持硬件广播或作为点对点链路的任何接口都将支持多播。

NOARP

此接口没有地址解析协议 (address resolution protocol, ARP)，该协议适用于没有广播地址的设备的所有接口。此标志特定于 IPv4。

NOFAILOVER

`in.mpathd` 可使用与此逻辑接口关联的地址，对关联的物理 IP 接口进行基于探测器的故障检测。

NOLOCAL

此接口没有地址，只是具有链路上的子网。

NONUD

对此接口禁用 NUD。节点会使用 NUD (neighbor unreachability detection，相邻节点不可访问性检测) 跟踪其相邻节点 (节点会主动向其发送包) 的可访问性状态，从而在检测到有相邻节点不可访问时执行任何恢复。此标志特定于 IPv6。

NORTEXCH

此接口不交换路由信息。对于 RIP-2，不通过此接口发送路由包。此外，似乎来自此接口的消息也不会收到响应。此接口的子网或地址不包含在通过其他接口发送到其他路由器的通告中。

NOXMIT

指示地址不会传送包。RIP-2 也不会通告此地址。

OFFLINE

接口处于脱机状态，因此不能收发 IP 数据通信。仅对 IPMP 组中的 IP 接口设置此标志。请参见 [if_mpadm\(1M\)](#) 和 [cfgadm\(1M\)](#)。

POINTOPOINT

指示地址是点对点链路。此标志与 BROADCAST 互斥

PREFERRED

此地址是首选 IPv6 源地址。该地址将用作与所有 IPv6 目标之间的 IPv6 通信的源地址，除非系统上有其他地址处于更合适的范围内。DEPRECATED 标志优先于 PREFERRED 标志。

PRIVATE

指示不通告此地址。对于 RIP-2，此接口用于发送通告。不过，子网和此地址都不包含在发送到其他路由器的通告中。

PROMISC

指示接口处于混杂模式的只读标志。与处于混杂模式的接口关联的所有地址都将显示 **PROMISC** 标志（例如，为了响应 `ifconfig -a`）。

ROUTER

指示可通过接口转发 IP 包。

RUNNING

指示已分配接口所需的资源。对于某些接口，此标志还指示链路已启动。对于 IPMP IP 接口，只要组中的一个 IP 接口处于活动状态，就会设置 **RUNNING**。

STANDBY

指示此物理接口不会用于数据通信，除非 IPMP 组中的其他接口无法使用。**INACTIVE** 和 **FAILED** 标志指示目前是否正在使用此物理接口。

TEMPORARY

指示这是 RFC 3041 中定义的临时 IPv6 地址。

UNNUMBERED

当链路的本地 IP 地址与系统中其他某一链路的本地地址相匹配时，将设置此标志。

UP

指示逻辑接口（及关联的物理接口）已启动。IP 模块将接受发往 **UP** 地址（该地址为零时除外）以及任何关联的多播和广播 IP 地址的包。同样，IP 模块还允许以 **UP** 地址作为源地址发送包。

VIRTUAL

指示物理接口没有底层硬件。无法通过虚拟接口发送或接收包。这些接口用于配置可用在多个接口上的本地地址。（另请参见 `usesrc` 选项。）

L3PROTECT

指示已使用 `dladm(1M)` 中的 `allowed-ips` 链路属性对物理接口实施第 3 层保护。

PROBER

指示 IPMP 组中的 **FAILED** 底层接口正在探测它是否已修复。**PROBER** 标志及其语义供 Solaris IPMP 实现内部使用，可能会有所更改。

逻辑接口

Solaris TCP/IP 允许多个逻辑接口与一个物理网络接口关联。这样可以向单台计算机分配多个 IP 地址，即使该计算机只能有一个网络接口也是如此。物理网络接口的名称采用 **驱动程序名称后跟物理单元号** 格式，而逻辑接口的名称采用 **驱动程序名称后跟物理单元号:逻辑单元号** 格式。物理接口是使用 `plumb` 命令配置到系统中的。例如：

```
example% ifconfig eri0 plumb
```

在物理接口已“激活”后，即可通过 `ifconfig` 命令的单独选项 `-plumb` 或 `-addif` 配置与该物理接口关联的逻辑接口。

```
example% ifconfig eri0:1 plumb
```

以上命令分配与物理接口 `eri0` 关联的特定逻辑接口。以下命令


```
example% ifconfig eri0 addif 192.168.200.1/24 up
```

在 `eri0` 物理接口上分配下一个可用的逻辑单元号，并分配 `address` 和 `prefix_length`。

可使用不同于关联物理接口的参数（`address`、`prefix_length` 等）来配置逻辑接口。也可以为与同一物理接口关联的多个逻辑接口指定不同的参数。每个逻辑接口都必须与一个现有的启动状态物理接口关联。例如，只有在已激活物理接口 `eri0` 后，才能配置逻辑接口 `eri0:1`。

要删除逻辑接口，请使用 `unplumb` 或 `removeif` 选项。例如，

```
example% ifconfig eri0:1 down unplumb
```

将删除逻辑接口 `eri0:1`。

Ip 多路径组

共享同一链路层广播域的各物理接口必须通过 `group` 子命令集中到单个 IP 多路径 (IP Multipathing, IPMP) 组中。每个 IPMP 组都有一个关联的 IPMP IP 接口；该接口可通过 `ipmp` 子命令显式创建（首选方法），也可通过 `ifconfig` 隐式创建，以便将 IP 接口放入新 IPMP 组。隐式创建的 IPMP 接口将命名为 `ipmpN`，其中 `N` 是不会与现有 IP 接口名称或 IPMP 组名发生冲突的最小整数。

每个创建的 IPMP IP 接口都具有匹配的 IPMP 组名，但可使用 `group` 子命令更改该名称。每个 IPMP IP 接口都承载一组高度可用的 IP 地址。只要组中至少有一个接口处于活动状态，这些地址就是始终可访问的，其中“活动”定义为至少有一个 UP 地址并已清除 INACTIVE、FAILED 和 OFFLINE。承载在 IPMP IP 接口中的 IP 地址可以静态进行配置，也可以通过 `dhcp` 子命令由 DHCP 配置。

分配给同一 IPMP 组的接口将被视为等效的，并由 `in.mpathd` 监视故障情况。如果组中的活动接口保留，将以对基于套接字的应用程序透明的方式处理 IP 接口故障（及任何后续修复）。此外，IPMP 还与动态重新配置框架集成（请参见 [cfgadm\(1M\)](#)），允许通过对基于套接字的应用程序不可见的方式更换网络适配器。

IP 模块自动在 IPMP 组中的所有活动接口之间对所有传出通信进行负荷分配。同样，也会在活动接口之间分发承载在 IPMP IP 接口中的所有 UP 地址，从而增强传入负荷分配能力。使用 [ipmpstat\(1M\)](#) 实用程序可以观察 IPMP 子系统的许多方面，包括 IP 数据地址与 IP 接口的当前绑定。

在将接口放入 IPMP 组后，所有 UP 逻辑接口都将“迁移”至 IPMP IP 接口，以供该组使用，但以下情况下除外：

- 逻辑接口标记为 NOFAILOVER；
- 逻辑接口承载 IPv6 链路本地地址；
- 逻辑接口承载 IPv4 0.0.0.0 地址。

同样，当接口位于组中后，如果对逻辑接口进行了更改，使它成为 UP 状态并且不属于上述任一排除情况，它也会迁移到关联的 IPMP IP 接口。逻辑接口从不回迁，即使从组中删除了提供地址的物理接口也是如此。

放入 IPMP 组中的每个接口都可以选择配置“测试”地址，`in.mpathd` 将使用该地址进行基于探测器的故障检测；请参见 [in.mpathd\(1M\)](#)。这些地址在标记为 UP 之前必须标记为 NOFAILOVER（使用 `-failover` 子命令）。测试地址也可以通过 `dhcp` 子命令由 DHCP 获取。

有关 IPMP 的更多背景知识，请参见《[Oracle Solaris 管理：网络接口和网络虚拟化](#)》的 IPMP 相关章节。

配置 IPv6 接口

当 IPv6 物理接口已激活并通过 `ifconfig` 配置为启动时，它会自动分配有 IPv6 链路本地地址，该地址的后 64 位将基于该接口的 MAC 地址计算得出。

```
example% ifconfig eri0 inet6 plumb up
```

以下示例显示链路本地地址带有前缀 `fe80::/10`。

```
example% ifconfig eri0 inet6
ce0: flags=2000841<UP,RUNNING,MULTICAST,IPv6>
      mtu 1500 index 2
      inet6 fe80::a00:20ff:fe8e:f3ad/10
```

链路本地地址仅用于本地子网上的通信，对其他子网是不可见的。

如果链路通告前缀中存在通告 IPv6 路由器，则新激活的 IPv6 接口将根据前缀通告自动配置逻辑接口。例如，对于前缀通告 `2001:0db8:3c4d:0:55::/64`，自动配置的接口如下所示：

```
eri0:2: flags=2080841<UP,RUNNING,MULTICAST,ADDRCONF,IPv6>
      mtu 1500 index 2
      inet6 2001:0db8:3c4d:55:a00:20ff:fe8e:f3ad/64
```

即使链路上没有前缀通告，仍可以手动分配全局地址，例如：

```
example% ifconfig eri0 inet6 addif \
2001:0db8:3c4d:55:a00:20ff:fe8e:f3ad/64 up
```

配置 IP-over-IP 隧道接口

IP 隧道在概念上由两部分组成：一条连接两个或更多 IP 节点的虚拟链路和一个 IP 接口；该接口位于此链路之上，使系统可以发送和接收由底层链路封装的 IP 包。

可使用 [dladm\(1M\)](#) 命令配置隧道链路，使用 `ifconfig` 在这些隧道链路上配置 IP 接口。IPv4-over-IPv4 隧道是通过在 IPv4 隧道链路上激活 IPv4 接口来创建的。IPv6-over-IPv4 隧道是通过在 IPv4 隧道链路上激活 IPv6 接口来创建的。

在 IP 隧道链路上激活 IPv6 接口后，将自动设置其 IPv6 地址。对于 IPv4 和 IPv6 隧道，将配置以下格式的源和目标链路本地地址：`fe80::interface-id`。对于 IPv4 隧道，`interface-id` 是 IPv4 隧道源或目标地址。对于 IPv6 隧道，`interface-id` 是 IPv6 隧道源或目标地址的后 64 位。例如，对于位于 10.1.2.3 与 10.4.5.6 之间的 IPv4 隧道，IPv6 接口的 IPv6 链路本地源和目标地址将分别为 `fe80::a01:203` 和 `fe80::a04:506`。对于位于 `2000::1234:abcd` 与 `3000::5678:abcd` 之间的 IPv6 隧道，该接口的 IPv6 链路本地源和目标地址将分别为 `fe80::1234:abcd` 和 `fe80::5678:abcd`。就像任何其他点对点接口一样，也可通过显式指定地址来覆盖这些缺省链路本地地址。

对于 6to4 隧道，将配置 `2002:tsrc::1/16` 格式的 6to4 全局地址。`tsrc` 部分是隧道 IPv4 源地址。6to4 接口的前缀长度将自动设置为 16，因为所有 6to4 包（位于 `2002::/16` 范围内的目标）都将转发到 6to4 隧道接口。例如，对于隧道源为 75.1.2.3 的 6to4 链路，IPv6 接口将具有 `2002:4b01:203::1/16` 格式的地址。

可以使用 `addif` 选项或通过激活其他逻辑接口，添加其他 IPv6 地址。

为了实现向后兼容性，激活具有专用名称的隧道 IP 接口时，将会隐式创建隧道链路，而不会调用 `dladm create-iptun`。这些隧道名称包括：

<code>ip.tunN</code>	IPv4 隧道
<code>ip6.tunN</code>	IPv6 隧道
<code>ip.6to4tunN</code>	6to4 隧道

这些隧道是“隐式隧道”，在 `dladm show-iptun` 输出中以 `i` 标志表示。激活这些专用 IP 接口的隧道链路将会自动创建，当释放最后一个引用（也就是取消激活最后一个 IP 接口）时也会自动将其删除。

`ifconfig` 的 `tsrc`、`tdst`、`encaplim` 和 `hoplimit` 选项已过时，保留它们只是为了实现向后兼容性。这些选项等效于其 `dladm(1M)` 对应项。

显示隧道安全设置

IP 隧道接口的 `ifconfig` 输出指示是否为底层 IP 隧道链路配置 IPsec 策略。例如，如果存在 IPsec 策略，将会显示如下格式的一行内容：

```
tunnel security settings --> use 'ipsecconf -ln -i ip.tun1'
```

如果未使用 `ifconfig` 或 `ipsecconf(1M)` 设置安全策略，则不会显示隧道安全设置。

示例

示例1 使用 `ifconfig` 命令

如果未将工作站连接到以太网，则应将网络接口（例如，`eri0`）标记为“down”，如下所示：

```
example% ifconfig eri0 down
```

示例2 打印寻址信息

要打印出每个接口的寻址信息，请使用以下命令：

```
example% ifconfig -a
```

示例3 重置广播地址

要在正确设置网络掩码后重置每个接口的广播地址，请使用下一个命令：

```
example% ifconfig -a broadcast +
```

示例4 更改以太网地址

要更改接口 `ce0` 的以太网地址，请使用以下命令：

```
example% ifconfig ce0 ether aa:1:2:3:4:5
```

示例5 配置 IP-in-IP 隧道

要配置 IP-in-IP 隧道，请先创建 IP 隧道链路（`tunsrc` 和 `tundst` 是在 `/etc/hosts` 中有相应 IPv4 条目的主机名）：

```
example% dladm create-iptun -T ipv4 -s tunsrc -d tundst tun0
```

然后激活点对点接口，应提供源和目标地址（`mysrc` 和 `thedst` 是在 `/etc/hosts` 中有相应 IPv4 条目的主机名）：

```
example% ifconfig tun0 plumb mysrc thedst up
```

使用 [ipsecconf\(1M\)](#)（如上所述）配置隧道安全属性。

IPv6 隧道的配置是通过对 `create-iptun` 使用 `ipv6` 隧道类型来完成的。也可以在任何类型的隧道上激活 IPv6 接口。

示例6 配置 6to4 隧道

要配置 6to4 隧道，请先创建 6to4 隧道链路（`myv4addr` 是在 `/etc/hosts` 中有相应 IPv4 条目的主机名）：

```
example% dladm create-iptun -T 6to4 -s myv4addr my6to4tun0
```

然后在此链路上激活 IPv6 接口：

```
example% ifconfig my6to4tun0 inet6 plumb up
```

该接口的 IPv6 地址将会自动设置（如上所述）。

示例7 对接口配置 IP 转发

要对单个接口启用 IP 转发，请使用以下命令：

```
example% ifconfig eri0 router
```

要对单个接口禁用 IP 转发，请使用以下命令：

```
example% ifconfig eri0 -router
```

示例8 使用虚拟接口配置源地址选择

以下命令配置源地址选择，使本地生成的、没有绑定源地址且在 `qfe2` 上传出的每个包都首选 `vni0` 中承载的源地址。

```
example% ifconfig qfe2 usesrc vni0
```

示例8 使用虚拟接口配置源地址选择 (续)

qfe2 和 vni0 接口的 ifconfig -a 输出显示如下所示：

```
qfe2: flags=1100843<UP,BROADCAST,RUNNING,MULTICAST,ROUTER,IPv4> mtu
      1500 index 4
      usesrc vni0
      inet 1.2.3.4 netmask fffffff0 broadcast 1.2.3.255
      ether 0:3:ba:17:4b:e1
vni0: flags=20011100c1<UP,RUNNING,NOARP,NOXMIT,ROUTER,IPv4,VIRTUAL>
      mtu 0 index 5
      srcof qfe2
      inet 3.4.5.6 netmask ffffffff
```

注意观察以上 ifconfig 输出中的 usesrc 和 srcof 关键字。这些关键字也会显示在物理接口的逻辑实例中，即使这是每物理接口参数也不例外。用于配置接口的 ifconfig 没有 srcof 关键字。此信息通过设置了 usesrc 的接口集自动确定。

以下命令使用 none 关键字撤消了上述 ifconfig usesrc 命令的结果。

```
example% ifconfig qfe2 usesrc none
```

在执行此命令之后，会显示 ifconfig -a 输出，如下所示：

```
qfe2: flags=1100843<UP,BROADCAST,RUNNING,MULTICAST,ROUTER,IPv4> mtu
      1500 index 4
      inet 1.2.3.4 netmask fffffff0 broadcast 1.2.3.255
      ether 0:3:ba:17:4b:e1
vni0: flags=20011100c1<UP,RUNNING,NOARP,NOXMIT,ROUTER,IPv4,VIRTUAL>
      mtu 0 index 5
      inet 3.4.5.6 netmask ffffffff
```

请注意上述输出中没有 usesrc 和 srcof 关键字。

示例9 为 IPv6 地址配置源地址选择

以下命令为 IPv6 地址配置源地址选择，该命令选择了 vni0 中承载的源地址。

```
example% ifconfig qfe1 inet6 usesrc vni0
```

在执行此命令之后，会显示 ifconfig -a 输出，如下所示：

```
qfe1: flags=2000841<UP,RUNNING,MULTICAST,IPv6> mtu 1500 index 3
      usesrc vni0
      inet6 fe80::203:baff:fe17:4be0/10
      ether 0:3:ba:17:4b:e0
vni0: flags=2002210041<UP,RUNNING,NOXMIT,NUD,IPv6,VIRTUAL> mtu 0
      index 5
      srcof qfe1
      inet6 fe80::203:baff:fe17:4444/128
```

示例9 为 IPv6 地址配置源地址选择 (续)

```
vni0:1: flags=2002210040<RUNNING,NOXMIT,NUD,IPv6,VIRTUAL> mtu 0
  index 5
  srcof qfe1
  inet6 fec0::203:baff:fe17:4444/128
vni0:2: flags=2002210040<RUNNING,NOXMIT,NUD,IPv6,VIRTUAL> mtu 0
  index 5
  srcof qfe1
  inet6 2000::203:baff:fe17:4444/128
```

根据在 qfe1 上传出的包的目标范围，将从 vni0 及其别名中选择具有适当范围的源地址。

示例10 对共享 IP 区域使用源地址选择

以下示例说明在 Solaris 中如何将 usesrc 功能与 zones(5) 工具结合使用。在全局区域中会调用以下各命令：

```
example% ifconfig hme0 usesrc vni0
example% ifconfig eri0 usesrc vni0
example% ifconfig qfe0 usesrc vni0
```

在执行上一个命令之后，将显示虚拟接口的 ifconfig -a 输出，如下所示：

```
vni0: flags=20011100c1<UP,RUNNING,NOARP,NOXMIT,ROUTER,IPv4,VIRTUAL>
  mtu 0 index 23
  srcof hme0 eri0 qfe0
  inet 10.0.0.1 netmask ffffffff
vni0:1:
  flags=20011100c1<UP,RUNNING,NOARP,NOXMIT,ROUTER,IPv4,VIRTUAL> mtu 0
  index 23
  zone test1
  srcof hme0 eri0 qfe0
  inet 10.0.0.2 netmask ffffffff
vni0:2:
  flags=20011100c1<UP,RUNNING,NOARP,NOXMIT,ROUTER,IPv4,VIRTUAL> mtu 0
  index 23
  zone test2
  srcof hme0 eri0 qfe0
  inet 10.0.0.3 netmask ffffffff
vni0:3:
  flags=20011100c1<UP,RUNNING,NOARP,NOXMIT,ROUTER,IPv4,VIRTUAL> mtu 0
  index 23
  zone test3
  srcof hme0 eri0 qfe0
  inet 10.0.0.4 netmask ffffffff
```

示例 10 对共享 IP 区域使用源地址选择 (续)

每个区域 (test1、test2 和 test3) 都有一个虚拟接口别名。将从同一区域的虚拟接口别名中选择源地址。虚拟接口别名是使用 [zonecfg\(1M\)](#) 创建的，如下所示：

```
example% zonecfg -z test1
zonecfg:test1> add net
zonecfg:test1:net> set physical=vni0
zonecfg:test1:net> set address=10.0.0.2
```

test2 和 test3 区域接口及地址的创建方法相同。

示例 11 禁用 DHCPv6

以下示例说明如何在所有接口上禁止自动使用 DHCPv6，以及如何直接在名为 hme0 的接口上关闭 DHCPv6。有关 DHCPv6 自动配置机制的更多信息，请参见 [in.ndpd\(1M\)](#) 和 [ndpd.conf\(4\)](#)。

```
example% echo ifdefault StatefulAddrConf false >> /etc/inet/ndpd.conf
example% pkill -HUP -x in.ndpd
example% ifconfig hme0 dhcp release
```

文件

/etc/netmasks
网络掩码数据。

/etc/default/inet_type
缺省 Internet 协议类型。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
命令行选项的接口稳定性	Committed (已确定)
命令输出的接口稳定性	Uncommitted (未确定)

另请参见

[dhcpinfo\(1\)](#)、[cfgadm\(1M\)](#)、[dhcpagent\(1M\)](#)、[dladm\(1M\)](#)、[if_mpadm\(1M\)](#)、[in.mpathd\(1M\)](#)、[in.](#)
《Oracle Solaris 管理：IP 服务》

诊断

- ifconfig 发送的消息会指出：
- 指定的接口是否不存在
 - 请求的地址是否未知
 - 用户是否未经授权就尝试更改接口的配置

附注

在选择主机名时，请不要选择 `broadcast`、`down`、`private`、`trailers`、`up` 名称或其他可能的选项名称。如果选择其中任一名称作为主机名，可能会导致很难诊断的异常问题。

引用名	if_mpadm – administer interfaces in an IP multipathing group
用法概要	<code>if_mpadm -d -r ifname</code>
描述	The <code>if_mpadm</code> utility administers IP interfaces that are part of an IP Multipathing (IPMP) group. Currently, administration is limited to offlining IP interfaces and undoing previous offline operations. When an IP interface is taken offline, all IP data traffic that was flowing over the IP interface is moved to another IP interface in the IPMP group. In addition, all UP IP addresses hosted on the IP interface are brought down, causing <code>in.mpathd(1M)</code> to stop probe-based failure detection on the IP interface. As a result, an offline IP interface will not be used for any inbound or outbound IP traffic. Only IP interfaces that are in an IPMP group may be brought offline. If the IP interface is the last functioning interface in the IPMP group, the offline operation will fail. When an offline operation is undone, any IP addresses hosted on that IP interface are brought UP and will be considered by <code>in.mpathd</code> for probe-based failure detection. In addition, provided the IP interface is otherwise active (see <code>in.mpathd(1M)</code>), it will again be used to send and receive IP data traffic for the IPMP group. Note that not all offline operations can be undone. For instance, <code>in.mpathd</code> may have offlined an IP interface because its hardware address was not unique within its IPMP group. The <code>ipmpstat</code> utility can be used to determine why an IP interface is offline, identify which IP interfaces in a group are being used for inbound and outbound IP traffic, and more; see <code>ipmpstat(1M)</code>.
选项	The <code>if_mpadm</code> utility supports the following options: <code>-d ifname</code> Offline the IP interface specified by <i>ifname</i>. If <i>ifname</i> is not in an IPMP group, or the offline would cause the IPMP group to lose network connectivity, the operation will fail. <code>-r ifname</code> Undo a previous offline of the IP interface specified by <i>ifname</i>. If <i>ifname</i> is not offline, the operation will fail.
示例	示例 1 Offlining an IP Interface The following command offlines the IP interface <code>under0</code>, causing any IP packets that were being sent and received through it to be handled by another IP interface in its group. <pre>example% if_mpadm -d under0</pre> 示例 2 Undoing a Previous Offline Operation Use the following command to undo the operation in the previous example: <pre>example% if_mpdadm -r under0</pre>

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Uncommitted

另请参见

[ifconfig\(1M\)](#), [in.mpathd\(1M\)](#), [ipmpstat\(1M\)](#), [attributes\(5\)](#)

诊断

cannot offline: no other functioning interfaces are in its IPMP group.

描述: This message means that offlining the IP interface would leave the IPMP group without network connectivity.

cannot offline: not a physical interface or not in an IPMP group

描述: This means that the IP interface is not an underlying interface in an IPMP group, and therefore is not eligible to be offlined.

引用名	ifparse – parse ifconfig command line
用法概要	<code>/usr/sbin/ifparse [-fs] addr_family commands</code>
描述	Use the ifparse command to parse the ifconfig(1M) command line options and output substrings, one per line, as appropriate. If no options are specified, ifparse returns the entire ifconfig command line as a series of substrings, one per line.
选项	The ifparse command supports the following options: -f Lists only substrings of the ifconfig command line that are relevant to IP network multipath failover -s Lists only substrings of the ifconfig command line that are not relevant to IP network multipath failover
操作数	The ifparse command <i>does not</i> support the <i>interface</i> operand of the ifconfig command.
示例	示例 1 Parsing Command Line Options Relevant to Failover The following example shows the use of the ifparse command to parse the command line options relevant to IP network multipath failover: <pre>example# ifparse -f inet 1.2.3.4 up group one addif 1.2.3.5 -failover up set 1.2.3.4 up</pre> 示例 2 Parsing Command Line Options That Are Not Relevant to Failover The following example shows the use of the ifparse command to parse the command line options that are not relevant to IP network multipath failover: <pre>example# ifparse -s inet 1.2.3.4 up group one addif 1.2.3.5 -failover up group one addif 1.2.3.5 -failover up</pre> 示例 3 Parsing the Command Line For All Options The following example shows the use of the ifparse command to parse the command line for all ifconfig options: <pre>example# ifparse inet 1.2.3.4 up group one addif 1.2.3.5 -failover up group one set 1.2.3.4 up addif 1.2.3.5 -failover up</pre>
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Obsolete

另请参见 [ifconfig\(1M\)](#), [attributes\(5\)](#)

诊断

usage: -fs <addr_family> <commands>
描述: This message indicates an invalid command line.

ifparse: Not enough space
描述: This message indicates insufficient memory.

ifparse: dhcp not supported for inet6
描述: DHCP operations are not supported for the inet6 address family.

ifparse: Operation <operation> not supported for <addr_family>
描述: Most operations cannot be used with all address families. For example, the broadcast operation is not supported on the inet6 address family.

ifparse: no argument for <operation>
描述: Some operations, for example broadcast, require an argument.

附注

The ifparse command is classified as an obsolete interface. It will likely be removed in a future release. You should not develop applications that depend upon this interface.

引用名	iiadm – 用于控制 Sun StorageTek Availability Suite Point-in-Time Copy 操作的命令行接口
用法概要	<pre> iiadm -e {ind shd} master_vol shadow_vol bitmap_vol iiadm -ne ind master_vol shadow_vol bitmap_vol iiadm [-p] [-n] {-c -u} {s m} volume_set iiadm [-adDiLR] volume_set iiadm [-p] [-n] -w volume_set iiadm [-hilLv] iiadm -P delay units volume_set iiadm -P volume_set iiadm -A overflow_vol volume_set iiadm [-OQ] overflow_vol iiadm -E volume_set iiadm [-IJ] volume_set bitmap iiadm -g group_name [-aAcDDeiLLmPRuw] iiadm [-C] cluster_tag [options] </pre>
描述	Point-in-Time Copy 软件是 Solaris 操作系统的时间点快照功能。 Point-in-Time Copy 快照是暂时停止的卷的复制视图，此视图可立即使用，且时间固定。创建快照后，Point-in-Time Copy 软件允许对主卷数据和影子卷数据立即进行读取/写入访问。 Point-in-Time Copy 软件从建立快照之时起跟踪主卷和影子卷之间的差异（由写入所致）。此功能允许访问主卷数据的应用程序独立于访问影子卷数据的应用程序继续运行，反之亦然。 Point-in-Time Copy 软件对主卷和影子卷之间的差异的跟踪功能有助于在以后进行快速重新同步或完整复制。可以从影子卷重新同步到主卷，也可以从主卷重新同步到影子卷。 在（重新）建立时间点之后（当返回 CLI 提示符或读取下一 shell 脚本命令时），可以立即重新挂载主卷或恢复使用主卷的应用程序。此外，还可以挂载并立即访问影子卷。 iiadm 命令行实用程序在每次调用命令时仅执行一个操作。为此，除组合下列所有命令修饰符以外，不能组合多个选项： ■ 如果未输入操作项，iiadm 显示当前配置的 Point-in-Time Copy 集（未暂停）的列表。如果输入了多个操作项或错误指定的操作项，iiadm 向标准错误显示具体错误消息，然后显示简要的用法汇总。

- 对于 Point-in-Time Copy 选项 ENABLE (-e)、COPY (-c) 和 UPDATE (-u)，存在两个关联的影子卷选择限定符 ({ind|dep})，用于指定要创建的 Point-in-Time Copy 卷集的类型。

独立 (ind) 快照导致 Point-in-Time Copy 软件执行从主卷到影子卷的完整卷复制操作。复制完成后，影子卷数据在建立时与主卷数据相同。如果您需要两个物理数据副本，请创建单独的影子卷。单独的影子卷的大小必须等于或大于主卷的大小。对于要考虑从影子卷重新同步到主卷的环境，Oracle 建议主卷和影子卷应采用相同大小。

从属 (dep) 快照导致 Point-in-Time Copy 软件不会执行完整的卷复制。对于所有未修改的数据块（除非请求，否则不会复制这些数据块），生成的影子卷依赖于主卷。如果您不需要两个物理数据副本，请创建从属影子卷。从属影子卷的大小可以等于或小于主卷的大小。较小的影子卷称为**压缩从属影子卷**，当对 Point-in-Time Copy 卷集所做的更改量小于主卷的整个大小时，通常会使用压缩从属影子卷。

使用下面的语法格式，您可以在 Sun Cluster 环境中创建可导出的独立影子卷：

```
# iiadm -ne ind master shadow bitmap
```

由于压缩从属影子卷的大小在启用 Point-in-Time Copy 卷集时建立，因此使用压缩从属影子卷时会出现问题。如果在使用卷集时对整个卷集所做的更改量超出了为影子卷分配的空间，则会将影子卷标记为空间不足。即使影子卷空间不足，也可以从影子卷进行读取，直到请求没有相应空间的部分数据为止。一旦发生此情况，读取将失败，并且会将影子卷标记为脱机。

为了解决此问题，Point-in-Time Copy 支持将**溢出**卷关联到现有 Point-in-Time Copy 从属卷集的功能。因此，如果压缩从属影子卷的大小太小，或者对卷集执行了计划外的更改量，更改后的数据可重定向到关联的溢出卷。为了帮助有效使用此溢出卷，可以根据需要将溢出卷与多个 Point-in-Time Copy 卷集相关联。

注意事项

在调用 Point-in-Time Copy enable、copy 或 update 操作之前，Point-in-Time Copy 确保未挂载影子卷，以免出现文件系统故障。同时，建议您仅在创建时间点快照时卸载或暂停（停止）使用主卷的所有应用程序。这确保创建根本一致的时间点快照。

如果已暂停（而非卸载）主卷，建议使用 **fsck(1M)** 等卷验证实用程序验证新时间点影子卷的完整性。其原因在于，Point-in-Time Copy 在**已卸载**影子卷上创建了**已挂载**主卷的时间点副本。挂载影子卷时，文件系统检测到影子卷处于**已挂载**状态。通常，此状态仅在系统崩溃时出现，因此文件系统尝试验证此卷的完整性，并假定发生了系统故障，而不是 Point-in-Time Copy。

环境选项

/usr/kernel/drv/ii.conf 配置文件中的 ii_bitmap 变定位图卷操作语义，如下所示：

- 0 表示仅在内存或恢复操作中保存位图。
- 1 表示在内存和磁盘中保存位图。这是缺省值。

如果在使用 ii_bitmap=0 时发生系统故障，影子卷可能会不一致，并且无法进行快速重新同步。

如果 Point-in-Time Copy 与网络存储组件远程镜像配合使用或在 Sun Cluster 中使用，请设置 `ii_bitmap=1`。

`/usr/kernel/drv/ii.conf` 配置文件中的 `ii_debug` 变量确定在 Point-in-Time Copy 处理过程中输出到系统控制台 `/dev/console` 的信息日志量。

- 0 表示不会向系统控制台发送日志。
- 1 表示向系统控制台发送信息日志。
- 2 表示向系统控制台发送开发日志。

选项

`iiadm` 实用程序支持下列选项。

`-e{ind|dep} master_vol shadow_vol bitmap_vol`

针对指定主卷、影子卷和位图卷启用 Point-in-Time Copy。

启用投影集处理功能确保可访问指定卷，未挂载 `shadow_vol`，并为创建的投影集类型正确确定 `bitmap_vol` 的大小。此外，它还确保这些卷可由 SV 驱动程序控制（否则，则将其放置在此位置），初始化位图卷，并且如果卷集是独立投影集，还会启动完整复制操作。

成功启用后，Point-in-Time Copy 将指定 `master_vol`、`shadow_vol` 和 `bitmap_vol` 名称以及启用类型（`ind` 或 `dep`）存储到 Point-in-Time Copy 配置存储中。配置存储包含当前配置的所有 Point-in-Time Copy 卷集及其相关配置属性。（有关独立和从属影子卷的语义，请参见上面的讨论。）

`master_vol` 是创建时间点快照所基于的卷。

`shadow_vol` 是包含时间点快照的卷。

`bitmap_vol` 用于跟踪影子卷和主卷之间的区别。暂停或恢复 Point-in-Time Copy 投影操作后，可在永久性存储中存储或检索位图卷（保存在内核内存中）。与位图卷关联的存储应当像影子卷存储一样冗余。

`shadow_vol` 名称是需要指定 `volume_set` 名称的所有 `iiadm` 选项已知的 Point-in-Time Copy 投影集的名称。

`-d volume_set`

禁用与指定 `volume_set` 关联的 Point-in-Time Copy 卷集。

如果 Point-in-Time Copy 在上述 `-e ind` 选项指定的独立模式下运行，影子卷数据包含与禁用前相同的数据（假定未进行写入）。现在，主卷和影子卷均为单独的时间点副本，因此用户可以访问主卷和影子卷。

在完整副本处于活动状态期间，独立卷将按照从属卷的方式运行。为了确保卷不再处于完整复制模式，请发出以下命令以便等待完整复制完成：

```
# iiadm -w volume_set
```

[-p] -u s volume_set
从主卷更新影子卷。

将主卷的时间点副本更新到影子卷。*volume_set* 是包含主卷和影子卷的 Point-in-Time Copy 投影集。此选项提供影子卷的快速重新同步，并创建主卷的增量副本。此更新复制在主卷和影子卷之间标记为不同的所有 32 KB 段。它不会复制所有主卷数据，而仅复制更改的数据。复制数据时，影子卷依赖于主卷。

使用此选项之前，暂时停止卷的工作负荷；并使主机应用程序停止写入到卷。这可以确保时间点数据一致。在使用目标卷执行任何其他操作之前，您可以使用 **iiadm -i volume_set** 以可视方式检查此副本或更新操作的状态，或者使用 **iiadm -w volume_set** 以交互方式检查状态（通过 shell 或脚本）。

通过使用选项 **-p**、**iiadm -p -u s**，此命令支持 PID（Process Identifier，进程标识符）锁定。启用此选项可防止其他进程创建新的时间点快照，从而使时间点之前的数据无效。

[-p] [-n] -u m volume_set
从影子卷更新主卷的时间点副本。*volume_set* 是包含主卷和影子卷的 Point-in-Time Copy 卷集。此选项提供主卷的快速重新同步，并创建影子卷的增量副本。此更新复制在主卷和影子卷之间标记为不同的所有 32 KB 段。它不会复制所有影子卷数据，而仅复制更改的数据。复制数据时，主卷依赖于影子卷。

使用此选项之前，暂时停止卷的工作负荷；并使主机应用程序停止写入到卷。这可以确保时间点数据一致。在使用目标卷执行任何其他操作之前，您可以使用 **iiadm -i volume_set** 以可视方式检查此副本或更新操作的状态，或者使用 **iiadm -w volume_set** 以交互方式检查状态（通过 shell 或脚本）。

此命令支持查询，以免意外覆盖主卷上的数据。当在脚本中使用此命令选项时，添加 **-n** 选项可以防止出现查询。

通过使用选项 **-p**、**iiadm -p -u m**，此命令支持 PID（Process Identifier，进程标识符）锁定。启用此选项可防止其他进程创建新的时间点快照，从而使时间点之前的数据无效。

[-p] -c s volume_set
将主卷复制到影子卷。

在影子卷上创建主卷的时间点副本。*volume_set* 是包含主卷和影子卷的 Point-in-Time Copy 卷集。此选项将主卷的时间点副本中的所有数据写入到影子卷。在将数据从主卷复制到影子卷时，影子卷依赖于主卷。

此选项执行完整卷复制。除非您怀疑独立影子卷上的数据的完整性，请使用 **iiadm -u s**。否则，使用此选项可同步主卷和影子卷；也即，使每个卷上的数据相一致。

使用此选项之前，暂时停止卷的工作负荷；并使主机应用程序停止写入到卷。这可以确保时间点数据一致。在使用目标卷执行任何其他操作之前，您可以使用 **iiadm -i volume_set** 以可视方式检查此副本或更新操作的状态，或者使用 **iiadm -w volume_set** 以交互方式检查状态（通过 shell 或脚本）。

通过使用 `-p` 选项、`iiadm -p -c s`，此命令支持 PID（Process Identifier，进程标识符）锁定。启用此选项可防止其他进程创建新的时间点快照，从而使时间点之前的数据无效。

`-c m volume_set`

将影子卷复制到主卷。

在主卷上创建影子卷的时间点副本。`volume_set` 是包含主卷和影子卷的 Point-in-Time Copy 卷集。此选项将影子卷的时间点副本中的所有数据写入到主卷。在将数据从影子卷复制到主卷时，主卷依赖于影子卷。

此选项执行完整卷复制。除非您怀疑独立主卷上的数据的完整性，请使用 `iiadm -u m`。否则，使用此选项可同步主卷和影子卷；也即，使每个卷上的数据相一致。

使用此选项之前，暂时停止卷的工作负荷；并使主机应用程序停止写入到卷。这可以确保时间点数据一致。在使用目标卷执行任何其他操作之前，您可以使用 `iiadm -i volume_set` 以可视方式检查此副本或更新操作的状态，或者使用 `iiadm -w volume_set` 以交互方式检查状态（通过 shell 或脚本）。

此命令支持查询，以免意外覆盖主卷上的数据。当在脚本中使用此命令选项时，添加 `-n` 选项可以防止出现查询。

通过使用 `-p` 选项、`iiadm -p -c m`，此命令支持 PID（Process Identifier，进程标识符）锁定。启用此选项可防止其他进程创建新的时间点快照，从而使时间点之前的数据无效。

`-a volume_set`

中止可能在主卷和影子卷之间处于活动状态的任何当前复制操作。`volume_set` 是包含主卷和影子卷的 Point-in-Time Copy 卷集。执行 `iiadm -a` 后，不会完成对目标卷（主卷或影子卷）的更新或复制操作。目标卷现在已成为源卷的从属副本。重新发出更新或复制命令选项，以便重新同步卷。

`[-p] [-n] -w volume_set`

等待任何正在进行的复制或更新操作完成或中止。`volume_set` 是包含主卷和影子卷的 Point-in-Time Copy 卷集。

此选项等待当前 Point-in-Time Copy 操作完成，从而防止执行后续 `iiadm` 命令（通过 shell 或脚本）。如果您希望确定复制或更新操作是否已完成，请使用此命令选项。

此命令支持 PID（Process Identifier，进程标识符）解除锁定。如果使用 `-p` 选项调用使用命令 `iiadm -p {-c|-u} {m|s}` 的早期复制或更新，完成等待处理后，如果当前 PID 是锁定时间点数据的 PID，此选项将解除锁定数据。

`-i volume_set`

显示 Point-in-Time Copy 当前启用或暂停的卷集的状态。`volume_set` 是包含主卷和影子卷的 Point-in-Time Copy 卷集。如果未指定 `volume_set`，则会显示配置的所有 Point-in-Time Copy 卷集的状态。

-l

列出当前配置的所有 Point-in-Time Copy 卷。

-O *overflow_vol*

此选项导致 Point-in-Time Copy 初始化指定 *overflow_vol*，以便在以后作为溢出卷与压缩从属影子卷配合使用。为了帮助有效共享使用此溢出卷，可以根据需要将溢出卷与多个 Point-in-Time Copy 卷集相关联。

在初始化 *overflow_vol* 时，此选项的启动器必须回答以下问题“是否要初始化此溢出卷？是/否”。必须回答“是/否”才能继续。

此选项支持 -n 选项，因此将执行请求的操作，而不会显示提示。此选项对包含在脚本中非常有用。必须首先指定 -n 选项。例如，`iiadm -nO vol` 有效；`iiadm -On vol` 无效。

确保您希望初始化指定 *overflow_vol* 上的数据，特别是当使用 -n 选项时。

-A *overflow_vol volume_set*

此选项支持在压缩从属影子卷太小或对卷集执行计划外的更改的情况下，在后续操作中将指定 *overflow_vol* 用作溢出卷。更改的溢出数据将重定向到关联的溢出卷。*volume_set* 是包含主卷和影子卷的 Point-in-Time Copy 卷集。

如果未初始化 *overflow_vol*，此选项将初始化 *overflow_vol*（请参见 -O 选项），然后将 *overflow_vol* 附加到 *volume_set*。

如果以前已初始化 *overflow_vol*，此选项会将 *overflow_vol* 附加到 *volume_set*。

此选项支持 -n 选项，因此将执行请求的操作，而不会显示提示。此选项对包含在脚本中非常有用。必须首先指定 -n 选项。例如，`iiadm -nA vol` 有效；`iiadm -An vol` 无效。

确保您希望初始化指定 *overflow_vol* 上的数据，特别是当使用 -n 选项时。

-D *volume_set*

此选项删除当前与指定 *volume_set* 相关联的溢出卷。如果 *volume_set* 当前正在使用溢出卷，此操作将失败，并返回“溢出卷仍在使用”错误消息。要解决此情况，请在 *volume_set* 中执行下面介绍的操作之一。这些操作会暂时清除与此卷集关联的所有溢出写入。

`abort(-a)`

中止复制操作。

`disable(-d)`

分解卷集。

`update(-u)`

更新卷集。

-L

此选项列出与一个或多个卷集关联的所有溢出卷。

-Q *overflow_vol*

此选项显示 *overflow_vol* 的当前状态。

-E *volume_set*

导出 *volume_set* 指定的 Point-in-Time Copy 卷集的独立影子卷。通过多端口设备等支持的技术，影子卷可供其他主机进行读取/写入访问。当影子卷重新加入到其主卷时，上述其他主机负责维护差异位图，此位图用于合并在本地记录的主卷的差异。导出影子卷时，影子卷不能受更新或复制操作的影响。请在调用导出命令之前执行 `iiadm -w volume_set` 命令。

-I *volume_set bitmap_vol*

导入 *volume_set* 指定的 Point-in-Time Copy 卷集的独立影子卷。通过多端口设备等支持的技术，以前必须已从主机导出影子卷。导入操作导致此主机在修改卷时开始维护差异位图。*bitmap_vol* 不应与影子卷在最初形成投影组时所使用的 *bitmap_vol* 相同。

在其他节点不再需要导出/导入的独立影子卷之后，在执行联接操作之前，必须输入禁用命令以使 *bitmap_vol* 及其关联 *shadow_vol* 保持一致。例如，

```
# iiadm -d volume_set
```

-J *volume_set bitmap_vol*

使用 *bitmap_vol* 联接 *volume_set* 和 Point-in-Time Copy 卷集的主卷集。读取提供的位图卷，并将其与原始卷合并，以便重建由主卷、影子卷和位图卷组成的原始卷集。要合并的 *bitmap_vol* 是在已导入独立影子卷的节点上获取的 *bitmap_vol*。自复制位图之时起，不能对导入计算机上的影子卷执行写入活动，直到重新导入影子卷为止。

-g *group_name* -m *volume_set* [*volume_set2* ...]

向用户指定的 *group_name* 添加一个或多个现有 Point-in-Time Copy *volume_set*。通过组中的一个或多个 Point-in-Time Copy 卷集的这种关联，可以将 *group_name* 中的所有卷集作为一个整体，统一执行下面显示的 `iiadm` 选项的列表。

以原子方式仅对组中的所有 Point-in-Time Copy 集执行命令 `COPY (-c)` 和 `UPDATE (-u)`。按顺序对每个组成员执行所有其他 `iiadm` 组命令。

`iiadm` 组命令的语法如下所示：

```
iiadm -g group_name [options]
```

options 如下所示：

-a

对 *group_name* 中的所有集中止复制操作。

-A

将 *overflow_vol* 附加到 *group_name* 中的所有集。

-c {s | m}

针对 *group_name* 中的所有集复制影子卷/主卷。

- D
从 *group_name* 中的所有集分离 *overflow_vol*。
- d
禁用 *group_name* 中的所有集。
- M
导出 *group_name* 中的所有卷集。
- i
group_name 中的所有卷集的状态。
- l
列出 *group_name* 中的所有卷集。
- L
列出所有组。
- n
不询问更新主卷是否是用户希望执行的实际操作。
- P
对 *group_name* 中的所有卷集设置参数。
- R
重置 *group_name* 中的所有卷集。
- u {s | m}
针对 *group_name* 中的所有集更新影子卷/主卷。
- w
等待 *group_name* 中的所有卷集。
- g - " " -m *volume_set* [*volume_set2* ...]
从卷集的当前关联 *group_name* 中删除一个或多个现有 Point-in-Time Copy *volume_set*。缺省情况下，或者在移至用户指定的 *group_name* 之前，所有 Point-in-Time Copy *volume_set* 都位于空 (" ") 组中。此关联允许针对空 (" ") *iiadm* *group_name* 执行以前记录的所有 *iiadm* 组命令。

-C *cluster_tag*
此 Point-in-Time Copy 选项是一个修饰符，可将配置操作限定为仅包含在 Sun Cluster 资源组或磁盘组中的卷。

在卷管理器可识别 Sun Cluster 的 Sun Cluster 中，*iiadm* 自动获取正确的磁盘组信息，因此，除非卷是包含的资源组的一部分，否则通常不需要使用此选项。

在只能在本地节点上访问卷的 Sun Cluster 中，*local* 的特殊 *cluster_tag* 用于表示不是 Sun Cluster 资源组或磁盘组的一部分的卷。

如果将“-L”指定为 *cluster_tag* 参数，*iiadm* 将列出与 Point-in-Time Copy 关联的所有群集标记。

当在未安装或配置 Sun Cluster 软件包的 Solaris 系统上使用时，此选项无效。

- h
输出 iiadm 用法汇总。
- v
显示 Point-in-Time Copy 软件组件的当前版本。

有关使用此部分中的其余命令的帮助，请与 Oracle Enterprise Services 联系。

- P *delay unit volume_set*
将指定 *volume_set* 的 Point-in-Time Copy 卷集调整参数更改为每个 *unit* I/O 的 *delay* 周期。*delay* 范围为 2 至 10000（含 2 和 10000）；*unit* 范围为 100 至 60000（含 100 和 60000）。
- R *volume*
卷失败后，Point-in-Time Copy 会将其置于脱机状态。在替换卷之后，请使用此选项将其恢复为联机状态。同时，还会将 Point-in-Time Copy 卷集中的关联从属卷置于联机状态。在将卷置于联机状态之后，此命令还会启动任何必要的时间点卷更新。

退出状态 0 命令已成功完成。
 >0 出现错误。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	storage/avs/avs-point-in-time-copy
接口稳定性	Committed（已确定）

另请参见 [dscfg\(1M\)](#)、[svadm\(1M\)](#)、[ds.log\(4\)](#)、[rdc.cf\(4\)](#)、[attributes\(5\)](#)、[ii\(7D\)](#)、[sv\(7D\)](#)

引用名	iicpbmp – 复制 Availability Suite Point-In-Time 位图卷
用法概要	iicpbmp [-c] <i>old_bitmap new_bitmap</i> ...
描述	iicpbmp 命令复制 Availability Suite Point-in-Time 位图卷，并重写位图标头以使其与新的位图卷名称一致。将重写投影集的配置条目以反映新位图的位置。 不会检查当前使用的是旧位图卷，还是新位图卷。仅当使用旧位图的 Point-In-Time Copy 投影集被挂起时，才应运行 iicpbmp 命令。
选项	-c 不尝试为使用该位图的 Point-in-Time 投影集更新 Availability Suite 配置。此选项将生成位图的副本，但不会影响使用旧位图卷的投影集。
操作数	<i>old_bitmap new_bitmap</i> 旧的和新的 Point-In-Time 位图卷。
警告	仅应在系统处于单用户模式时运行 iicpbmp。iicpbmp 在进行复制时不会尝试检查 Point-In-Time Copy 集是否正在使用中。如果在 Point-In-Time Copy 正在使用投影集时在未使用 -c 标志的情况下运行了 iicpbmp，这将导致在 Point-In-Time Copy 下次启动时投影集中发生冲突。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	storage/avs/avs-point-in-time-copy
接口稳定性	Committed（已确定）

另请参见 [iiadm\(1M\)](#)、[iicpshd\(1M\)](#)、[attributes\(5\)](#)

- 引用名

iicpshd – 复制 Availability Suite Point-in-Time 影子卷
- 用法概要

iicpshd [-s] *old_shadow new_shadow...*
- 描述

iicpshd 命令复制 Availability Suite Instant Image 影子卷，并更新位图标头和 Availability Suite 配置以反映新的影子卷。

不会检查当前使用的是旧影子卷，还是新影子卷。仅当使用旧影子卷的 Instant Image 投影集被挂起时，才应运行 iicpshd 命令。
- 选项

iicpshd 命令支持以下选项：

-s 为 Point-in-Time 投影集更新 StorageTek 配置信息，但不将旧影子卷的数据复制到新影子卷中。
- 操作数

iicpshd 命令行具有以下操作数：

old_shadow new_shadow

iicpshd 将旧 Availability Suite Instant Image 影子卷的数据复制到新影子卷并更新位图标头和配置数据。
- 属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	storage/avs/avs-point-in-time-copy
接口稳定性	Committed（已确定）

另请参见

[iiadm\(1M\)](#)、[iicpbmp\(1M\)](#)、[attributes\(5\)](#)

警告

仅应在系统处于单用户模式时运行 iicpshd。当运行 iicpshd 时，该命令不会尝试检查 Point-in-Time 集是否正在使用中。如果在 Point-in-Time 正在使用旧影子卷时在未使用 -s 标志的情况下运行了 iicpshd，这将导致影子卷数据丢失。如果使用了 -s 选项，必须手动将旧影子卷上的数据复制到新影子卷中。

引用名	ikedadm – 处理 Internet 密钥交换 (Internet Key Exchange, IKE) 参数和状态
用法概要	<pre>ikedadm [-np] ikedadm [-np] get [debug priv stats defaults] ikedadm [-np] set [debug priv] [level] [file] ikedadm [-np] [get del] [pl rule preshared] [id] ikedadm [-np] add [rule preshared] { <i>description</i> } ikedadm [-np] token [login logout] <i>PKCS#11_Token_Object</i> ikedadm [-np] [read write] [rule preshared certcache] <i>file</i> ikedadm [-np] dump [pl rule preshared certcache groups encrals authalgs] ikedadm [-np] flush [pl certcache] ikedadm help [get set add del read write dump flush token]</pre>
描述	ikedadm 实用程序从 Internet 密钥交换 (Internet Key Exchange, IKE) 协议守护进程 in.iked(1M) 检索信息，并处理此守护进程的配置。 ikedadm 支持一系列操作，可对支持的一个或多个对象类型执行这些操作。当在不带任何参数的情况下调用时，ikedadm 将进入交互模式，这会在标准输出中输出一个提示，并接受来自标准输入的命令，直到到达文件结尾为止。 由于 ikedadm 处理敏感加密信息，因此您必须是超级用户才能使用此命令。此外，某些可用命令还要求在特权模式下运行此守护进程，此模式在启动守护进程时建立。 有关如何安全地使用此命令的详细信息，请参见安全。
选项	所支持的选项如下： -n 防止在报告操作时尝试以符号方式输出主机和网络名称。此选项非常有用，例如，当关闭或无法访问所有名称服务器时。 -p 偏执。不输出任何加密材料，即使在保存安全关联时。启用此标志时输出 x，而非实际的十六进制数字。
用法	
命令	所支持的命令如下：

add

添加指定对象。此选项可用于向当前（正在运行的）`in.iked` 配置添加新的策略规则或新的预先共享的密钥。当添加新的预先共享的密钥时，由于此命令将包含加密材料，因此无法从命令行调用此命令。添加的规则或密钥使用相应的 ID-值对指定，如 **ID 格式** 部分中所述。

del

从 `in.iked` 的当前配置中删除一个或多个特定对象。此操作适用于 IKE（阶段 1）SA、策略规则和预先共享的密钥。按照 **Id 格式** 中的说明指定要删除的对象。

dump

显示 `in.iked` 已知的指定类型的所有对象。此选项可用于显示所有阶段 1 SA、策略规则、预先共享的密钥、实现的 Diffie-Helman 组、可用于阶段 1 的加密和验证算法或证书高速缓存。此命令可能会生成大量输出。

flush

从 `in.iked` 中删除所有 IKE（阶段 1）SA 或缓存证书。

请注意，刷新**证书高速缓存**还将使用添加或删除的任何新证书更新 IKE（作为负面影响）。

get

查找并显示指定对象。可用于查看当前调试或特权级别、守护进程的全局统计信息和缺省值、特定 IKE（阶段 1）SA、策略规则或预先共享的密钥。后面三种对象类型要求传入标识信息；下面介绍了每种对象类型的相应规范。

help

输出简要的命令摘要，或者当后跟有命令时，输出有关此命令的信息。

read

通过从缺省位置或指定文件读取策略规则或预先共享的密钥，更新当前 `in.iked` 配置。

set

调整当前调试或特权级别。如果要修改调试级别，可以选择性地指定输出文件；如果守护进程正在后台运行，并且当前未输出到文件，**必须**指定输出文件。更改特权级别时，调整只能是降低访问级别；不能使用 `ikeadm` 提高访问级别。

write

将当前 `in.iked` 策略规则集或预先共享的密钥集写入到指定文件。必须指定目标文件。此命令不能用于覆盖现有配置文件。

token

登录 PKCS#11 令牌对象并授予对加密材料的访问，或者注销并使对加密材料的访问无效。

`token` 可作为具有下列授权的普通用户运行：

- 令牌登录：`solaris.network.ipsec.ike.token.login`
- 令牌注销：`solaris.network.ipsec.ike.token.logout`

对象类型

debug
指定守护进程的调试级别。此选项确定守护进程提供的有关其操作的输出的数量和类型。调试级别实际为位掩码，并且各个位支持不同类型的信息。

说明	标志	昵称
凭证签发管理	0x0001	cert
密钥管理	0x0002	key（用户密钥）
可操作的	0x0004	op
创建阶段 1 SA	0x0008	phase1
创建阶段 2 SA	0x0010	phase2
PF_KEY 接口	0x0020	pfkey
策略管理	0x0040	策略
提议的建构	0x0080	prop
门接口	0x0100	door
配置文件处理	0x0200	配置
标签处理	0x0400	label
所有调试标志	0x07ff	all

当指定调试级别时，可指定编号（十进制或十六进制）或昵称字符串。例如，88、0x58 和 phase1+phase2+policy 均等效，并且将启用调试以便创建 phase 1 sa、phase 2 sa 和策略管理。此外，还可以使用昵称字符串删除特定类型的信息；all-op 的作用是启用除操作消息以外的所有调试；它等效于编号 1019 或 0x3fb。

priv
指定守护进程的访问特权级别。可能值有：

Description	Level	Nickname
Base level	0	base
Access to preshared key info	1	modkeys
Access to keying material	2	keymat

缺省情况下，在基本级别启动 in.iked。命令行选项可用于在更高级别启动守护进程。ikeadm 可用于降低级别，但不能用于提升级别。

数字级别或昵称可用于指定目标特权级别。

为了获取、添加、删除、转储、读取或写入预先共享的密钥，特权级别必须至少能够授予对预先共享的密钥信息的访问。但是，当查看预先共享的密钥（使用获取或

转储命令) 时, 仅当特权级别能够授予对加密材料的访问时, 才能使用密钥本身。当查看阶段 1 SA 时, 也会出现此种情况。

stats

守护进程的全局统计信息, 涵盖成功和失败的阶段 1 SA 创建。

报告的统计信息包括:

- 本地实体启动的当前 P1 SA 的计数
- 本地实体作为响应器的当前 P1 SA 的计数
- 本地实体自引导以来启动的所有 P1 SA 的计数。
- 本地实体自引导以来作为响应器的所有 P1 SA 的计数
- 自引导以来尝试的所有 P1 SA 的计数 (其中本地实体是启动器); 包括失败的尝试
- 自引导以来尝试的所有 P1 SA 的计数 (其中本地实体是响应器); 包括失败的尝试
- 启动 P1 SA 的所有失败尝试的计数, 出现失败的原因是因为同级设备未响应
- 启动 P1 SA 的所有失败尝试的计数, 其中统计设备已响应
- 同级设备作为启动器的所有失败 P1 SA 的计数
- 是否使用了 PKCS#11 库以及装载的 PKCS#11 库 (如果适用)。请参见示例 11。

defaults

显示 `in.iked` 守护进程使用的缺省值。某些值可在守护进程配置文件中覆盖 (请参见 `ike.config(4)`); 对于这些值, 令牌名称显示在 `get defaults` 输出中。此输出将反映配置令牌已更改缺省值的位置。

如果同级系统提出有效的备用提案, 或者可根据 `ike.config` 中建立的基于规则的值覆盖缺省值, 则可以忽略缺省值。此种情况下, `get defaults` 命令继续显示缺省值, 而不是用于覆盖缺省值的值。

p1

IKE 阶段 1 SA。p1 对象通过 IP 地址对或 cookie 对标识; 标识格式如下所述。

rule

IKE 策略规则, 定义位于指定的本地标识和远程标识之间的阶段 1 SA 的可接受安全特征。规则通过其标签标识; 标识格式如下所述。

preshared

预先共享的密钥, 包括本地和远程标识以及适用的 IKE 模式。预先共享的密钥通过 IP 地址对或标识对标识; 标识格式如下所述。

Id 格式

`add`、`del` 和 `get` 等命令要求在命令行中指定其他信息。对于删除和获取命令, 所有所需操作仅为以最简单的方式标识给定对象; 对于添加命令, 必须指定完整对象。

在大多数情况下, 使用一对值可以实现最简单的标识。对于 IP 地址, 以点记法 (适用于 IPv4 地址)、冒号分隔的十六进制格式 (适用于 IPv6 地址) 或主机名称数据库中不存在的主机名的形式先后指定本地地址和远程地址。如果指定扩展为多个地址的主机名, 则将多次执行请求的操作 (对每个可能的地址组合各执行一次)。

标识对由后跟有远程类型-值对的本地类型-值对组成。有效类型包括：

prefix

地址前缀。

fqdn

全限定域名。

domain

域名，fqdn 的同义词。

user_fqdn

user@fqdn 格式的用户标识。

mailbox

user_fqdn 的同义词。

cookie 对由两个 cookie 组成，在创建阶段 1 安全关联 (Security Association, SA) 时会把这两个 cookie 分配给此关联；第一个是启动器的 cookie，然后是响应器的 cookie。cookie 是一个 64 位数字。

最后，标签（用于标识策略规则）是在创建规则时分配给规则的字符串。

为添加命令格式化规则或预先共享的密钥时应遵循 in.iked 配置文件的格式规则。两种规则都由一系列包含在花括号（{ 和 }）中的 ID-值对组成。有关格式化规则和预先共享的密钥的详细信息，请参见 [ike.config\(4\)](#) 和 [ike.preshared\(4\)](#)。

安全

ikeadm 命令允许特权用户输入加密密钥信息。如果对手获取了对此类信息的访问，则会影响 IPsec 通信的安全。使用 ikedadm 命令时，应考虑下列问题。

- TTY 是否会通过网络（交互模式）？

如果是，加密材料的安全即为此 TTY 的通信的网络路径的安全。使用 ikedadm（而非明文 telnet 或 rlogin 会话）具有一定的风险。甚至本地窗口也可能受到读取存在的窗口事件的隐藏程序的攻击。

- 文件是否能通过网络访问，或者是否可供用户读取（读/写命令）？

挂载在网络中的文件可在读取时被对手发现。网络中带有加密材料的用户可读文件同样具有一定的风险。

如果您的源地址是可通过网络查找的主机，并且您的命名系统本身已泄密，使用的所有名称都将不再值得信任。

安全漏洞通常在工具的不当应用中产生，而不存在于工具自身中。建议管理员在使用 ikedadm 命令时应小心谨慎。最安全的运行模式可能是在控制台或其他硬连接 TTY 中运行。

有关此主题的其他信息，请参见 Matt Blaze 在 Bruce Schneier 编著的《Applied Cryptography: Protocols, Algorithms, and Source Code in C.》中撰写的后记。

示例

示例1 清空所有阶段1安全关联

下面的命令清空所有阶段1安全关联：

```
example# ikeadm flush p1
```

示例2 显示所有阶段1安全关联

下面的命令显示所有阶段1安全关联：

```
example# ikeadm dump p1
```

示例3 删除特定阶段1安全关联

下面的命令删除指定阶段1安全关联：

```
example# ikeadm del p1 local_ip remote_ip
```

示例4 从文件中添加规则

下面的命令用于从文件中添加规则：

```
example# ikeadm add rule rule_file
```

示例5 添加预先共享的密钥

下面的命令添加预先共享的密钥：

```
example# ikeadm  
ikeadm> add preshared { localidtype ip localid local_ip  
         remoteidtype ip remoteid remote_ip ike_mode main  
         key 1234567890abcdef1234567890abcdef }
```

示例6 将所有预先共享的密钥保存到文件

下面的命令将所有预先共享的密钥保存到文件：

```
example# ikeadm write preshared target_file
```

示例7 查看特定规则

下面的命令用于查看特定规则：

```
example# ikeadm get rule rule_label
```

示例8 从ike.config读入新规则

下面的命令从ike.config文件读入新规则：

```
example# ikeadm read rules
```

示例9 降低特权级别

下面的命令降低特权级别：

```
example# ikeadm set priv base
```

示例 10 查看调试级别
下面的命令显示当前调试级别：

```
example# ikeadm get debug
```

示例 11 使用统计信息检验硬件加速器
下面的示例显示统计信息如何在末尾包含可选行，以便指示 IKE 是否使用 PKCS#11 库来加快公钥操作的速度（如果适用）。

```
example# ikeadm get stats
Phase 1 SA counts:
Current:  initiator:      0    responder:      0
Total:    initiator:    21    responder:    27
Attempted: initiator:    21    responder:    27
Failed:   initiator:      0    responder:      0
           initiator fails include 0 time-out(s)
PKCS#11 library linked in from /opt/system/core-osonn/lib/libpkcs11.so
example#
```

示例 12 显示证书高速缓存
下面的命令显示证书高速缓存以及关联私钥的状态（如果适用）：

```
example# ikeadm dump certcache
```

示例 13 登录 PKCS#11 令牌
下面的命令显示如何登录 PKCS#11 令牌对象以及如何解除锁定私钥：

```
example# ikeadm token login "Sun Metaslot"
Enter PIN for PKCS#11 token:
ikeadm: PKCS#11 operation successful
```

退出状态 将返回下列退出值：

- 0 成功完成。
- 非零 出现错误。在标准错误中写入相应错误消息。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Not an Interface（不是接口）

另请参见 [in.iked\(1M\)](#)、[ike.config\(4\)](#)、[ike.preshared\(4\)](#)、[attributes\(5\)](#)、[ipsec\(7P\)](#)

Bruce Schneier 编著的《Applied Cryptography: Protocols, Algorithms, and Source Code in C》第 2 版，John Wiley & Sons 出版，纽约，1996 年。

附注

由于 `in.iked` 只能在全局区域和专用 IP 区域中运行，因此此命令在共享 IP 区域中不起作用。

引用名

ikecert – 处理计算机文件系统上的公钥证书数据库

用法概要

```
ikecert certlocal
    [-a | -e | -h | -k | -l | -r | -U | -C | -L]
    [[-p] -T PKCS#11 token identifier]
    [option_specific_arguments]...

ikecert certdb [-a | -e | -h | -l | -r | -U | -C | -L]
    [[-p] -T PKCS#11 token identifier]
    [option_specific_arguments]...

ikecert certrldb [-a | -e | -h | -l | -r]
    [option_specific_arguments]...

ikecert tokens
```

描述

ikecert 命令处理计算机文件系统上的公钥证书数据库。请参见下面的“文件”部分。

ikecert 有三个子命令分别对应于三个主要系统信息库，还有一个子命令用于列出可用硬件令牌：

- certlocal，用于处理私钥系统信息库
- certdb，用于处理公钥系统信息库
- certrldb，用于处理证书撤销列表 (certificate revocation list, CRL) 系统信息库。
- tokens，用于显示给定 PKCS#11 库的可用 PKCS#11 令牌。

支持的唯一 PKCS#11 库和硬件是 Sun Cryptographic Accelerator 4000。

选项

除 tokens 以外，每个子命令都需要一个选项，后面可能跟有一个或多个特定于选项的参数。

tokens 子命令列出 /etc/inet/ike/config 中指定的 PKCS#11 库中的所有可用令牌。

所支持的选项如下：

-a

certlocal

当随 certlocal 子命令一起指定时，此选项将私钥安装（添加）到 Internet 密钥交换 (Internet Key Exchange, IKE) 本地 ID 数据库中。密钥数据从标准输入读取，并且采用仅限 Solaris 的格式或未加密的 PKCS#8 DER 格式。自动检测密钥格式。无法识别 PEM 格式的 PKCS#8 密钥文件和受口令保护的加密格式的文件，但可以使用 OpenSSL 中提供的工具对这些文件进行相应转换。

当对应的公共证书已不在 IKE 数据库中时，此选项不能用于 PKCS#11 硬件对象。当同时导入公共证书和私钥时，必须首先使用 `certdb` 子命令导入公共证书。

`certdb`

当随 `certdb` 子命令一起指定时，此选项从标准输入读取证书，并将其添加到 IKE 证书数据库。该证书必须是采用 PEM Base64 或 ASN.1 BER 编码的 X.509 证书。证书采用其标识名称。

此选项可采用下列两种方法之一将证书导入到 PKCS#11 硬件密钥存储：使用 `certlocal -kc` 选项创建匹配的公钥对象和现有私钥对象，或者如果明确指定了 PKCS#11 令牌，则使用 `-t` 选项。

`certrlb`

当随 `certrlb` 子命令一起指定时，此选项将 CRL 安装（添加）到 IKE 数据库中。CRL 从标准输入读取。

`-e [-f pkcs8] slot`

`certlocal`

当随 `certlocal` 子命令一起指定时，此选项从 IKE 本地 ID 数据库中提取私钥。密钥数据写入到标准输出。插槽指定要提取的私钥。私钥仅以二进制/ber 格式提取。

使用此选项时应格外小心。请参见下面的“安全”部分。

此选项不适用于 PKCS#11 硬件对象。

当与 `-f pkcs8` 结合使用时，将以未加密的 PKCS#8 格式提取私

`-e [-f output-format] certspec`

`-kc -m keysize -t keytype -D dtype -A altname [...]
 [-S validity start_time][-F validity end_time]
 [-T PKCS#11 token identifier]`

钥。

certdb

当随 `certdb` 子命令一起指定时，此选项从与 `certspec` 匹配的 IKE 证书数据库中提取证书，并将其写入到标准输出。 `output-format` 选项指定编码格式。有效选项包括 PEM 和 BER。它会提取第一个匹配的标识。缺省输出格式为 PEM。

certldb

当随 `certldb` 子命令一起指定时，此选项从 IKE 数据库提取 CRL。密钥数据写入到标准输出。 `certspec` 指定提取的 CRL。将提取数据库中的第一个匹配 CRL。有关 `certspec` 模式的详细信息，请参见下面的“说明”。

certlocal

当随 `certlocal` 子命令一起指定时，此选项生成一个 IKE 公钥/私钥对，并将其添加到本地 ID 数据库。此外，它还会生成证书请求，并将其发送到标准输出。有关上述选项的详细信息，请参见[附注](#)以获取有关 `dtype` 参数的详细信息，并参见“替代名称”了解有关此命令的 `altname` 参数的详细信息。

如果指定了 `-T`，硬件令牌将生成密钥对。

如果与 `-T` 一起指定了 `-p`，则会将 PKCS#11 令牌个人识别码存储在明确磁盘上，并具有受 `root` 保护的权限。如果未指

定，用户必须在运行 `in.iked(1M)` 时使用 `ikeadm(1M)` 解除令牌的锁定。

```
-ks -m keysize -t keytype -D dname -A altname[ ... ]
[-S validity_start_time][-F validity_end_time]
[-f output-format][[-p] -T PKCS#11 token identifier]
```

certlocal

当随 `certlocal` 子命令一起指定时，此选项生成一个公钥/私钥对，并将其添加到本地 ID 数据库。此外，此选项还会生成一个自签名证书，并将其安装到证书数据库。有关此命令的 `dname` 和 `altname` 参数的详细信息，请参见下面的“说明”部分。

如果指定了 `-T`，硬件令牌将生成密钥对，并将自签名证书存储到此硬件中。

```
-l [-v] [slot]
```

certlocal

当随 `certlocal` 子命令一起指定时，此选项列出本地 ID 数据库中的私钥。`-v` 选项将输出切换到详细模式，在此模式下将输出完整证书。

使用 `-v` 选项时应格外小心。请参见下面的“安全”部分。`-v` 选项不适用于 PKCS#11 硬件对象。

```
-l [-v] [certspec]
```

certdb

当随 `certdb` 子命令一起指定时，此选项列出 IKE 证书数据库中 `certspec` 匹配的证书（如果指定了任何模式）。此列表显示证书的标识字符串以及私钥（如果位于密钥数据库中）。`-v` 将输出切换到详细模式，在此模式下将输出完整证书。

如果匹配的证书位于硬件令牌上，则还会列出令牌 ID。

certldb

当随 **certldb** 子命令一起指定时，此选项列出 IKE 数据库中的 CRL 以及位于此数据库中且与颁发机构名称匹配的所有证书。**certspec** 可用于指定列出特定 CRL。**-v** 选项将输出切换到详细模式，在此模式下将输出完整证书。有关 **certspec** 模式的详细信息，请参见下文的“附注”部分。

-r slot

certlocal

当随 **certlocal** 子命令一起指定时，此选项删除指定插槽中的本地 ID。如果存在对应的公钥，则不会删除此 ID。如果此插槽被视为“已损坏”或无法识别，则也会删除此 ID。

如果对 PKCS#11 硬件对象调用了此选项，则还会删除 PKCS#11 公钥和私钥对象。如果已使用 **certdb -r** 删除公钥对象，则不会导致问题。

-r certspec

certdb

从 IKE 证书数据库删除证书。会删除与指定证书模式匹配的证书。不会删除 **certlocal** 数据库中与此些证书对应的任何私钥。这会删除第一个匹配的标识。

如果模式指定了一个插槽，并且此插槽被视为“已损坏”或无法识别，则也会删除该插槽。

如果对 PKCS#11 硬件对象调用了此选项，则还会删除证书和

PKCS#11 公钥对象。如果已使用 `certlocal -r` 删除公钥对象，则不会导致问题。

`certldb`

当随 `certldb` 子命令一起指定时，此选项删除具有给定 `certspec` 的 CRL。

`-U slot`

`certlocal`

当随 `certlocal` 子命令和 `-T` 标志一起指定时，此选项会解除 PKCS#11 私钥对象与 IKE 数据库之间的链接。系统不会尝试访问硬件 Keystore 或者验证或删除令牌上的私钥对象，而仅仅取消此对象与 IKE 数据库之间的关联。

`certdb`

当随 `certdb` 子命令和 `-T` 标志一起指定时，此选项会解除 PKCS#11 证书对象与 IKE 数据库之间的链接。系统不会尝试访问硬件 Keystore 或者验证或删除令牌上的证书或公钥对象，而仅仅取消这些对象与 IKE 数据库之间的关联。

`-C certspec`

`certlocal`

当随 `certlocal` 子命令一起指定时，此选项从磁盘上的 Keystore 将私钥（及其相应证书）和公钥复制到其 PKCS#11 令牌指定的硬件 Keystore。此子命令尝试创建其中每个组件，即使某个组件创建失败也是如此。在所有情况下，仍会保留磁盘上的原始私钥和公共证书，因此必须单独删除它们。某些硬件 Keystore（例如，符合 FIPS-140 的设备）可能不支持以此种方式迁移私钥对象。

-L pattern

certdb

当随 `certdb` 子命令一起指定时，此选项将与给定 `certspec` 匹配的证书及相应公钥从磁盘上的 `Keystore` 复制到其 `PKCS#11` 令牌指定的硬件 `Keystore`。系统仍会保留原始公共证书，因此必须根据需要单独删除这些证书。

如果指定了 `-p`，则会将 `PKCS#11` 令牌个人识别码存储在明确磁盘上，并具有受 `root` 保护的文件权限。如果未指定，用户必须在运行 `in.iked(1M)` 时使用 `ikeadm(1M)` 解除令牌的锁定。

certlocal

当随 `certlocal` 子命令一起指定时，此选项将现有令牌上的私钥对象链接到 `IKE` 数据库。此对象本身保留在令牌中。此选项只是让 `IKE` 基础结构了解对象是否存在，就像最初使用 `Solaris IKE` 实用程序在令牌上创建了此对象一样。

certdb

当随 `certdb` 子命令一起指定时，此选项将现有令牌上的证书对象链接到 `IKE` 数据库。此对象本身保留在令牌中。此选项只是让 `IKE` 基础结构了解对象是否存在，就像最初使用 `Solaris IKE` 实用程序在令牌上创建了此对象一样。

如果指定了 `-p`，则会将 `PKCS#11` 令牌个人识别码存储在明确磁盘上，并具有受 `root` 保护的文件权限。如果未指定，用户必须在运行 `in.iked(1M)` 时使用 `ikeadm(1M)` 解除令牌的锁定。

参数

支持的参数如下：

certspec

指定与证书规范匹配的模式。有效 certspec 包括“主题名称”、“颁发机构名称”和“主题替代名称”。

上述各项可指定与给定 certspec 值匹配且与其他 certspec 值匹配的证书。要指定不应存在于证书中的 certspec 值，请在该标记前加 !。

有效 certspec 包括：

```
<Subject Names>
SUBJECT=<Subject Names>
ISSUER=<Issuer Names>
SLOT=<Slot Number in the certificate database>
```

Example: "ISSUER=C=US, O=SUN" IP=1.2.3.4 !DNS=example.com

Example: "C=US, O=CALIFORNIA" IP=5.4.2.1 DNS=example.com

替代名称的有效参数如下：

```
IP=<IPv4 address>
DNS=<Domain Name Server address>
EMAIL=<email (RFC 822) address>
URI=<Uniform Resource Indicator value>
DN=<LDAP Directory Name value>
RID=<Registered Identifier value>
```

可以指定不带关键字标记的有效插槽编号。也可以发布具有关键字标记的替代名称。

-A

证书的主体替代名称。-A 选项后的参数应采用 **标记=值** 格式。有效标记包括 IP、DNS、EMAIL、URI、DN 和 RID（请参见下面的示例）。

-D

证书主题的 X.509 标识名。它通常具有以下格

式：C=country, O=organization, OU=organizational unit, CN=common name。有效标记包括：C、O、OU 和 CN。

-f

编码输出格式，包括 pem（表示 PEM Base64）或 ber（表示 ASN.1 BER）。如果未指定 -f，则采用 pem。

-F *validity_end_time*

证书有效期结束时间。如果未指定 -F 标志，有效期结束时间则为自有效期开始时间起的四年后。有关有效日期和时间语法的说明，请参见“说明”部分。

-m

密钥大小。此大小可以为 512、1024、2048、3072 或 4096。使用下面的命令可确定 Solaris 加密框架支持的密钥大小：

```
% cryptoadm list -vm
```

[pkcs11_softtoken\(5\)](#) 中介绍了上面的命令所显示的机制。如果您的系统具有硬件加速，该硬件支持的机制将在各提供器的单独部分中列出。机制可以是下列机制之一：

```
CKM_RSA_PKCS_KEY_PAIR_GEN
CKM_DSA_KEY_PAIR_GEN
CKM_DH_PKCS_KEY_PAIR_GEN
```

注 – 某些硬件仅支持某些密钥大小。例如，Sun Cryptographic Accelerator 4000 的 Keystore（当使用下面的 `-T` 选项时）最多仅对 RSA 支持 2048 位密钥，并对 DSA 支持 1024 位密钥。

`-S validity_start_time`

证书有效期开始时间。如果未指定 `-S` 标志，则将当前日期和时间用作有效期开始时间。有关有效日期和时间语法的说明，请参见下面的“说明”部分。

`-t`

密钥类型。密钥类型可以是 `rsa-sha1`、`rsa-md5` 或 `dsa-sha1`。

`-T`

硬件密钥存储的 PKCS#11 令牌标识符。此选项指定符合 PKCS#11 标准的硬件设备实例。PKCS#11 库必须在 `/etc/inet/ike/config` 中指定。（请参见 [ike.config\(4\)](#)。）

令牌标识符是包含 32 个字符的由空格补充的字符串。如果给定令牌的长度小 32 个字符，则会使用空格自动填充此令牌。

如果系统中存在多个 PKCS#11 库，请记住一次只能在 `/etc/inet/ike/config` 中指定一个 PKCS#11 库。一个 PKCS#11 库实例可能对应多个令牌（每个令牌具有单独的密钥存储）。

安全

此命令可将公钥-密钥对的私钥保存到文件中。如果恶意方通过某种方式获取了私钥，以任何途径泄露私钥都可能会导致攻击。

PKCS#11 硬件对象功能可克服磁盘上私钥的某些缺点。由于 IKE 是系统设备，因此不需要在引导时进行用户介入。但是，仍需要令牌的识别码。因此，PKCS#11 令牌的识别码存储在磁盘上的加密密钥通常所在的位置。通过使用硬件密钥存储，仍然无法拥有密钥，只有在主机面临安全风险时使用密钥才会导致问题，因此此设计决策可被接受。若不使用识别码，ikecert 安全性会降低为 PKCS#11 实现的安全性。因此，还应当仔细检查 PKCS#11 实现。

有关其他信息，请参见由 Matt Blaze 撰写的后文（摘自 Bruce Schneier 编著的 *Applied Cryptography: Protocols, Algorithms, and Source Code in C*）。

示例

示例 1 生成自签名证书

下面是自签名证书的示例：

示例1 生成自签名证书 (续)

```
example# ikecert certlocal -ks -m 512 -t rsa-md5 -D "C=US, O=SUN" -A
IP=1.2.3.4
Generating, please wait...
Certificate generated.
Certificate added to database.
-----BEGIN X509 CERTIFICATE-----
MIIBRDCB76ADAgECAGEBMA0GCSqGSIb3DQEBAUAMBsxCzAJBgNVBAYTAlVTMQww
CgYDVQQKEwNTVU4wHhcNMDEwMzE0MDEzMDEwMDEwMzE0MDEzMDEwMDEwMzE0MDEw
CQYDVQQGEwJVUzEMMAoGA1UEChMDU1VOMFowDQYJKoZIhvcNAQEBBQADSAwRgJB
APDhpqKgJgRoRUR6twTMTtSuNsReEnFoReVer!ztpXpQK6ybYLRH18JIQu/uCV/r
26R/cVXTy5qc5NbMwA40KzcCASOjIDAEmaSgA1UdDwQEAwIFoDAPBgNVHRECDAG
hwQBAgMEMA0GCSqGSIb3DQEBAUAA0EApTRD23KzN95GMvPD71hwwClukslKLVg8
f1xm9ZsHLPJLRxHFwsqqjAad4j4wwwriiUmGAHLTGB0lJm18xsxgag==
-----END X509 CERTIFICATE-----
```

示例2 生成 CA 请求

生成 CA 请求的语法与生成自签名证书的基本相同。二者之间的唯一区别在于使用选项 `-c` 而非 `-s`，并且证书数据为 CA 请求。

```
example# ikecert certlocal -kc -m 512 -t rsa-md5 \
-D "C=US, O=SUN" -A IP=1.2.3.4
```

示例3 使用硬件密钥存储的 CA 请求

下面的示例展示如何使用 `-T` 选项指定令牌。

```
example# # ikecert certlocal -kc -m 1024 -t rsa-md5 -T vca0-keystore \
-D "C=US, O=SUN" -A IP=1.2.3.4
```

退出状态

将返回下列退出值：

0

成功完成。

非零

出现错误。在标准错误中写入相应错误消息。

文件

/etc/inet/secret/ike.privatekeys/*

私钥。私钥**必须**在 /etc/inet/ike/publickeys/ 中具有匹配的同名公钥证书。

/etc/inet/ike/publickeys/*

公钥证书。这些名称仅对匹配私钥名称至关重要。

/etc/inet/ike/crls/*

公钥证书撤销列表。

/etc/inet/ike/config

参考此文件以获取 PKCS#11 库的路径名。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）

另请参见 [ikeadm\(1M\)](#)、[in.iked\(1M\)](#)、[getdate\(3C\)](#)、[ike.config\(4\)](#)、[attributes\(5\)](#)、[pkcs11_softtoken\(5\)](#)

由 Bruce Schneier 编著的Applied Cryptography: Protocols, Algorithms, and Source Code in C第 2 版。John Wiley & Sons 出版。纽约。1996 年。

由 RSA Labs 编著，PKCS#11 v2.11：Cryptographic Token Interface Standards，2001 年 11 月。

附注 下面提供了使用 -F 或 -S 标志时的有效日期和时间语法。

对于相对日期，此语法如下所示：

{+, -}[Ns][Nm][Nh][Nd][Nw][NM][Ny]

其中：

N
表示整数

s
表示秒

m
表示分钟

h
表示小时

d
表示天

w
表示星期

M
表示月

y
表示年

这些参数可按任意顺序指定。例如，“+3d12h”表示从现在开始三天半后的时间，“-3y2M”表示三年两个月以前。

所有具有固定值的参数都可以以绝对秒数的形式相加。具有可变秒数的月和年使用日历时间计算。长度不固定的月和年定义为，加一年或一个月表示下一年或下个月的同一天。例如，如果现在为 2005 年 1 月 26 日，并且证书应在从今天起的 3 年零 1 个月 after 过期，那么到期日期（有效期结束时间）将为 2008 年 2 月 26 日。系统会相应地处理溢出。例如，从 2005 年 1 月 31 日起一个月后的时间为 2005 年 3 月 3 日，这是因为 2 月仅包含 28 天。

对于绝对日期，接受文件 `/etc/datemsk` 中包含的日期格式的语法（有关详细信息，请参见 [getdate\(3C\)](#)）。前面带有“+”或“-”的任何日期字符串均被视为相对于当前时间的时间，而其他字符串则视为绝对日期。此外，系统还会执行合理性检查，确保有效期结束日期大于有效期开始日期。例如，下面的命令将创建一个证书，使其开始日期为 1 天零 2 小时以前，结束日期为当地时间 2007 年 1 月 22 日 12:00:00。

```
# ikecert certlocal -ks -t rsa-sha1 -m 1024 \  
-D "CN=mycert, O=Sun, C=US" \  
-S -1d2h -F "01/22/2007 12:00:00"
```

由于 [in.iked\(1M\)](#) 只能在全局区域和专用 IP 区域中运行，因此该命令在共享 IP 区域中不起作用。

引用名

ilbadm – 建立和处理负载均衡规则

用法概要

```

ilbadm create-rule [-e] [-p] -i vip=value,port=value[,protocol=value]
    -m lbalg=value,type=value[,proxy-src=ip-range][,pmask=mask]
    [-h hc-name=value[,hc-port=value]]
    [-t [conn-drain=N][,nat-timeout=N],[persist-timeout=N]]
    -o servergroup=value name

ilbadm show-rule [-e|-d] [-f |[-p] -o key[,key ...]] [name ...]

ilbadm delete-rule -a | name ...

ilbadm enable-rule [name ...]

ilbadm disable-rule [name ...]

ilbadm show-statistics [-p] -o field[,field] [-thAdvi]
    [-r rulename] | [-s servername] [interval [count]]

ilbadm create-servergroup [-s server=hostspect[:portspec...]] groupname

ilbadm delete-servergroup groupname

ilbadm show-servergroup [-s|-f|[-p] -o field[,field]] [[-v] name]

ilbadm enable-server server ...

ilbadm disable-server server ...

ilbadm show-server [[-p] -o field[,field...]] [rulename...]

ilbadm add-server -s server=value[,value ... ] name

ilbadm remove-server -s server=value[,value ... ] name

ilbadm create-healthcheck [-n] -h hc-test=value
    [,hc-timeout=value][,hc-count=value][,hc-interval=value] hcname

ilbadm delete-healthcheck hcname

ilbadm show-healthcheck [hcname ...]

ilbadm show-hc-result [rule-name]

ilbadm show-nat [count]

ilbadm show-persist [count]

ilbadm export-config filename

ilbadm import-config [-p] filename

```

描述

ilbadm 命令使用下面介绍的子命令处理或显示有关集成负载均衡器 (Integrated Load Balancer, ILB) 规则的信息。

规则名称不区分大小写，但会保留输入时的大小写。规则名称的长度限制为 19 个字符。服务器名称不能超过 14 个字符。

所有可解析的输出（使用 `-p` 选项调用）都要求使用 `-o` 选项指定要输出或显示的字段。字段将按它们在命令行中的相同显示顺序显示。使用冒号 (:) 字符分隔多个字段。如果显示的字符串本身包含冒号或反斜杠 (\)，则会在字符串前面添加一个反斜杠。可解析的输出不会显示任何标头。

当使用 `create-servergroup` 或 `add-server` 子命令添加服务器时，系统将生成服务器 ID。

确保服务器 ID 在服务器组中唯一。一个规则只能附加到一个服务器组，因此服务器 ID 对规则也是唯一的。请注意，由于同一服务器组可以附加多个规则，因此单凭服务器 ID 不足以表示规则。

为了区分服务器 ID 和主机名，服务器 ID 带有前导下划线 (_) 前缀。

如下所述，必须先定义服务器组和运行状况检查实体，然后才能在 `create-rule` 子命令中使用它们。

子命令

下面列出了 `ilbadm` 子命令及其相关选项和操作数。请注意，子命令具有标准和简捷两种形式；例如，`create-rule` 和 `create-rl`，这样您不必键入其他几个字符。

```
create-rule|create-rl [-e] [-p] -i incoming -m method_attributes -o outgoing_spec [-h
healthcheck] [-t timers] name
```

创建具有一组指定特征的规则 *name*。*incoming* 和 *method_attributes* 均指定为一组键=值对。如果 *name* 已存在，该命令将失败。如果给定元组（虚拟 IP 地址、端口和协议）与另一规则匹配，该命令也将失败。`create-rule` 具有下列选项，用于控制此命令的整体效果：

- `-e` 启用 `create-rule` 功能。缺省值为禁用 `create-rule`。
- `-p` 创建持久性（粘性）规则。缺省值为规则仅存在于当前会话中。

键和值通过包含一个字母的标识符引入。这些标识符及其相关键和可接受的值如下所示。

- `-i` 引入传入包的匹配条件。
 - `vip`
（虚拟）目标 IP 地址
 - `port[-port]`
端口号或名称，例如，`telnet` 或 `dns`。端口可以使用端口号或符号名称（如在 `/etc/services` 中）指定。此外，还支持端口号范围。
 - `protocol`
TCP（缺省值）或 UDP（请参见 `/etc/services`）。
- `-m` 指定说明包的处理方式的键。

lbalg

缺省值为 **roundrobin** 或其简捷形式 **rr**。其他备用项包括：**hash-ip**（简捷形式：**hip**）、**hash-ip-port**（简捷形式：**hipp**）、**hash-ip-vip**（简捷形式：**hipv**）。

type

表示网络拓扑。可以为 **DSR**（或 **dss** 或 **d**）、**NAT**（或 **n** 或 **nat**）、**HALF-NAT**（或 **h** 或 **half-nat**）。

proxy-src

只有完全 **NAT** 才需要此选项。指定用作代理源地址范围的 IP 地址范围。此范围限定为 10 个 IP 地址。

pmask

可选。别名为：**stickiness**。指定此规则将用作持久性规则。此参数是采用 **CIDR** 表示法的前缀长度；对于 **IPv4**，为 0-32；对于 **IPv6** 为 0-128。使用 **-p** 选项指定此关键字。

-o

指定与 **-i**“子句”指定的条件相符的包的目标。此标识符具有一个已知参数：

servergroup 指定一个服务器组作为目标。必须已创建服务器组。

-h

运行状况检查选项具有两个参数：

hc-name

指定预定义的运行状况检查方法的名称

hc-port

指定要检查的 **HC** 测试程序的端口。该值可以为关键字 **ALL** 或 **ANY**，也可以为服务器组的端口范围中的特定端口号。

-t

指定自定义计时器（以秒为单位）。值 **0** 表示使用系统缺省值。下面列出了 **-t** 的有效修饰符：

conn-drain

如果服务器的 **type** 为 **NAT** 或 **HALF-TYPE**，**conn-drain** 即为超时，在此超时后，从规则中删除服务器之后，将删除此服务器的连接状态。即使服务器未处于空闲状态，也会执行此删除。

适用于 **TCP** 的缺省值为：连接状态保持稳定，直到以正常方式关闭此连接为止。适用于 **UDP** 的缺省值为：连接状态保持稳定，直到此连接在时段 **nat-timeout** 内保持空闲为止。

nat-timeout

仅适用于 **NAT** 和半 **NAT** 类型连接。如果此类连接在 **nat-timeout** 时段内处于空闲状态，则将删除连接状态。适用于 **TCP** 的缺省值为 **120**，适用于 **UDP** 的缺省值为 **60**。

persist-timeout

启用持久性映射后，如果在 **persist-timeout** 秒内未使用仅限数字的映射，则将删除此映射。缺省为 60。

请注意，必须先定义服务器组和运行状况检查，然后才能在 **create-rule** 中使用它们。

delete-rule|delete-rl -a name[...]

删除与规则 *name* 相关的所有信息。如果 *name* 不存在，该命令将失败。**delete-rule** 具有一个选项：

-a

删除所有规则。（忽略 *name*。）

enable-rule|enable-rl name[...]

启用指定规则，或者如果未指定名称，则启用所有规则。启用已启用的规则不起作用。

disable-rule|disable-rl name[...]

禁用指定规则，或者如果未指定名称，则禁用所有规则。禁用已禁用的规则不起作用。

show-statistics|show-stats [[-p] -o field[,...]] [-tv] [-A | -d] [[-i] -r rulename] -s servername] [interval [count]]

显示统计信息，统计信息输出易受下述选项的影响。此子命令的语法和语义以 **vmstat(1M)** 的语法和语义作为模型。

-t

在每个样例前面附加一个时间戳。

-d

显示在整个时间间隔内的增量。缺省值为每秒钟的更改。无法与 **-a** 选项配合使用。

-A

显示绝对数字。即，自初始化模块、创建规则和添加服务器以来的数字。无法与 **-d** 选项配合使用。

-r rulename

仅显示指定 *rulename* 的统计信息。与 **-i** 选项结合使用，为每台服务器各显示一行。

-s servername

仅显示 *server* 的统计信息。与 **-i** 选项结合使用，为每个规则各显示一行。

-i

逐条列出 **-r** 和 **-s** 选项显示的信息。这些选项是可与 **-i** 有效配合使用的唯一选项。无法与 **-v** 选项配合使用。

-v

显示删除操作的其他详细信息。请注意，当指定规则名称时，将根据规则（而非服务器）计算删除的计数。无法与 -i 选项配合使用。

-p

显示可解析的格式。需要使用 -o 选项。

-o *field*

可以为下表中的一个或多个项。*field* 可以为大写或小写。

PKT_P	已处理的包。
BYTES_P	已处理的字节。
PKT_U	未处理的包。
BYTES_U	未处理的字节。
PKT_D	已删除的包。
BYTES_D	已删除的字节。
ICMP_P	已处理的 ICMP 回显请求。
ICMP_D	已删除的 ICMP 回显请求。
ICMP2BIG_P	需要的 ICMP 分段；已处理的消息。
ICMP2BIG_D	需要的分段；已删除的消息。
NOMEMP_D	因内存不足条件而删除的包。
NOPORTP_D	因源端口不可用而在 NAT 模式下删除的包。

请注意，当问号 (?) 显示为列条目时，这表示无法确定正确值，其原因多半是因为已添加或删除规则或服务器。

请注意，仅对每 10 个样例显示一次标头。时间戳格式遵循 C 语言环境的 `date(1)` 格式。未检测到规则的添加和删除操作。

`show-rule[show-rl [-d|-e] [-f|[-p] -o field[,...]] [name...]`

显示指定规则的特征，或者如果未指定规则，则显示所有规则的特征。此子命令具有以下选项：

-d

仅显示已禁用的规则。

-e

仅显示已启用的规则。

-f

显示完整列表。

-o *field*[,...]

显示 *field* 的输出。无法与 **-f** 选项配合使用。

-p

以“说明”中介绍的格式显示可解析的输出。需要使用 **-o** 选项。

请注意，**-o**（带或不带 **-p**）和 **-f** 选项是互斥的。

show-nat *count*

显示 NAT 表信息。如果指定了 *count*，则显示 NAT 表中的 *count* 个条目。如果未指定 *count*，则显示整个 NAT 表。

count

在连续运行此命令的过程中，不对元素的相对位置做出任何假定。例如，执行两次 **show-nat 10** 不能保证两次显示的 10 个项相同，特别是在繁忙系统中。

显示格式：

```
T: IP1 > IP2 >>> IP3 > IP4
```

这些项如下所述：

T 在此条目中使用的传输协议。

IP1 客户机的 IP 地址和端口。

IP2 VIP 和端口。

IP3 如果为半 NAT 模式，则为客户机的 IP 地址和端口。如果为完全 NAT 模式，则为经过 NAT 的客户机的 IP 地址和端口。

IP4 后端服务器的 IP 地址和端口。

show-persist|show-pt *count*

显示持久性表信息。如果指定了 *count*，则显示此表中的 *count* 个条目。如果未指定 *count*，则显示整个持久性表。

在连续运行此命令的过程中，不对元素的相对位置做出任何假定。例如，执行两次 **show-persist 10** 不能保证两次显示的 10 个项相同，特别是在繁忙系统中。

显示格式：

```
R: IP1 --> IP2
```

这些项如下所述：

R

此持久性条目绑定到的规则。

IP1

客户机的 IP 地址和端口。

IP2

后端服务器的 IP 地址。

`export-config|export-cf [filename]`

以适用于使用 `ilbadm import` 重新导入的格式导出当前配置。如果未指定 `filename`，此子命令写入到标准输出。

`import-config|import-cf [-p] [filename]`

读取文件的配置内容。缺省情况下，这会覆盖所有现有配置。如果未指定 `filename`，此子命令从标准输入读取。此子命令具有以下选项：

`-p`

保留现有配置并执行增量导入。

`create-servergroup|create-sg [-s server=hostspec[:portspec...]] groupname`

创建服务器组。以后可以使用 `add-server` 子命令添加其他服务器。服务器组是在规则创建期间使用的用于表示后端服务器的唯一实体。如果指定服务器组与一个或多个规则相关联，则会在添加此服务器时启用它。此子命令具有以下选项和操作数：

`-s server=hostspec[:portspec...]`

指定要添加到服务器组的服务器的列表。

hostspec 为主机名或 IP 地址。IPv6 地址必须括在方括号 ([]) 中，以便与 *:portspec* 区分开来。

portspec 为服务名称或端口号。如果未指定端口号，则使用范围 1–65535 中的数字。

`disable-server|disable-srv server`

禁用一台或多台服务器。也就是说，通知内核不要将通信转发到此服务器。`disable-server` 适用于包含此服务器的服务器组所附加的所有规则。

server 为服务器 ID。

`enable-server|enable-srv server...`

重新启用已禁用的服务器。

`show-server|show-srv [[-p] -o field[,field...]] [rulename...]`

显示与指定规则关联的服务器，或者如果未指定规则名称，则显示所有服务器。此子命令具有以下选项。

`-o field[,field...]`

仅显示指定字段。

`-p`

以可解析的格式显示字段。需要使用 `-o` 选项。

`delete-servergroup|delete-sg groupname`

删除服务器组。

`show-servergroup|show-sg [[-p] -o field[,...]] [name]`

列出某个服务器组，或者如果未指定 *name*，则列出所有服务器组。此子命令具有以下选项：

`-o field[,...]`

显示 *field* 的输出。

`-p`

以“说明”中介绍的格式显示可解析的输出。需要使用 `-o` 选项。

`add-server|add-srv -s server=value[, value...] servergroup`

向 *servergroup* 添加指定服务器。有关 *value* 的定义，请参见 `create-servergroup` 的说明。

`-s`

请参见 `create-servergroup`。

在对服务器组执行 `remove-server` 之后立即对该服务器组执行 `add-server` 可能会失败，这是因为未完成连接清除。有关如何避免此故障的说明，请参见 `remove-server` 子命令的说明。

`remove-server|remove-srv -s server=value[, value...] servergroup`

从 *servergroup* 删除指定服务器。

`-s`

一个或多个服务器 ID。

如果 NAT/半 NAT 规则正在使用服务器，则建议在删除前禁用此服务器（使用 `disable-server`）。通过禁用服务器，服务器将进入连接清除状态。在清除所有连接后，可以使用 `remove-server` 删除服务器。如果设置了 `conn-drain` 超时值，在超时时段结束时，连接清除状态将结束。请注意，缺省 `conn-drain` 超时为 0，这表示将在正常关闭连接之前一直等待。

`create-healthcheck|create-hc [-n] -h hc-test=value,hc-timeout=value,`

`hc-count=num_value,hc-interval=value hcname`

为要使用的规则设置运行状况检查对象。使用相同测试检查与规则关联的所有服务器。服务器的运行状况检查事件包括执行 1 至 `hc-count` 次 `hc-test`。如果 `hc-test` 的结果显示服务器无响应，则会进一步执行 `hc-test` 检查，并最多进行 `hc-count` 次调用，然后才会将服务器视为已关闭。

`-h`

执行 `hc-count` 次 `hc-test`，直到其成功或 `hc-timeout` 已过期为止。对于给定规则，使用相同测试检查所有服务器。测试如下所示：

`hc-test`

PING、TCP、外部方法（脚本或二进制）。应使用全路径名指定外部方法。

`hc-timeout`

在 `hc-test` 临时失败后将测试视为失败时的阈值。如果中止 `hc-test` 测试，则会将结果视为失败。缺省值为 5 秒。

hc-count

在将服务器标记为关闭之前尝试运行 **hc-test** 的最大次数。缺省值为重复 3 次。

hc-interval

两次调用 **hc-test** 之间的时间间隔。此值必须大于 **hc-timeout** 乘以 **hc-count** 的值。缺省值为 30 秒。

将下列参数传递给外部方法：

\$1

VIP（IPv4 或 IPv6 文本地址）。

\$2

服务器 IP（IPv4 或 IPv6 文本地址）。

\$3

协议（作为字符串的 UDP、TCP）。

\$4

负载均衡模式 DSR、NAT、HALF_NAT。

\$5

数字端口。

\$6

方法在返回故障之前应等待的最长时间（以秒为单位）。如果此方法的运行时间长于此时间，则会中止此方法，并将测试视为失败。

如果成功，外部方法应返回 0 [或者向后端服务器返回往返时间（以微秒为单位）]，如果服务器视为已关闭，则应返回 -1。

在开始较高层的运行状况检查、TCP、UDP 和外部测试之前，必须首先执行缺省的 ping 测试。如果 ping 失败，则不会执行较高层的测试。通过使用 **-n**，您可以针对上述较高层的运行状况检查禁用缺省的 ping 检查。

-n

针对较高层的运行状况检查测试禁用缺省的 ping 测试。

delete-healthcheck|delete-hc *hcname*...

删除指定的运行状况检查对象 (*hcname*)。如果给定运行状况检查对象与启用的规则相关联，删除此对象将失败。

show-healthcheck|show-hc [*hcname*...]

列出指定的运行状况检查 (*hcname*) 的运行状况检查信息。如果未指定运行状况检查，则列出所有现有运行状况检查的信息。

show-hc-result|show-hc-res [*rule-name*]

列出与 *rule-name* 关联的服务器的运行状况检查结果。如果未指定 *rule-name*，则会显示所有服务器的运行状况检查结果。

示例

示例1 配置 NAT 模式

下面的命令创建设置了运行状况检查和计时器的规则（端口范围移动和会话持久性）。

```
# ilbadm create-healthcheck -h hc-test=tcp,hc-timeout=2,hc-count=3, \
    hc-interval=10 hc1
# ilbadm create-servergroup -s \
    server=60.0.0.10:6000-6009,60.0.0.11:7000-7009 sg1
# ilbadm create-rule -e -i vip=81.0.0.10,port=5000-5009,protocol=tcp \
-m lbalg=rr,type=NAT,proxy-src=60.0.0.101-60.0.0.104, \
    pmask=24 \
-h hc-name=hc1 \
-t conn-drain=180,nat-timeout=180,persist-timeout=180 \
-o servergroup=sg1 rule1
```

下面的命令创建带有缺省计时器值而无运行状况检查的规则。

```
# ilbadm create-servergroup -s server=60.0.0.10 sg1
# ilbadm create-rule -e -i vip=81.0.0.10,port=5000 \
-m lbalg=rr,type=NAT,proxy-src=60.0.0.105 \
-o servergroup=sg1 rule1
# ilbadm add-server -e -s server=60.0.0.11sg1
# ilbadm enable-rule rule1
```

示例2 配置半 NAT 模式

下面的命令配置半 NAT 模式，并使用端口范围折叠作为示例。

```
# ilbadm create-servergroup sg1
# ilbadm create-rule -e -i vip=81.0.0.10,port=5000-5009 \
-m lbalg=rr,type=h -o servergroup=sg1 rule1
# ilbadm add-server -s server=60.0.0.10:6000,60.0.0.11:7000 sg1
```

示例3 配置 DSR 模式并准备两组规则

下面的命令建立两组规则，以便在 HTTP 和 FTP 通信之间实现负载平衡。请注意两种类型的通信遍历接口 60.0.0.10。

```
# ilbadm create-servergroup -s servers=60.0.0.9,60.0.0.10 websg
# ilbadm create-servergroup -s servers=60.0.0.10,60.0.0.11 ftpgroup

# ilbadm create-rule -e -i vip=81.0.0.10,port=80 \
-m lbalg=hash-ip-port,type=DSR \
-o servergroup=websg webrule
# ilbadm create-rule -e -i vip=81.0.0.10,port=ftp \
-m lbalg=hash-ip-port,type=DSR,pmask=24 \
-o servergroup=ftpgroup ftprule
# ilbadm create-rule -e -p -i vip=81.0.0.10,port=ftp-data \
-m lbalg=hash-ip-port,type=DSR,pmask=24 \
-o servergroup=ftpgroup ftpdatarule
```

示例4 删除规则、服务器组和运行状况检查

下面的命令删除在第一个示例中建立的规则、服务器组和运行状况检查。

```
# ilbadm ilbadm delete-rule -a
# ilbadm delete-servergroup sg1
# ilbadm delete-healthcheck hc1
```

示例5 显示规则列表。

下面的命令显示规则列表。

```
# ilbadm show-rule
RULENAME          STATUS LBALG      TYPE    PROTOCOL VIP  PORT
r2                 E      hash-ip      NAT     TCP 45.0.0.10 81
r1                 E      hash-ip      NAT     TCP 45.0.0.10 80

# ilbadm show-rule -f
      RULENAME: rule1
      STATUS: E
      PORT: 80
      PROTOCOL: TCP
      LBALG: roundrobin
      TYPE: HALF-NAT
      PROXY-SRC: --
      PERSIST: --
      HC-NAME: hc1
      HC-PORT: ANY
      CONN-DRAIN: 0
      NAT-TIMEOUT: 120
      PERSIST-TIMEOUT: 60
      SERVERGROUP: sg1
      VIP: 80.0.0.2
      SERVERS: _sg1.0,_sg1.1
```

示例6 导出和导入规则

下面的命令显示如何在标准输出和文件中导出和导入规则。

```
# ilbadm export-config

create-servergroup ftpgroup
add-server -s server=10.1.1.3:21 ftpgroup
add-server -s server=10.1.1.4:21 ftpgroup
create-servergroup webgroup_v6
add-server -s server=[2000::ff]:80 webgroup_v6
create-rule -e protocol=tcp,VIP=1.2.3.4,port=ftp \
-m lbalg=roundrobin,type=DSR \
-o servergroup=ftpgroup rule4
create-rule protocol=tcp,VIP=2003::1,port=ftp \
-m lbalg=roundrobin,type=DSR \
```

示例6 导出和导入规则 (续)

```
-o servergroup=ftpgroup6 rule3
create-rule -e protocol=tcp,VIP=2002::1,port=http \
-m lbalg=roundrobin,type=DSR \
-o servergroup=webgrp_v6 RULE-all
```

下面的命令将规则导出到文件。

```
# ilbadm export-config /tmp/ilbrules
```

执行此命令之后，/tmp/ilbrules 包含在上一命令中显示的输出。

下面的命令从文件导入规则。

```
# ilbadm import-config /tmp/ilbrules
```

此命令使用 /tmp/ilbrules 的内容替换了原有的任何规则。

下面的命令从**标准输入**导入规则。

```
# cat /tmp/ilbrules | ilbadm import-config
```

此命令的作用与上一命令的作用相同。

示例7 创建一个运行状况检查

下面的命令创建一个运行状况检查。

```
# ilbadm create-healthcheck -h hc-timeout=3,hc-count=2,hc-interval=8,\
hc-test=tcp hc1
```

示例8 列出所有运行状况检查

下面的命令列出所有现有的运行状况检查。

```
# ilbadm show-healthcheck
```

HCNAME	TIMEOUT	COUNT	INTERVAL	DEF_PING	TEST
hc1	2	1	10	Y	tcp
hc2	2	1	10	N	/usr/local/bin/probe

示例9 删除一个运行状况检查

下面的命令删除一个运行状况检查。

```
# ilbadm delete-healthcheck hc1
```

示例10 显示统计信息

下面的命令以 1 秒钟为时间间隔重复显示三次统计信息。

```
# ilbadm show-stats -A 1 3
```

PKT_P	BYTES_P	PKT_U	BYTES_U	PKT_D	BYTES_D
0	0	0	0	4	196

示例 10 显示统计信息 (续)

```
0      0      0      0      4      196
0      0      0      0      4      196
```

下面的命令用于以 1 秒钟为时间间隔在详细模式下显示统计信息。输出太宽，无法包含在页界限中。

```
# ilbadm show-stats -v 1
```

下面的命令以 1 秒钟为时间间隔重复显示 3 次规则 r1 的统计信息。

```
# ilbadm show-stats -A -r r1 1 3
PKT_P  BYTES_P  PKT_U  BYTES_U  PKT_D  BYTES_D
0      0      0      0      4      196
0      0      0      0      4      196
0      0      0      0      4      196
```

下面的命令以 1 秒钟为时间间隔为每台服务器显示 3 次规则 r1 的统计信息。

```
# ilbadm show-stats -A -r r1 -i 1 3
SERVERNAME      PKT_P  BYTES_P
_sg1.0          0      0
_sg1.1          0      0
_sg1.2          0      0
_sg1.0          0      0
_sg1.1          0      0
_sg1.2          0      0
_sg1.0          0      0
_sg1.1          0      0
_sg1.2          0      0
```

下面的命令以 1 秒钟为时间间隔显示 3 次服务器 _sg1.0 逐条列出的统计信息和时间戳。

```
# ilbadm show-stats -A -s _sg1.0 -it 1 3
RULENAME      PKT_P  BYTES_P  TIME
r1            0      0      2009-07-20:16.10.20
r1            0      0      2009-07-20:16.10.21
r1            0      0      2009-07-20:16.10.22
```

下面的命令以 1 秒钟为时间间隔显示 3 次统计信息和特定选项字段。

```
# ilbadm show-stats -o BYTES_D,TIME 1 3
BYTES_D  TIME
196      2009-07-20:16.14.25
0        2009-07-20:16.14.26
0        2009-07-20:16.14.27
```

示例 11 显示运行状况检查结果

下面的命令显示运行状况检查的结果。

示例 11 显示运行状况检查结果 (续)

```
# ilbadm show-hc-result rule1
RULENAME  HCNAME      SERVERID  STATUS  FAIL  LAST      NEXT      RTT
rule1     hc1         _sg1.0   dead    6     04:45:17 04:45:30 698
rule1     hc1         _sg1.1   alive   0     04:45:11 04:45:25 260
rule1     hc1         _sg1.2   unreach 6     04:45:17 04:45:30 0
```

示例 12 显示 NAT 表
下面的命令显示 NAT 表。

```
# ilbadm show-nat 5
UDP: 124.106.235.150.53688>85.0.0.1.1024>>>82.0.0.39.4127>82.0.0.56.1024
UDP: 71.159.95.31.61528> 85.0.0.1.1024>>> 82.0.0.39.4146> 82.0.0.55.1024
UDP: 9.213.106.54.19787> 85.0.0.1.1024>>> 82.0.0.40.4114> 82.0.0.55.1024
UDP: 118.148.25.17.26676> 85.0.0.1.1024>>>82.0.0.40.4112> 82.0.0.56.1024
UDP: 69.219.132.153.56132>85.0.0.1.1024>>>82.0.0.39.4134> 82.0.0.55.1024
```

在实际 ilbadm 输出中，为了实现更好的可读性，使用空格进行分隔。

示例 13 显示持久性表
下面的命令显示持久性表。

```
# ilbadm show-persist 5
rule2: 124.106.235.150 --> 82.0.0.56
rule3: 71.159.95.31 --> 82.0.0.55
rule3: 9.213.106.54 --> 82.0.0.55
rule1: 118.148.25.17 --> 82.0.0.56
rule2: 69.219.132.153 --> 82.0.0.55
```

示例 14 显示服务器组
下面的命令显示有关服务器组的基本信息。

```
# ilbadm show-servergroup
sg1: id:sg1.2 35.0.0.4:80
sg1: id:sg1.1 35.0.0.3:80
sg1: id:sg1.0 35.0.0.2:80
sg2: id:sg2.3 35.0.0.5:81
sg2: id:sg2.2 35.0.0.4:81
sg2: id:sg2.1 35.0.0.3:81
sg2: id:sg2.0 35.0.0.2:81
```

下面的命令显示有关服务器组的所有可用信息。

```
# ilbadm show-servergroup -o all
sgname      serverID      minport maxport IP_address
sg1         _sg1.0       --      --      1.1.1.1
sg2         _sg2.1       --      --      1.1.1.6
```

示例 14 显示服务器组 (续)

sg3	_sg3.0	9001	9001	1.1.1.1
sg3	_sg3.1	9001	9001	1.1.1.2
sg3	_sg3.2	9001	9001	1.1.1.3
sg3	_sg3.3	9001	9001	1.1.1.4
sg3	_sg3.4	9001	9001	1.1.1.5
sg3	_sg3.5	9001	9001	1.1.1.6
sg3	_sg3.6	9001	9001	1.1.1.11
sg3	_sg3.7	9001	9001	1.1.1.12
sg3	_sg3.8	9001	9001	1.1.1.13
sg3	_sg3.9	9001	9001	1.1.1.14
sg3	_sg3.10	9001	9001	1.1.1.15
sg3	_sg3.11	9001	9001	1.1.1.16
sg4	_sg4.0	9001	9006	1.1.1.1
sg4	_sg4.1	9001	9006	1.1.1.6

示例 15 列出与规则关联的服务器
下面的命令列出规则关联的服务器。

```
# ilbadm show-server r1
SERVERID      ADDRESS      PORT RULENAME  STATUS SERVERGROUP
_sg1.0        35.0.0.10    80  rule1      E      sg1
_sg1.1        35.0.0.11    80  rule1      E      sg1
_sg1.2        35.0.0.12    80  rule1      D      sg1
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/network/load-balancer/ilb
接口稳定性	Committed（已确定）

另请参见 [inetadm\(1M\)](#)、[vmstat\(1M\)](#)、[attributes\(5\)](#)

引用名	ilbd – ILB 守护进程
用法概要	<code>/usr/lib/inet/ilbd</code>
描述	集成负载均衡器 (Integrated Load Balancer, ILB) 守护进程 <code>ilbd</code> 为 ILB 提供用户级服务。它是由服务管理工具启动的一个系统守护进程（请参见 smf(5)）。其故障管理资源标识符 (fault management resource identifier, FMRI) 为： <pre>svc:/network/loadbalancer/ilb:default</pre> 请注意，<code>ilbd</code> 是一个 Consolidation Private（整合专用）接口。请参见 attributes(5)。 <code>ilbd</code> 不提供管理接口。所有 ILB 管理都应通过 ilbadm(1M) 或程序库 libilb(3LIB) 进行。
选项	<code>ilbd</code> 守护进程没有选项。
示例	示例1 启用 ILB 服务 以下命令启用 ILB 服务： <pre># svcadm enable svc:/network/loadbalancer/ilb:default</pre> 示例2 禁用 ILB 服务 以下命令禁用 ILB 服务： <pre># svcadm disable svc:/network/loadbalancer/ilb:default</pre>
错误	<code>ilbd</code> 守护进程使用 syslog(3C) 来报告状态和错误消息。所有的消息由 <code>LOG_DAEMON</code> 工具记录。错误消息是以 <code>LOG_ERR</code> 和 <code>LOG_NOTICE</code> 优先级记录的，信息性消息是以 <code>LOG_INFO</code> 优先级记录的。 <code>/etc/syslog.conf</code> 文件中的缺省项是将所有 <code>ilbd</code> 守护进程错误消息记录到 <code>/var/adm/messages</code> 日志中。
文件	<code>/usr/lib/inet/ilbd</code> ILB 守护进程二进制文件
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	<code>service/network/load-balancer/ilb</code>
接口稳定性	Project Private（项目专用）

另请参见 [svcs\(1\)](#)、[ilbadm\(1M\)](#)、[svcadm\(1M\)](#)、[syslog\(3C\)](#)、[libilb\(3LIB\)](#)、[syslog.conf\(4\)](#)、[attributes](#)

附注 `ilbd` 服务是由服务管理工具 [smf\(5\)](#) 管理的，其故障管理资源标识符 (fault management resource identifier, FMRI) 为：

```
svc:/network/loadbalancer/ilb:default
```

可以使用 `svcadm(1M)` 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 `svcs(1)` 命令来查询服务的状态。要管理该服务，必须授予 `solaris.smf.manage.ilb` 授权，或者授予权限配置文件 "ILB Management"（ILB 管理）。

必须启用 ILB 服务器，`ilbadm(1M)` 命令和任何其他 `libilb(3LIB)` 客户机应用程序才能正常运行。

引用名	ilomconfig – ILOM LAN configuration utility
用法概要	<pre>ilomconfig list subcommand [options] ilomconfig modify subcommand [options] ilomconfig enable subcommand [options] ilomconfig disable subcommand [options]</pre>
描述	ilomconfig is an Oracle Integrated Lights Out Management (ILOM) configuration utility that allows an administrator to configure the LAN interface between the Service Processor (SP) and host. The administrator can view the LAN configuration settings by means of the <code>list</code> command and configure the interface using the <code>enable</code> , <code>disable</code> , and <code>modify</code> commands.
子命令	The following subcommands are available under the <code>list</code> command: <pre>system-summary</pre> List product summary information and ILOM version. <pre>interconnect</pre> List Host-to-ILOM interconnect settings. The following subcommands are available under the <code>modify</code> command: <pre>interconnect</pre> Modifies Host-to-ILOM interconnect settings. The following subcommand is available under the <code>enable</code> command: <pre>interconnect</pre> Enables the Host-to-ILOM interconnect. The following subcommand is available under the <code>disable</code> command: <pre>interconnect</pre> Disables the Host-to-ILOM interconnect.
选项	The following options are available for the <code>enable</code> and <code>modify</code> commands: <pre>-ipaddress=ipaddress</pre> Set the IP address for the SP side of the LAN interconnect. <pre>-hostipaddress=ipaddress</pre> Set the IP address for the host side of the LAN interconnect. <pre>-netmask=netmask</pre> Set the network mask for the IP addresses used for the LAN. The following are general options: <pre>-, -h, --help</pre> Display a brief usage message and then exit.

- q, --quiet
Suppress informational message output and return only error codes.
- V, --version
Display version information for ilomconfig and then exit.

文件
/var/log/ilomconfig/ilomconfig.log
ilomconfig-specific log file.

属性
See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/management/ilomconfig
Interface Stability	Committed

另请参见 [attributes\(5\)](#)

引用名	imqadmin – launch the Message Queue administration console
用法概要	<pre> /usr/bin/imqadmin [-javahome path] /usr/bin/imqadmin -h /usr/bin/imqadmin -v </pre>
描述	imqadmin launches the graphical user interface application that performs most Message Queue administration tasks. These tasks include managing broker instances (including physical destinations) and administered objects.
选项	The following options are supported: -h Display usage help. The application is not launched. -javahome <i>path</i> Specify a path to an alternate Java 2 compatible runtime. -v Display version information.
环境变量	The following environment variables affect the execution of this command: IMQ_JAVAHOME Specify the Java 2 compatible runtime. When this environment variable is not set it defaults to <code>/usr/j2se</code>.
退出状态	The following exit values are returned: 0 Successful completion. >0 An error occurred.
文件	<code>\$HOME/.imq/admin/brokerlist.properties</code> Contains user settings, a list of broker instances being managed. <code>\$HOME/.imq/admin/objectstorelist.properties</code> Contains user settings, a list of object stores being managed.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWiq

另请参见 [imqbrokerd\(1M\)](#), [imqcmd\(1M\)](#), [imqdbmgr\(1M\)](#), [imqkeytool\(1M\)](#), [imqobjmgr\(1M\)](#), [imqusermgr\(1M\)](#), [attributes\(5\)](#)

Sun Java System Message Queue Administrator's Guide

引用名	imqbrokerd – start a Message Queue broker instance												
用法概要	<pre>/usr/bin/imqbrokerd [option...] /usr/bin/imqbrokerd -h</pre>												
描述	imqbrokerd starts an instance of the Message Queue broker. The Message Queue broker is the main component of a Message Queue message server. The broker performs reliable delivery of messages to and from Java Message Service (JMS) clients. imqbrokerd uses command line options to specify broker configuration properties.												
选项	The following options are supported: <table><tr><td>-backup <i>fileName</i></td><td>Back up a Master Broker's configuration change record to <i>fileName</i>. This option only applies to broker clusters.</td></tr><tr><td>-cluster <i>brokerList</i></td><td> Specify the list of broker instances which are connected in a cluster. This list is merged with the list in the imq.cluster.brokerlist property. This option only applies to broker clusters. <i>brokerList</i> is a comma-separated list of broker instances, each specified by <i>hostName:port</i> (the host on which the broker instance is running and the port number it is using) If you don't specify a value for <i>hostName</i>, <i>localhost</i> is used. If you don't specify a value for <i>port</i>, the value of 7676 is used. For example: <i>host1:8899,host2,:7878</i>. </td></tr><tr><td>-dbpassword <i>password</i></td><td>Specify the password for a plugged-in JDBC-compliant database used as a Message Queue data store.</td></tr><tr><td>-dbuser <i>userName</i></td><td>Specify the user name for a plugged-in JDBC-compliant data store.</td></tr><tr><td>-D<i>property</i>=<i>value</i></td><td>Set the specified broker configuration property to the <i>value</i>. The system does not validate either the configuration property or <i>value</i>. Therefore, spelling, formatting, and case is important. Message Queue can not set incorrect values passed using the -D option.</td></tr><tr><td>-force</td><td>Perform action without user confirmation. This option only applies when you use the -remove <i>instance</i> option, which normally requires confirmation.</td></tr></table>	-backup <i>fileName</i>	Back up a Master Broker's configuration change record to <i>fileName</i> . This option only applies to broker clusters.	-cluster <i>brokerList</i>	Specify the list of broker instances which are connected in a cluster. This list is merged with the list in the imq.cluster.brokerlist property. This option only applies to broker clusters. <i>brokerList</i> is a comma-separated list of broker instances, each specified by <i>hostName:port</i> (the host on which the broker instance is running and the port number it is using) If you don't specify a value for <i>hostName</i>, <i>localhost</i> is used. If you don't specify a value for <i>port</i>, the value of 7676 is used. For example: <i>host1:8899,host2,:7878</i>.	-dbpassword <i>password</i>	Specify the password for a plugged-in JDBC-compliant database used as a Message Queue data store.	-dbuser <i>userName</i>	Specify the user name for a plugged-in JDBC-compliant data store.	-D <i>property</i> = <i>value</i>	Set the specified broker configuration property to the <i>value</i> . The system does not validate either the configuration property or <i>value</i> . Therefore, spelling, formatting, and case is important. Message Queue can not set incorrect values passed using the -D option.	-force	Perform action without user confirmation. This option only applies when you use the -remove <i>instance</i> option, which normally requires confirmation.
-backup <i>fileName</i>	Back up a Master Broker's configuration change record to <i>fileName</i> . This option only applies to broker clusters.												
-cluster <i>brokerList</i>	Specify the list of broker instances which are connected in a cluster. This list is merged with the list in the imq.cluster.brokerlist property. This option only applies to broker clusters. <i>brokerList</i> is a comma-separated list of broker instances, each specified by <i>hostName:port</i> (the host on which the broker instance is running and the port number it is using) If you don't specify a value for <i>hostName</i>, <i>localhost</i> is used. If you don't specify a value for <i>port</i>, the value of 7676 is used. For example: <i>host1:8899,host2,:7878</i>.												
-dbpassword <i>password</i>	Specify the password for a plugged-in JDBC-compliant database used as a Message Queue data store.												
-dbuser <i>userName</i>	Specify the user name for a plugged-in JDBC-compliant data store.												
-D <i>property</i> = <i>value</i>	Set the specified broker configuration property to the <i>value</i> . The system does not validate either the configuration property or <i>value</i> . Therefore, spelling, formatting, and case is important. Message Queue can not set incorrect values passed using the -D option.												
-force	Perform action without user confirmation. This option only applies when you use the -remove <i>instance</i> option, which normally requires confirmation.												

-h	Display usage help. Execute nothing else on the command line.
-javahome <i>path</i>	Specify the path to an alternate Java 2-compatible Java Development Kit (JDK) or Java Runtime Environment (JRE) The default is to use the runtime bundled with the operating system.
-ldappassword <i>password</i>	Specify the password for accessing an LDAP user repository when using an LDAP server (as opposed to a built-in flat-file repository) to authenticate users of a Message Queue message server.
-license [<i>name</i>]	Specify the license to load, if different from the default for your Message Queue product edition. If you don't specify a license name, this lists all licenses installed on the system. Depending on the installed Message Queue edition, the values for <i>name</i> are pe (Platform Edition-basic features), t ry (Platform Edition-90-day trial enterprise features), and unl (Enterprise Edition).
-loglevel <i>level</i>	Specify the logging level. Valid values for <i>level</i> are NONE, ERROR, WARNING, or INFO. The default value is INFO.
-metrics <i>int</i>	Report metrics at a specific interval. Specify <i>int</i> as the number of seconds.
-name <i>brokerName</i>	Specify the instance name of this broker and use the corresponding instance configuration file. If you do not specify a broker name, the name of the file is set to imqbroker. If you run more than one instance of a broker on the same host, each must have a unique name.
-passfile <i>filename</i>	Specify the name of the file from which to read the passwords for the SSL keystore, LDAP user repository, or JDBC-compliant database.
-password <i>keypassword</i>	Specify the password for the SSL certificate keystore.
-port <i>number</i>	Specify the broker's Port Mapper port number. By default, this is set to 7676. To run two instances of a broker on the same server, each broker's Port

	Mapper must have a different port number. JMS clients connect to the broker instance using this port number.								
- remove <i>instance</i>	Remove the broker instance. Delete the instance configuration file, log files, data store, and other files and directories associated with the broker instance. This option requires user confirmation unless you also specify the -force option.								
- reset store messages durables props	Reset the data store (or a subset of the store) or resets the configuration properties of the broker instance when the broker instance is started. The action depends on the argument provided. <table><tr><td>store</td><td>Clear all persistent data in the data store, including messages, durable subscriptions, and transaction information store.</td></tr><tr><td>messages</td><td>Clear all persistent messages durable.</td></tr><tr><td>durables</td><td>Clear all durable subscriptions.</td></tr><tr><td>props</td><td>Clear all configuration information in the config.props instance configuration file. All properties assume default values.</td></tr></table>	store	Clear all persistent data in the data store, including messages, durable subscriptions, and transaction information store.	messages	Clear all persistent messages durable.	durables	Clear all durable subscriptions.	props	Clear all configuration information in the config.props instance configuration file. All properties assume default values.
store	Clear all persistent data in the data store, including messages, durable subscriptions, and transaction information store.								
messages	Clear all persistent messages durable.								
durables	Clear all durable subscriptions.								
props	Clear all configuration information in the config.props instance configuration file. All properties assume default values.								
- restore <i>filename</i>	Replace the Master Broker's configuration change record with the specified backup file. This file must have been previously created using the -backup option. This option only applies to broker clusters.								
- shared	Specify that the jms connection service be implemented using the shared threadpool model, in which threads are shared among connections to increase the number of connections supported by a broker instance.								
- silent	Turn off logging to the console.								
- tty	Display all messages be to the console. WARNING and ERROR level messages are displayed on the console by default.								

<code>-upgrade-store-nobackup</code>	Specify that an earlier, incompatible version Message Queue data store is automatically removed when migrating to Message Queue 3.5 format. If you do not use this option, you must manually delete the earlier data store. This applies to both built-in (flat-file) persistence and plugged-in (JDBC-compliant) persistence. Migration of the earlier data store to a Message Queue 3.5 data store takes place the first time you start a Message Queue 3.5 broker instance on an earlier version data store.
<code>-version</code>	Display the version number of the installed product.
<code>-vmargs are [[arg]...]</code>	Specify arguments to pass to the Java VM. Separate arguments with spaces. If you want to pass more than one argument or if an argument contains a space, use enclosing quotation marks. For example: <code>imqbrokerd -tty -vmargs "-Xmx128m -Xincgc"</code>

环境变量

The following environment variables affect the execution of this command:

`IMQ_JAVAHOME` Specify the Java 2 compatible runtime. When this environment variable is not set it defaults to `/usr/j2se`.

退出状态

The following exit values are returned:

`0` Successful completion.

`>0` An error occurred.

文件

`/etc/init.d/imq`

Shell script for starting `imqbrokerd`. This file looks at the `/etc/imq/imqbrokerd.conf` file.

`/etc/imq/imqbrokerd.conf`

Configuration file which controls the behavior of the broker startup script.

`/etc/imq/passwd`

Flat file user repository for authenticating users.

`/etc/imq/accesscontrol.properties`

Controls client access to broker functionality.

`/etc/imq/passfile.sample`

Sample passfile used by the `-passfile` option.

`/var/imq/instances/brokerName/props/config.properties`
Broker instance configuration file.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWiq

另请参见 [imqadmin\(1M\)](#), [imqcmd\(1M\)](#), [imqdbmgr\(1M\)](#), [imqkeytool\(1M\)](#), [imqobjmgr\(1M\)](#), [imqusermgr\(1M\)](#), [attributes\(5\)](#)

Sun Java System Message Queue Administrator's Guide

引用名	imqcmd – manage Message Queue brokers
用法概要	<pre>/usr/bin/imqcmd subcommand argument [option...]</pre> <pre>/usr/bin/imqcmd [-h -H]</pre> <pre>/usr/bin/imqcmd -v</pre>
描述	imqcmd manages the Message Queue broker, including resources such as connection services, physical destinations, durable subscriptions, and transactions. The utility provides a number of subcommands for managing these resources. imqcmd supports many subcommands. Basic connection and authentication is required for the execution of every imqcmd subcommand. Use the <code>-secure</code> option to specify secure connections. Subcommands and their corresponding arguments and options follow the imqcmd command on the command line. See <code>USAGE</code> and <code>OPTIONS</code>.
选项	The following options are supported: <code>-b hostName:port</code> Specify the name of the host on which the broker instance is running and the port number it is using. The default value is <code>localhost:7676</code>. If you do not specify the <code>-b</code> option, imqcmd uses the default. To specify port only, use: <code>-b :7878</code>. This is equivalent to <code>-b localhost:7878</code> To specify name only, use: <code>-b somehost</code>. This is equivalent to <code>-b somehost:7676</code>. <code>-c clientID</code> Specify the ID of the durable subscriber to a topic. <code>-d topicName</code> Specify the name of the topic. Use this option with the <code>list dur</code> and <code>destroy dur</code> subcommands. <code>-f</code> Perform action without user confirmation. Use this option with any subcommand. <code>-h</code> Display usage help. Execute nothing else on the command line. <code>-H</code> Display usage help, attribute list, and examples. Execute nothing else on the command line. <code>-int interval</code> Specify the interval, in seconds, at which the <code>metrics bkr</code>, <code>metrics dst</code>, and <code>metrics svc</code> subcommands display metrics output. Use this option with the <code>metrics</code> subcommand. <code>-javahome</code> Specify an alternate Java 2 compatible runtime to use.

<code>-m <i>metricType</i></code>	Specify the type of metric information to diplay. Use this option with the <code>metrics bkr</code>, <code>metrics dst</code>, and <code>metrics svc</code> subcommands. The value of <i>metricType</i> depends on whether the metrics are generated for a destination, a service, or a broker. Use one of the following values to specify <i>metricType</i>: <table><tr><td><code>ttl</code></td><td>Total of messages in and out of the broker (default)</td></tr><tr><td><code>rts</code></td><td>Provides the same information as <code>ttl</code>, but specifies the number of messages per second</td></tr><tr><td><code>cxn</code></td><td>Connections, virtual memory heap, threads</td></tr></table> The following command displays connection, VM heap, and threads metric information for the default broker instance (<code>localhost:7676</code>) every five seconds: <pre>imqcmd metrics bkr -m cxn -int 5</pre>	<code>ttl</code>	Total of messages in and out of the broker (default)	<code>rts</code>	Provides the same information as <code>ttl</code> , but specifies the number of messages per second	<code>cxn</code>	Connections, virtual memory heap, threads
<code>ttl</code>	Total of messages in and out of the broker (default)						
<code>rts</code>	Provides the same information as <code>ttl</code> , but specifies the number of messages per second						
<code>cxn</code>	Connections, virtual memory heap, threads						
<code>-msp <i>numSamples</i></code>	Specify the number of samples the <code>metrics bkr</code> , <code>metrics dst</code> , and <code>metrics svc</code> subcommands display in the metrics output.						
<code>-n <i>argumentName</i></code>	Specify the name of the subcommand argument. Depending on the subcommand, this might be the name of a service, a physical destination, a durable subscription, or a transaction ID.						
<code>-o <i>attribute=value</i></code>	Specify the value of an attribute. Depending on the subcommand argument, this might be the attribute of a broker, service, or destination.						
<code>-p <i>password</i></code>	Specify the administrator password.						
	This option is deprecated. Use the <code>-passfile</code> option instead.						
<code>-passfile</code>	Specify the administrator password.						
<code>-pst <i>pauseType</i></code>	Specify whether producers, consumers, or both are paused when pausing a destination. Use this option with the <code>pause dst</code> subcommand. Use one of the following values: <table><tr><td>CONSUMERS</td><td>Pause delivery of messages to consumers.</td></tr><tr><td>PRODUCERS</td><td>Pause delivery of messages from producers.</td></tr><tr><td>ALL</td><td>Pause delivery of messages to consumers and from producers.</td></tr></table>	CONSUMERS	Pause delivery of messages to consumers.	PRODUCERS	Pause delivery of messages from producers.	ALL	Pause delivery of messages to consumers and from producers.
CONSUMERS	Pause delivery of messages to consumers.						
PRODUCERS	Pause delivery of messages from producers.						
ALL	Pause delivery of messages to consumers and from producers.						

	If the <code>-pst</code> option is not specified, pauses both consumers and producers (the equivalent of <code>-pst ALL</code>).
<code>-rtm timeout</code>	Specify the timeout period in seconds of an <code>imqcmd</code> subcommand. The default value is 10.
<code>-rtr numRetries</code>	Specify the number of retries attempted after an <code>imqcmd</code> subcommand times out. The default value is 5.
<code>-s</code>	Silent mode. No output is displayed.
	Use this option with any subcommand.
<code>-secure</code>	Specify a secure administration connection to the broker instance. You must first configure the broker to enable a secure connection service. Use this option whenever you want a secure communication with the broker.
<code>-svn serviceName</code>	Specify the service for which the connections are listed. Use this option with the <code>list cxn</code> subcommand.
<code>-t destinationType</code>	Specify the type of a destination: <code>t</code> (topic) or <code>q</code> (queue).
<code>-tmp</code>	Include temporary destinations when listing destinations using the <code>list dst</code> subcommand.
<code>-u name</code>	Specify the administrator user name. If you omit this value, you are prompted for it.
<code>-v</code>	Display version information. Execute nothing else on the command line.

用法

Subcommands and Options	The following subcommands and associated arguments and options are supported: <code>compact dst [-t type -n destName]</code> Compact the flat-file data store for the destination of the specified type and name. If no type and name are specified, all destinations are compacted. Destinations must be paused before they can be compacted. <code>commit txn -n transaction_id</code> Commit the specified transaction

`create dst -t destinationType -n destName [-o attribute=value] [-o attribute=value1]...`

Create a destination of the specified type, with the specified name, and the specified attributes. Destination names must contain only alphanumeric characters (no spaces) and can begin with an alphabetic character or the underscore character (`_`).

`destroy dst -t destinationType -n destName`

Destroy the destination of the specified type and name.

`destroy dur -n subscrName -c client_id`

Destroy the specified durable subscription for the specified Client Identifier.

`list cxn [-s svn serviceName] [-b hostName:port]`

List all connections of the specified service name on the default broker or on a broker at the specified host and port. If the service name is not specified, all connections are listed.

`list dst [-t tmp]`

List all destinations, with option of listing temporary destinations as well.

`list dur -d destination`

List all durable subscriptions for the specified destination.

`list svc`

List all connection services on the broker instance.

`list txn`

List all transactions, being tracked by the broker.

`metrics bkr [-m metricType] [-i int interval] [-msp numSamples]`

Display broker metrics for the broker instance.

Use the `-m` option to specify the type of metric to display. Use one of the following values to specify *metricType*:

`ttl` Specifies the total of messages in and out of the broker (default).

`rts` Provides the same information as `ttl`, but specifies the number of messages per second.

`cxn` Connections, virtual memory heap, threads.

Use the `-i` option to specify the interval (in seconds) at which to display the metrics. The default is 5 seconds.

Use the `-msp` option to specify the number of samples displayed in the output. A value of `-1` means an unlimited number. The default value is `-1`.

`metrics dst -t type -n destName [-m metricType] [-i int interval] [-msp numSamples]`

Displays metrics information for the destination of the specified type and name.

Use the `-m` option to specify the type of metrics to display. Use one of the following values to specify *metricType*:

- `ttl` Specifies the number of messages flowing in and out of the broker and residing in memory.
- `rts` Provides the same information as `ttl`, but specifies the number of messages per second.
- `con` Displays consumer related metrics.
- `dsk` Displays disk usage metrics.

Use the `-int` option to specify the interval (in seconds) at which to display the metrics. The default is 5 seconds.

Use the `-msp` option to specify the number of samples displayed in the output. A value of `-1` means an unlimited number. The default value is 5.

`metrics svc -n serviceName [-m metricType] [-int interval] [-msp numSamples]`

List metrics for the specified service on the broker instance. Use the `-m` option to specify the type of metric to display. Use one of the following values to specify *metricType*:

- `ttl` Total of messages in and out of the broker (default)
- `rts` Provides the same information as `ttl`, but specifies the number of messages per second
- `cxn` Connections, virtual memory heap, threads

Use the `-int` option to specify the interval (in seconds) at which to display the metrics. The default is 5 seconds.

Use the `-msp` option to specify the number of samples displayed in the output. A value of `-1` means an unlimited number. The default value is `-1`.

`pause bkr`

Pause the broker instance.

`pause dst [-t type -n destName] [-pst pauseType]`

Pause the delivery of messages to consumers (`-pst CONSUMERS`), or from producers (`-pst PRODUCERS`), or both (`-pst ALL`), for the destination of the specified type and name. If no destination type or name are specified, all destinations are paused.

`pause svc -n serviceName`

Pause the specified service running on the broker instance. You cannot pause the administrative service.

`purge dst -t destinationType -n destName`

Purge messages at the destination with the specified type and name.

`purge dur -n subscrName -c client_id`

Purge all messages for the specified client identifier.

query bkr

List the current settings of properties of the broker instance. Show the list of running brokers (in a multi-broker cluster) that are connected to the specified broker.

query dst -t *destinationType* -n *destName*

List information about the destination of the specified type and name.

query svc -n *serviceName*

Display information about the specified service running on the broker instance.

query txn -n *transaction_id*

List information about the specified transaction.

reload cls

Forces all the brokers in a cluster to reload the `imq.cluster.brokerlist` property and update cluster information. This subcommand only applies to broker clusters.

restart bkr

Shut down and restart the broker instance. This command restarts the broker using the options specified when the broker was first started. If you want different options to be in effect, you must shut down the broker and then start it again, specifying the options you want.

resume bkr

Resume the broker instance.

resume dst [-t *type*] [-n -destName]

Resumes the delivery of messages for the paused destination of the specified type and name. If no destination type and name are specified, all destinations are resumed.

resume svc -n *serviceName*

Resume the specified service running on the broker instance.

rollback txn -n *transaction_id*

Roll back the specified transaction.

shutdown bkr

Shut down the broker instance

update bkr -o *attribute=value* [-o *attribute=value*]...

Change the specified attributes for the broker instance.

update dst -t *destinationType* -n *destName* -o *attribute=value* [-o *attribute=value1*]...

Update the value of the specified attributes at the specified destination..

update svc -n *serviceName* -o *attribute=value* [-o *attribute=value1*]...

Update the specified attribute of the specified service running on the broker instance.

Attribute Value Pairs

You can specify attributes with the create and update subcommands. Applicable attributes depend on the subcommand arguments.

The following attributes are supported:

Queue (dst):

maxTotalMsgBytes	Value: Integer (maximum total size of messages, in bytes) Default: 0 (unlimited)
maxBytesPerMsg	Value: Integer (maximum size of a single message, in bytes) Default: 0 (unlimited)
maxNumMsgs	Value: Integer (maximum total number of messages) Default: 0 (unlimited)
consumerFlowLimit	Value: Integer Initial number of queued messages sent to active consumers before load-balancing starts A value of -1 means an unlimited number. Default: 1000
isLocalOnly	Value: Boolean (destination limited to delivering messages to local consumers only) Default: false
limitBehavior	Value: Specify how broker responds when memory-limit is reached. Use one of the following values: FLOW_CONTROL Slows down producers REMOVE_OLDEST Purges oldest messages REJECT_NEWEST Rejects the newest messages Default: REJECT_NEWEST
localDeliveryPreferred	Value: Boolean Specify messages be delivered to remote consumers only if there are no consumers on the local broker. Requires that the destination not be restricted to local-only delivery (isLocalOnly = false) Default: false
maxNumActiveConsumers	Value: Integer (maximum number of active consumers in load-balanced delivery) A value of -1 means an unlimited number. Default: 1
maxNumBackupConsumers	Value: Integer (maximum number of backup consumers in load-balanced delivery) A value of -1 means an unlimited number. Default: 0

maxNumProducers	Value: (maximum total number of producers) A value of - 1 means an unlimited number. Default: - 1
useDMQ	Specify whether a destination's dead messages are discarded or put on the dead message queue. Default: true
Topic (dst):	
consumerFlowLimit	Value: Integer Maximum number of messages delivered to a consumer in a single batch. A value of - 1 means an unlimited number. Default: 1000
isLocalOnly	Value: Boolean (destination limited to delivering messages to local consumers only) Default: false
limitBehavior	Value: Specify how broker responds when memory-limit is reached. Use one of the following values: FLOW_CONTROL Slows down producers REMOVE_OLDEST Purges the oldest messages REJECT_NEWEST Rejects the newest messages Default: REJECT_NEWEST
maxBytesPerMsg	Value: Integer (maximum size of a single message, in bytes) Default: 0 (unlimited)
maxNumMsgs	Value: Integer (maximum total number of messages) A value of - 1 means an unlimited number. Default: - 1
maxNumProducers	Value: (maximum total number of producers) Default: 0 (unlimited)
maxTotalMsgBytes	Value: Integer (maximum total size of messages, in bytes) A value of - 1 means an unlimited number. Default: - 1

useDMQ	Specify whether a destination's dead messages are discarded or put on the dead message queue. Default: <code>true</code>
Broker (bkr):	
imq.autocreate.destination.useDMQ	Value: Boolean. Set the useDMQ attribute to <code>true</code> to enable all autocreated physical destinations on a broker to use the dead message queue. Set the useDMQ attribute to <code>false</code> to disable all autocreated physical destinations on a broker from using the dead message queue. Default: <code>true</code>
imq.autocreate.queue	Value: Boolean Default: <code>true</code>
imq.autocreate.queue.maxNumActiveConsumers	Value: Integer (maximum number of consumers that can be active in load-balanced delivery from an autocreated queue destination) A value of -1 means an unlimited number. Default: 1
imq.autocreate.queue.maxNumBackupConsumers	Value: Integer (maximum number of backup consumers that can take the place of active consumers) A value of -1 means an unlimited number. Default: 0
imq.autocreate.topic	Value: Boolean Default: <code>true</code>
imq.cluster.url	Value: String (location of cluster configuration file) Default: none
imq.log.file.rolloverbytes	Value: Integer (maximum size of a log file, in bytes)

	Default: 0 (no rollover based on size)
imq.log.file.rolloversecs	Value: Integer (maximum age of a log file, in seconds)
	Default: 0 (no rollover based on age)
imq.log.level	Value: String (NONE, ERROR, WARNING, INFO)
	Default: INFO
imq.message.max_size	Value: Integer (maximum size of a single message, in bytes)
	Default: 70m
imq.portmapper.port	Value: Integer
	Default: 7676
imq.system.max_count	Value: Integer (maximum total number of messages)
	Default: 0 (no limit)
imq.system.max_size	Value: Integer (maximum total size of messages, in bytes)
	Default: 0 (no limit)
Service (svc):	
maxThreads	Value: Integer (maximum threads assigned)
	Default: Depends on service
minThreads	Value: Integer (minimum threads assigned)
	Default: Depends on service
port	Value: Integer
	Default: 0 (dynamically allocated)

示例

示例 1 Shutting Down a Broker

The following command shuts down a broker for hostname myserver on port 7676:

```
mqcmd shutdown bkr -b myserver:7676
```

示例 2 Restarting a Broker

The following command restarts a broker for hostname `myserver`:

```
imqcmd restart bkr -b myserver
```

示例 3 Pausing a Service

The following command pauses a broker for hostname `localhost` on port `7676`, with a *serviceName* of `jms`:

```
imqcmd pause svc -n jms -b :7676
```

示例 4 Resuming a Service

The following command resumes a service for hostname `localhost` on port `7676`, with a *serviceName* of `jms`:

```
imqcmd resume svc -n jms -b myserver:7676
```

示例 5 Creating a Queue Destination

The following command creates a queue destination for hostname `myserver` on port `7676`, with a *destName* of `myFQ`, a *queueDeliveryPolicy* of `Failover`, and a *maxBytesPerMsg* of `10000`:

```
imqcmd create dst -n myFQ -t q -o "queueDeliveryPolicy=f" \
    -o "maxBytesPerMsg=10000" -b myserver:7676
```

示例 6 Purging a Queue Destination

The following command purges a queue destination for hostname `myserver` on port `7676`, with a *destName* of `myFQ`:

```
imqcmd purge dst -n myFQ -t q -b myserver:7676
```

示例 7 Listing Destinations on a Broker

The following command lists destinations for hostname `myserver` on port `7676`:

```
imqcmd list dst -b myserver:7676
```

示例 8 Updating a Portmapper Port

The following command updates a portmapper port on hostname `myserver` from port `7676` to `7878`:

```
imqcmd update bkr -o "imq.portmapper.port=7878"
```

示例 9 Updating the Maximum Number of Messages in the Queue

The following command updates the maximum number of messages in the queue to `2000` for `myserver` on port `8080` with a *destName* of `TestQueue`:

```
imqcmd update dst -b myserver:8080 -n TestQueue -t q -o "maxNumMsgs=2000"
```

示例 10 Updating the Maximum Threads

The following command updates the maximum threads `jms` connection service to `200` for hostname `localhost` on port `7676`:

```
imqcmd update svc -n jms -o "minThreads=200"
```

示例 11 Listing Durable Subscriptions

The following command lists durable subscriptions for a topic with hostname `localhost` on port `7676` with a *destName* of `myTopic`:

```
imqcmd list dur -d myTopic
```

示例 12 Destroying Durable Subscriptions

The following command destroys subscriptions for hostname `localhost` on port `7676` with a *dursubName* of `myDurSub` and a *client_ID* of `111.222.333.444`:

```
imqcmd destroy dur -n myDurSub -c "111.222.333.444"
```

示例 13 Listing All Transactions

The following command lists all transactions on a broker with hostname `localhost` on port `7676`:

```
imqcmd list txn
```

示例 14 Displaying Information About a Transaction

The following command displays information about a transaction with hostname `localhost` on port `7676`, and a *transactionID* of `1234567890`

```
imqcmd query txn -n 1234567890
```

示例 15 Committing a Transaction

The following command commits a transaction with hostname `localhost` on port `7676`, and a *transactionID* of `1234567890`:

```
imqcmd commit txn -n 1234567890
```

环境变量

The following environment variables affect the execution of this command:

IMQ_JAVAHOME Specify the Java 2 compatible runtime. When this environment variable is not set it defaults to `/usr/j2se`.

退出状态

The following exit values are returned:

- 0** Successful completion.
- >0** An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWiq

另请参见

[imqadmin\(1M\)](#), [imqbrokerd\(1M\)](#), [imqdbmgr\(1M\)](#), [imqkeytool\(1M\)](#), [imqobjmgr\(1M\)](#), [imqusermgr\(1M\)](#), [attributes\(5\)](#)

Sun Java System Message Queue Administrator's Guide

引用名	imqdbmgr – manage a plugged-in JDBC-compliant Message Queue data store																
用法概要	<pre>/usr/bin/imqdbmgr subcommand argument [[option...]] /usr/bin/imqdbmgr -h -help /usr/bin/imqdbmgr -v -version</pre>																
描述	The <code>imqdbmgr</code> utility creates and manages a Java DataBase Connectivity (JDBC) compliant database used for Message Queue persistent storage. The database can be either embedded or external. To use a JDBC-compliant database (and the <code>imqdbmgr</code> utility), you need to first set a number of JDBC-related properties in the broker instance configuration file. See the Sun Java System Message Queue Administrator's Guide for additional information. <code>imqdbmgr</code> supports four management subcommands. These <i>subcommands</i>, and their corresponding <i>arguments</i> and <i>options</i> follow the <code>imqdbmgr</code> command on the command line. See USAGE and OPTIONS. The following subcommands are supported: <table><tr><td>create</td><td>Create a Message Queue database schema.</td></tr><tr><td>delete</td><td>Delete Message Queue database tables in the current data store.</td></tr><tr><td>recreate</td><td>Delete Message Queue database tables and recreate Message Queue database schema in the curent data store.</td></tr><tr><td>reset</td><td>Reset the database table lock to allow other processes to access database tables.</td></tr></table> The <code>imqdbmgr</code> subcommands support the following arguments: <table><tr><td>all</td><td>Indicates the subcommand applies to the data store, as well as the database tables.</td></tr><tr><td>lck</td><td>Indicates the subcommand applies to the database table lock.</td></tr><tr><td>oldtbl</td><td>Indicates the subcommand applies to an older version of the database tables.</td></tr><tr><td>tbl</td><td>Indicates the subcommand applies to the database tables only.</td></tr></table>	create	Create a Message Queue database schema.	delete	Delete Message Queue database tables in the current data store.	recreate	Delete Message Queue database tables and recreate Message Queue database schema in the curent data store.	reset	Reset the database table lock to allow other processes to access database tables.	all	Indicates the subcommand applies to the data store, as well as the database tables.	lck	Indicates the subcommand applies to the database table lock.	oldtbl	Indicates the subcommand applies to an older version of the database tables.	tbl	Indicates the subcommand applies to the database tables only.
create	Create a Message Queue database schema.																
delete	Delete Message Queue database tables in the current data store.																
recreate	Delete Message Queue database tables and recreate Message Queue database schema in the curent data store.																
reset	Reset the database table lock to allow other processes to access database tables.																
all	Indicates the subcommand applies to the data store, as well as the database tables.																
lck	Indicates the subcommand applies to the database table lock.																
oldtbl	Indicates the subcommand applies to an older version of the database tables.																
tbl	Indicates the subcommand applies to the database tables only.																
选项	The following options are supported: <table><tr><td>-b <i>brokerName</i></td><td>Specify the broker instance name and corresponding instance configuration properties. If this option is not specified, the default broker instance is assumed.</td></tr><tr><td></td><td>Use this option with the <code>create</code>, <code>delete</code>, <code>recreate</code>, or <code>reset</code> subcommands.</td></tr><tr><td>-D <i>property=value</i></td><td>Set system property <i>property</i> to <i>value</i>.</td></tr></table>	-b <i>brokerName</i>	Specify the broker instance name and corresponding instance configuration properties. If this option is not specified, the default broker instance is assumed.		Use this option with the <code>create</code> , <code>delete</code> , <code>recreate</code> , or <code>reset</code> subcommands.	-D <i>property=value</i>	Set system property <i>property</i> to <i>value</i> .										
-b <i>brokerName</i>	Specify the broker instance name and corresponding instance configuration properties. If this option is not specified, the default broker instance is assumed.																
	Use this option with the <code>create</code> , <code>delete</code> , <code>recreate</code> , or <code>reset</code> subcommands.																
-D <i>property=value</i>	Set system property <i>property</i> to <i>value</i> .																

	Use this option with the <code>create</code> , <code>delete</code> , <code>recreate</code> , or <code>reset</code> subcommands.
<code>-h</code> <code>-help</code>	Display usage help. Execute nothing else on the command line.
<code>-p password</code>	Specify the database password.
	Use this option with the <code>create</code> , <code>delete</code> , <code>recreate</code> , or <code>reset</code> subcommands.
<code>-u userName</code>	Specify the database user name.
	Use this option with the <code>create</code> , <code>delete</code> , <code>recreate</code> , or <code>reset</code> subcommands.
<code>-v</code> <code>-version</code>	Display version information. Execute nothing else on the command line.

用法

The following subcommands and associated arguments are supported:

<code>create all</code>	Create a new embedded data store and Message Queue database schema for a specified or default broker instance.
<code>create tbl [-u userName] [-p password]</code>	Create Message Queue database schema in an existing data store for a specified or default broker instance.
<code>delete tbl [-u userName] [-p password]</code>	Delete Message Queue database tables in the current data store for a specified or default broker instance.
<code>delete oldtbl [-u userName] [-p password]</code>	Delete the earlier version of Message Queue database tables. Used after the data store has been automatically migrated to the current version of Message Queue.
<code>recreate tbl [-u userName] [-p password]</code>	Delete Message Queue database tables and recreate Message Queue database schema in the current data store for a specified or default broker instance.
<code>reset lck</code>	Reset the database table lock to allow other processes to access database tables.

环境变量

The following environment variables affect the execution of this command:

<code>IMQ_JAVAHOME</code>	Specify the Java 2 compatible runtime. When this environment variable is not set it defaults to <code>/usr/j2se</code> .

- 退出状态

The following exit values are returned:

0

Successful completion.

>0

An error occurred.
- 文件

/var/imq/instances/*brokerName*/dbstore

Recommended directory in which to create an embedded database.
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWiq

另请参见

[imqadmin\(1M\)](#), [imqbrokerd\(1M\)](#), [imqcmd\(1M\)](#), [imqusermgr\(1M\)](#), [imqkeytool\(1M\)](#), [imqobjmgr\(1M\)](#), [attributes\(5\)](#)

Sun Java System Message Queue Administrator's Guide

引用名	imqkeytool – generate a self-signed certificate for secure communication						
用法概要	<pre>/usr/bin/imqkeytool [-broker] [-servlet <i>keystore_location</i>] /usr/bin/imqkeytool -h</pre>						
描述	The <code>imqkeytool</code> utility generates a self-signed certificate for secure communication. The certificate can be used by a broker instance to establish a secure connection with a client, or by a Message Queue-supplied HTTPS servlet to establish a secure connection with a broker instance. An HTTPS servlet is an SSL-enabled variant of the HyperText Transfer Protocol that establishes a secure connection with a broker instance. Without an option, <code>imqkeytool</code> generates a self-signed certificate for a broker instance. <code>imqkeytool</code> uses command line options to specify whether the certificate is used by a broker instance or by a servlet.						
选项	The following options are supported: <table><tr><td>-broker</td><td>Generate a self-signed certificate for the broker and places it in the Message Queue keystore. All broker instances running on a system must use the same certificate.</td></tr><tr><td>-h</td><td>Display usage help. Do not execute anything else on the command line.</td></tr><tr><td>-servlet <i>keystore_location</i></td><td>Generate a self-signed certificate for an HTTPS servlet and places it in <i>keystore_location</i>. <i>keystore_location</i> refers to the location of the keystore. You should move this keystore to a location where it is accessible and readable by the Message Queue HTTPS servlet to establish a secure connection with a broker.</td></tr></table>	-broker	Generate a self-signed certificate for the broker and places it in the Message Queue keystore. All broker instances running on a system must use the same certificate.	-h	Display usage help. Do not execute anything else on the command line.	-servlet <i>keystore_location</i>	Generate a self-signed certificate for an HTTPS servlet and places it in <i>keystore_location</i> . <i>keystore_location</i> refers to the location of the keystore. You should move this keystore to a location where it is accessible and readable by the Message Queue HTTPS servlet to establish a secure connection with a broker.
-broker	Generate a self-signed certificate for the broker and places it in the Message Queue keystore. All broker instances running on a system must use the same certificate.						
-h	Display usage help. Do not execute anything else on the command line.						
-servlet <i>keystore_location</i>	Generate a self-signed certificate for an HTTPS servlet and places it in <i>keystore_location</i> . <i>keystore_location</i> refers to the location of the keystore. You should move this keystore to a location where it is accessible and readable by the Message Queue HTTPS servlet to establish a secure connection with a broker.						
环境变量	The following environment variables affect the execution of this command: <table><tr><td>IMQ_JAVAHOME</td><td>Specify the Java 2 compatible runtime. When this environment variable is not set it defaults to <code>/usr/j2se</code>.</td></tr></table>	IMQ_JAVAHOME	Specify the Java 2 compatible runtime. When this environment variable is not set it defaults to <code>/usr/j2se</code> .				
IMQ_JAVAHOME	Specify the Java 2 compatible runtime. When this environment variable is not set it defaults to <code>/usr/j2se</code> .						
退出状态	The following exit values are returned: <table><tr><td>0</td><td>Successful completion.</td></tr><tr><td>>0</td><td>An error occurred.</td></tr></table>	0	Successful completion.	>0	An error occurred.		
0	Successful completion.						
>0	An error occurred.						
文件	<table><tr><td>/etc/imq/keystore</td><td>Contains Message Queue keystore in which <code>imqkeytool</code> stores a self-signed certificate for brokers.</td></tr></table>	/etc/imq/keystore	Contains Message Queue keystore in which <code>imqkeytool</code> stores a self-signed certificate for brokers.				
/etc/imq/keystore	Contains Message Queue keystore in which <code>imqkeytool</code> stores a self-signed certificate for brokers.						

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWiQ

```
imqadmin(1M),imqbrokerd(1M),imqcmd(1M),imqdbmgr(1M),imqobjmgr(1M),
imqusermgr(1M),attributes(5)
```

手册页第 1M 部分：系统管理命令 • 最后修订 11 Mar 2004

引用名 imqobjmgr – manage Message Queue administered objects

用法概要

```
/usr/bin/imqobjmgr subcommand [[option]...]

/usr/bin/imqobjmgr -i fileName

/usr/bin/imqobjmgr -h | [-H] | -help | -Help

/usr/bin/imqobjmgr -v
```

描述 imqobjmgr manages Message Queue administered objects in an object store accessible using JNDI. Administered objects allow JMS clients to be provider-independent by insulating them from provider-specific naming and configuration formats.

imqobjmgr supports five management subcommands. These *subcommands*, and their corresponding *options* follow the imqobjmgr command on the command line. See USAGE and OPTIONS.

The following subcommands are supported:

add	Add a new administered object
delete	Delete an administered object
list	Display a list of administered objects
query	Display information about administered objects
update	Update administered objects

You can use the -i option to specify the name of an input file that uses java property file syntax to represent all or part of any imqobjmgr subcommand clause. The -f, -s, and -pre options can be used with any imqobjmgr subcommand.

选项 The following options are supported:

-f	Perform action without user confirmation.
-h -help	Display usage help. Execute nothing else on the command line.
-H -Help	Display usage help, attribute list, and examples. Execute nothing else on the command line.
-i <i>fileName</i>	Specify the name of an input file containing all or part of the subcommand clause, specifying object type, lookup name, object attributes, object store attributes, or other options. Use this option for repetitive information, such as object store attributes.
-j <i>attribute=value</i>	Specify attributes necessary to identify and access a JNDI object store.
-javahome	Specify an alternate Java 2 compatible runtime to use. imqobjmgr uses the runtime bundled with the operating system by default.

-l <i>lookupName</i>	Specify the JNDI lookup name of an administered object. This name must be unique in the object store's context.
-o <i>attribute=value</i>	Specify the attributes of an administered object.
-pre	Run command in preview mode. Preview mode indicates what will be done without performing the command.
-r <i>read-only_state</i>	Specify if an administered object is a read-only object. A value of true indicates the administered object is a read-only object. JMS clients cannot modify the attributes of read-only administered objects. The read-only state is set to false by default.
-s	Silent mode. No output is displayed.
-t <i>type</i>	Specify the type of an administered object: q = queue t = topic cf = ConnectionFactory qcf = queueConnectionFactory tcf = topicConnectionFactory xcf = XA ConnectionFactory (distributed transactions) xqcf = XA queueConnectionFactory (distributed transactions) xtcf = XA topicConnectionFactory (distributed transactions) e = SOAP endpoint (used to support SOAP messaging)
-v	Display version information. Execute nothing else on the command line.

用法	This section provides information on subcommands, options, and attribute value pairs.
Subcommands and Options	The following subcommands and corresponding options are supported: <code>add -t <i>type</i> -l <i>lookupName</i> [-o <i>attribute=value</i>]... -j <i>attribute=value</i>...</code> Add a new administered object of the specified type, lookup name , and object attributes to an object store. <code>delete -t <i>type</i> -l <i>lookupName</i> -j <i>attribute=value</i>...</code> Delete an administered object, of the specified type and lookup name from an object store. <code>list [-t <i>type</i>] -j <i>attribute=value</i>...</code> Display a list of administered objects of a specified type, or all administered objects, in an object store. <code>query -l <i>lookupName</i> -j <i>attribute=value</i>...</code> Display information about an administered object of a specified lookup name in an object store.

update -l *lookupName* [-o *attribute=value*]... -j *attribute=value*...

Update the specified attribute values of an administered object of the specified lookup name in an object store.

Attribute Value Pairs

The following attribute value pairs are supported for the specified administered object types:

Type = ConnectionFactories: ConnectionFactory, TopicConnectionFactory, QueueConnectionFactory, XAConnectionFactory, XATopicConnectionFactory, and XAQueueConnectionFactory

imqAckOnAcknowledge

Value: String (true, false, not specified)

Default: not specified

imqAckOnProduce

Value: String (true, false, not specified)

Default: not specified

imqAckTimeout

Value: String (time in milliseconds)

Default: 0 (no timeout)

imqAddressList

Value: String

Default: not specified

imqAddressListBehavior

Value: String

Default: Priority

imqAddressListIterations

Value: Integer

Default: 1

imqBrokerHostName

Used if imqConnectionType is TCP or TLS. This attribute type is only supported in Message Queue 3.0.

Value: String

Default: localhost

imqBrokerHostPort

Used if imqConnectionType is TCP or TLS. This attribute type is only supported in Message Queue 3.0.

Value: Integer

Default: 7676

imqBrokerServicePort

Used if imqConnectionType is TCP or TLS. This attribute type is only supported in Message Queue 3.0.

Value: Integer

Default: 0

imqConfiguredClientID

Value: String (ID number)

Default: no ID specified

imqConnectionFlowCount

Value: Integer

Default: 100

imqConnectionFlowLimit

Value: Integer

Default: 1000

imqConnectionFlowLimitEnabled

Value: Boolean

Default: false

imqConnectionType

This attribute type is only supported in Message Queue 3.0.

Value: String (TCP, TLS, HTTP).

Default: TCP

imqConnectionURL

Used if imqConnectionType is HTTP. This attribute type is only supported in Message Queue 3.0.

Value: String

Default: `http://localhost/imq/tunnel`

imqConsumerFlowLimit

Value: Integer

Default: 1000

imqConsumerFlowThreshold
Value: Integer
Default: 50

imqDefaultPassword
Value: String
Default: guest

imqDefaultUsername
Value: String
Default: guest

imqDisableSetClientID
Value: Boolean
Default: false

imqJMSDeliveryMode
Value: Integer (1=non-persistent, 2=persistent)
Default: 2

imqJMSExpiration
Value: Long (time in milliseconds)
Default: 0 (does not expire)

imqJMSPriority
Value: Integer (0 to 9)
Default: 4

imqLoadMaxToServerSession
Value: Boolean
Default: true

imqOverrideJMSDeliveryMode
Value: Boolean
Default: false

imqOverrideJMSExpiration
Value: Boolean
Default: false

imqOverrideJMSHeadersToTemporaryDestinations
Value: Boolean
Default: false

imqOverrideJMSPriority
Value: Boolean

Default: false

imqQueueBrowserMaxMessagesPerRetrieve
Value: Integer

Default: 1000

imqBrowserRetrieveTimeout
Value: Long (time in milliseconds)

Default: 60,000

imqReconnectAttempts
Value: Integer

Default: 0

imqReconnectEnabled
Value: Boolean

Default: false

imqReconnectInterval
Value: Long (time in milliseconds)

Default: 3000

imqSetJMSXAppID
Value: Boolean

Default: false

imqSetJMSXConsumerTXID
Value: Boolean

Default: false

imqSetJMSXProducerTXID
Value: Boolean

Default: false

imqSetJMSXRcvTimestamp
Value: Boolean

Default: false

imqSetJMSXUserID
Value: Boolean

Default: false

imqSSLIsHostTrusted

Used if imqConnectionType is TLS. This attribute type is only supported in Message Queue 3.0.

Value: Boolean

Default: true

Type = Destinations: Topic and Queue

imqDestinationDescription Value: String

Default: no description

imqDestinationName Value: String

Default: Untitled_Destination_Object

Type = Endpoint (SOAP Endpoint)

imqEndpointDescription Value: String

Default: A description for the endpoint object

imqEndpointName Value: String

Default: Untitled_Endpoint_Object

imqSOAPEndpointList Value: String (one or more space-separated URLs)

Default: no url

示例

示例 1 Adding a Topic Administered Object to an Object Store

Where JNDI lookup name=myTopic and imqDestinationName=MyTestTopic, the following command adds to an LDAP server object store:

```
imqobjmgr add -t t -l "cn=myTopic"\
-o "imqDestinationName=MyTestTopic"\
-j "java.naming.factory.initial=com.sun.jndi.ldap.LdapCtxFactory"\
-j "java.naming.provider.url=ldap://mydomain.com:389/o=imq"
```

Where JNDI lookup name=myTopic and imqDestinationName=MyTestTopic, the following command adds to a file system object store:

```
imqobjmgr add -t -l "cn=myTopic"\
-o "imqDestinationName=MyTestTopic"\
-j \
  "java.naming.factory.initial=com.sun.jndi.fscontext.RefFSCtxFactory"\
-j "java.naming.provider.url=file:/home/foo/imq_admin_objects"
```

示例 1 Adding a Topic Administered Object to an Object Store (续)

Where JNDI lookup name=myTopic and imqDestinationName=MyTestTopic, the following command adds to a file system object store, using an input file:

```
imqobjmgr -i inputfile
```

The associated input file consists of the following:

```
cmdtype=add
obj.type=t
obj.lookupName=cn=myTopic
obj.attrs.imqDestinationName=MyTestTopic
objstore.attrs.java.naming.factory.initial=com.sun.jndi.fscontext.\
    RefFSContextFactory
objstore.attrs.java.naming.provider.url=file:/home/foo/imq_admin_objects
```

示例 2 Adding a QueueConnectionFactory Administered Object to an Object Store

Where JNDI lookup name=myQCF, read-only state=true, imqAddressList=mq://foohost:777/jms, the following command adds to an LDAP server object store:

```
imqobjmgr add -t qf -l "cn=myQCF" -r true\
-o "imqAddressList=mq://foohost:777/jms"\
-j "java.naming.factory.initial=com.sun.jndi.ldap.LdapCtxFactory"\
-j "java.naming.provider.url=ldap://mydomain.com:389/o=imq"
```

Where JNDI lookup name=myQCF, read-only state=true, imqAddressList=mq://foohost:777/jms, the following command adds to an LDAP server object store using an input file:

```
imqobjmgr -i inputfile
```

The associated input file consists of the following:

```
cmdtype=add
obj.type=qf
obj.lookupName=cn=myQCF
obj.readOnly=true
obj.attrs.imqAddressList=mq://foohost:777/jms
objstore.attrs.java.naming.factory.initial=com.sun.jndi.\
    ldap.LdapCtxFactory
objstore.attrs.java.naming.provider.url=ldap://mydomain.com:389/o=imq
```

Where JNDI lookup name=myQCF, read-only state=true, imqAddressList=mq://foohost:777/jms, the following command adds to an LDAP server object store, using both an input file and command options:

示例 2 Adding a QueueConnectionFactory Administered Object to an Object Store (续)

```
imqobjmgr add -t qf -l "cn=myQCF"\
-o "imqAddressList=mq://foohost:777/jms"\
-i inputfile
```

The associated input file consists of the following:

```
objstore.attrs.java.naming.factory.initial=com.sun.jndi.ldap.LdapCtxFactory
objstore.attrs.java.naming.provider.url=ldap://mydomain.com:389/o=imq
```

示例 3 Deleting a Topic Administered Object from an Object Store

Where JNDI lookup name=myTopic and no confirmation is requested, the following command deletes from an LDAP server object store:

```
imqobjmgr delete -f -l "cn=myTopic"\
-j "java.naming.factory.initial=com.sun.jndi.ldap.LdapCtxFactory"\
-j "java.naming.provider.url=ldap://mydomain.com:389/o=imq"
```

示例 4 Querying Information About a Topic Administered Object

Where JNDI lookup name=myTopic, the following command queries from an LDAP server object store using simple authentication scheme:

```
imqobjmgr query -l "cn=myTopic"\
-j "java.naming.factory.initial=com.sun.jndi.ldap.LdapCtxFactory"\
-j "java.naming.provider.url=ldap://mydomain.com:389/o=imq"\
-j "java.naming.security.authentication=simple"\
-j "java.naming.security.principal=uid=foo,ou=imqobjmgr,o=imq"\
-j "java.naming.security.credentials=foo"
```

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWiq

另请参见

[imqadmin\(1M\)](#), [imqcmd\(1M\)](#), [imqbrokerd\(1M\)](#), [imqkeytool\(1M\)](#), [imqusermgr\(1M\)](#), [attributes\(5\)](#)

引用名	imqusermgr – command utility for managing a Message Queue user repository										
用法概要	<pre>/usr/bin/imqusermgr <i>subcommand</i> [[<i>option</i>]...] /usr/bin/imqusermgr -h /usr/bin/imqusermgr -v</pre>										
描述	The <code>imqusermgr</code> utility manages a file-based user repository to authenticate and authorize users of a Message Queue message server. <code>imqusermgr</code> provides subcommands for adding, deleting, updating, and listing user entries in the repository. <code>imqusermgr</code> supports four management subcommands. These <i>subcommands</i>, and their corresponding <i>options</i> follow the <code>imqusermgr</code> command on the command line. See <code>USAGE</code> and <code>OPTIONS</code>. The following subcommands are supported: <table><tr><td>add</td><td>Add a new user and associated password to the repository.</td></tr><tr><td>delete</td><td>Delete a user from the repository.</td></tr><tr><td>list</td><td>Display information users in the repository.</td></tr><tr><td>update</td><td>Update the password or state of a user in the repository.</td></tr></table>	add	Add a new user and associated password to the repository.	delete	Delete a user from the repository.	list	Display information users in the repository.	update	Update the password or state of a user in the repository.		
add	Add a new user and associated password to the repository.										
delete	Delete a user from the repository.										
list	Display information users in the repository.										
update	Update the password or state of a user in the repository.										
选项	The following options are supported: <table><tr><td>-a <i>active_state</i></td><td> Specify if user's state is active or inactive. An inactive user cannot create connections to the Message Queue message server. Valid values for <i>active_state</i> are <code>true</code> or <code>false</code>. Specify <code>true</code> for active or <code>false</code> for inactive. the default is <code>true</code>. Use this option with the <code>update</code> subcommmand. </td></tr><tr><td>-f</td><td> Perform action without user confirmation. Use this option with the <code>delete</code> and <code>update</code> subcommmands. </td></tr><tr><td>-g <i>group</i></td><td> Specify the group of the user. Valid values for <i>group</i> are <code>admin</code>, <code>user</code>, and <code>anonymous</code>. Use this option with the <code>add</code> subcommmand. </td></tr><tr><td>-h</td><td>Display usage help. Exceute nothing else on the command line.</td></tr><tr><td>-i <i>brokerName</i></td><td> Specify the broker instance user repository to which the command applied. If you do not specify <i>brokerName</i>, the default <i>brokerName</i> is assumed. </td></tr></table>	-a <i>active_state</i>	Specify if user's state is active or inactive. An inactive user cannot create connections to the Message Queue message server. Valid values for <i>active_state</i> are <code>true</code> or <code>false</code>. Specify <code>true</code> for active or <code>false</code> for inactive. the default is <code>true</code>. Use this option with the <code>update</code> subcommmand.	-f	Perform action without user confirmation. Use this option with the <code>delete</code> and <code>update</code> subcommmands.	-g <i>group</i>	Specify the group of the user. Valid values for <i>group</i> are <code>admin</code>, <code>user</code>, and <code>anonymous</code>. Use this option with the <code>add</code> subcommmand.	-h	Display usage help. Exceute nothing else on the command line.	-i <i>brokerName</i>	Specify the broker instance user repository to which the command applied. If you do not specify <i>brokerName</i>, the default <i>brokerName</i> is assumed.
-a <i>active_state</i>	Specify if user's state is active or inactive. An inactive user cannot create connections to the Message Queue message server. Valid values for <i>active_state</i> are <code>true</code> or <code>false</code>. Specify <code>true</code> for active or <code>false</code> for inactive. the default is <code>true</code>. Use this option with the <code>update</code> subcommmand.										
-f	Perform action without user confirmation. Use this option with the <code>delete</code> and <code>update</code> subcommmands.										
-g <i>group</i>	Specify the group of the user. Valid values for <i>group</i> are <code>admin</code>, <code>user</code>, and <code>anonymous</code>. Use this option with the <code>add</code> subcommmand.										
-h	Display usage help. Exceute nothing else on the command line.										
-i <i>brokerName</i>	Specify the broker instance user repository to which the command applied. If you do not specify <i>brokerName</i>, the default <i>brokerName</i> is assumed.										

	Use this option with the add, delete, list, and update subcommands.
-p <i>password</i>	Specify user password.
	Use this option with the add and update subcommands.
-s	Silent mode. Display no output
	Use this option with the add, delete, and update subcommands.
-u <i>userName</i>	Specify user name.
	<i>userName</i> cannot contain the following characters: asterisk (*), colon (:), NEWLINE, or RETURN.
	Use this option with the add, delete, update and list subcommands.
-v	Display version information. Execute nothing else on the command line.

用法

The following subcommands and corresponding options are supported:

add -u *userName* -p *password* [-g *group*] [-s] [-i *brokerName*]

Add a new user and associated password to the repository, and optionally specify the user's group.

delete -u *userName* [-s] [-f] [-i *brokerName*]

Delete a user from the repository.

list [-u *user_name*] [-i *brokerName*]

Display information about the specified user in the repository. If no user is specified, all users are displayed.

update -u *userName* -p *password* [-a *state*] [-s] [-f] [-i *brokerName*]

update -u *userName* -a *state* [-p *password*] [-s] [-f] [-i *brokerName*]

Update the password or state (or both) of a user.

环境变量

The following environment variables affect the execution of this command:

IMQ_JAVAHOME Specify the Java 2 compatible runtime. When this environment variable is not set, it defaults to `/usr/j2se`.

退出状态

The following exit values are returned:

0 Successful completion.

>0 An error occurred.

文件

`/etc/imq/passwd` Flat-file user repository.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWiq

另请参见 [imqadmin\(1M\)](#), [imqbrokerd\(1M\)](#), [imqcmd\(1M\)](#), [imqdbmgr\(1M\)](#), [imqkeytool\(1M\)](#), [imqobjmgr\(1M\)](#), [attributes\(5\)](#)

Sun Java System Message Queue Administrator's Guide

引用名 in.chargend – UDP or TCP character generator service daemon

用法概要 in.chargend

FMRI `svc:/internet/chargen:default`

描述 FMRI stands for Fault Management Resource Identifier. It is used to identify resources managed by the Fault Manager. See [fmd\(1M\)](#) and [smf\(5\)](#).

The `in.chargend` service provides the server-side of the character-generator protocol. This protocol is used for debugging and bandwidth measurement and is available on both TCP and UDP transports, through port 19.

The `in.chargend` service is an [inetd\(1M\)](#) [smf\(5\)](#) delegated service. The `in.chargend` detects which transport is requested by examining the socket it is passed by the `inetd` daemon.

TCP-based service Once a connection is established, the `in.chargend` generates a stream of data. Any data received is discarded. The server generates data until the client program terminates the connection. Note that the data flow is limited by TCP flow control mechanisms.

UDP-based service The `in.chargend` listens for UDP datagrams. When a datagram is received, the server generates a UDP datagram in response containing a random number of ASCII characters (ranging from 0 to 512 characters). Any received data is ignored.

The `in.chargend` data consists of a pattern of 72 character lines containing the printable, 7-bit ASCII characters. Each line is terminated with a carriage return and a line feed character.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/legacy-network-services
Interface Stability	Committed

另请参见 [inetd\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

RFC 864

引用名 in.comsat, comsat – biff server

用法概要 /usr/sbin/in.comsat

描述 comsat is the server process which listens for reports of incoming mail and notifies users who have requested to be told when mail arrives. It is invoked as needed by [inetd\(1M\)](#), and times out if inactive for a few minutes.

comsat listens on a datagram port associated with the biff service specification (see [services\(4\)](#)) for one line messages of the form

user@mailbox - offset

If the *user* specified is logged in to the system and the associated terminal has the owner execute bit turned on (by a biff y), the *offset* is used as a seek offset into the appropriate mailbox file, and the first 7 lines or 560 characters of the message are printed on the user's terminal. Lines which appear to be part of the message header other than the From, To, Date, or Subject lines are not printed when displaying the message.

文件 /var/adm/utmpx user access and administration information

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-servers

另请参见 [svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [svcadm\(1M\)](#), [services\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The message header filtering is prone to error.

The in.comsat service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/network/comsat:default

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名 in.daytimed – UDP or TCP daytime protocol service daemon

用法概要 in.daytimed

FMRI svc:/internet/daytime:default

描述 FMRI stands for Fault Management Resource Identifier. It is used to identify resources managed by the Fault Manager. See [fmd\(1M\)](#) and [smf\(5\)](#).

The in.daytimed service provides the server-side of the daytime protocol. This protocol is used for debugging and bandwidth measurement and is available on both TCP and UDP transports, through port 13.

The in.daytimed service is an [inetd\(1M\)](#) [smf\(5\)](#) delegated service. The in.daytimed detects which transport is requested by examining the socket it is passed by the inetd daemon.

TCP-based service Once a connection is established, the in.daytimed generates the current date and time in [ctime\(3C\)](#) format as 7-bit ASCII and sends it through the connection. The server then closes the connection. Any data received from the client side of the connection is discarded.

UDP-based service The in.daytimed listens for UDP datagrams. When a datagram is received, the server generates the current date and time in [ctime\(3C\)](#) format as 7-bit ASCII and inserts it in a UDP datagram sent in response to the client's request. Any received data is ignored.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/legacy-network-services
Interface Stability	Committed

另请参见 [inetd\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

RFC 867

引用名

in.dhcpd – Dynamic Host Configuration Protocol server

用法概要

```
/usr/lib/inet/in.dhcpd [-denv] [-h relay_hops] [-i interface, ...]
[-l syslog_local_facility] [-b automatic | manual]
[-o DHCP_offer_time] [-t dhcptab_rescan_interval]

/usr/lib/inet/in.dhcpd [-dv] [-h relay_hops] [-i interface,]...
[-l syslog_local_facility] -r IP_address | hostname, ...
```

描述

in.dhcpd is a daemon that responds to Dynamic Host Configuration Protocol (DHCP) requests and optionally to BOOTP protocol requests. The daemon forks a copy of itself that runs as a background process. It must be run as root. The daemon has two run modes, DHCP server (with optional BOOTP compatibility mode) and BOOTP relay agent mode.

The first line in the SYNOPSIS section illustrates the options available in the DHCP/BOOTP server mode. The second line in the SYNOPSIS section illustrates the options available when the daemon is run in BOOTP relay agent mode.

The DHCP and BOOTP protocols are used to provide configuration parameters to Internet hosts. Client machines are allocated their IP addresses as well as other host configuration parameters through this mechanism.

The DHCP/BOOTP daemon manages two types of DHCP data tables: the dhcptab configuration table and the DHCP network tables.

See [dhcptab\(4\)](#) regarding the dhcptab configuration table and [dhcp_network\(4\)](#) regarding the DHCP network tables.

The dhcptab contains macro definitions defined using a termcap-like syntax which permits network administrators to define groups of DHCP configuration parameters to be returned to clients. However, a DHCP/BOOTP server always returns hostname, network broadcast address, network subnet mask, and IP maximum transfer unit (MTU) if requested by a client attached to the same network as the server machine. If those options have not been explicitly configured in the dhcptab, in.dhcpd returns reasonable default values.

The dhcptab is read at startup, upon receipt of a SIGHUP signal, or periodically as specified by the -t option. A SIGHUP (sent using the command `svcadm refresh network/dhcp-server`) causes the DHCP/BOOTP daemon to reread the dhcptab within an interval from 0-60 seconds (depending on where the DHCP daemon is in its polling cycle). For busy servers, users should run `svcadm restart network/dhcp-server` to force the dhcptab to be reread.

The DHCP network tables contain mappings of client identifiers to IP addresses. These tables are named after the network they support and the datastore used to maintain them.

The DHCP network tables are consulted during runtime. A client request received from a network for which no DHCP network table exists is ignored.

This daemon is obsolete and is subject to removal in a future release of Oracle Solaris. Scripts, programs, or procedures that use this command will likely need modification when upgrading to future Solaris software releases. The command line options provided with the in.dhcpd

daemon are used only for the current session, and include only some of the server options you can set. The `dhcpsvc.conf(4)` contains all the server default settings, and can be modified by using the `dhcpgmr` utility. See `dhcpsvc.conf(4)` for more details.

选项

The following options are supported:

- b *automatic | manual* This option enables BOOTP compatibility mode, allowing the DHCP server to respond to BOOTP clients. The option argument specifies whether the DHCP server should automatically allocate permanent lease IP addresses to requesting BOOTP clients if the clients are not registered in the DHCP network tables (*automatic*) or respond only to BOOTP clients who have been manually registered in the DHCP network tables (*manual*). This option only affects DHCP server mode.
- d Debugging mode. The daemon remains as a foreground process, and displays verbose messages as it processes DHCP and/or BOOTP datagrams. Messages are displayed on the current TTY. This option can be used in both DHCP/BOOTP server mode and BOOTP relay agent mode.
- h *relay_hops* Specifies the maximum number of relay agent hops that can occur before the daemon drops the DHCP/BOOTP datagram. The default number of relay agent hops is 4. This option affects both DHCP/BOOTP server mode and BOOTP relay agent mode.
- i *interface,...* Selects the network interfaces that the daemon should monitor for DHCP/BOOTP datagrams. The daemon ignores DHCP/BOOTP datagrams on network interfaces not specified in this list. This option is only useful on machines that have multiple network interfaces. If this option is not specified, then the daemon listens for DHCP/BOOTP datagrams on all network interfaces. The option argument consists of a comma-separated list of interface names. It affects both DHCP/BOOTP server and BOOTP relay agent run modes.
- l *syslog_local_facility* The presence of this option turns on transaction logging for the DHCP server or BOOTP relay agent. The value specifies the `syslog` local facility (an integer from 0 to 7 inclusive) the DHCP daemon should use for tagging the transactions. Using a facility separate from the `LOG_DAEMON` facility allows the network administrator to capture these transactions separately from other DHCP daemon events for such

purposes as generating transaction reports. See [syslog\(3C\)](#), for details about local facilities. Transactions are logged using a record with 9 space-separated fields as follows:

1. Protocol:

Relay mode: "BOOTP"
Server mode: "BOOTP" or "DHCP" based upon client type.

2. Type:

Relay mode: "RELAY-CLNT", "RELAY-SRVR"
Server mode: "ASSIGN", "EXTEND", "RELEASE",
"DECLINE", "INFORM", "NAK" "ICMP-ECHO."

3. Transaction time: absolute time in seconds (unix time)

4. Lease time:

Relay mode: Always 0.
Server mode: 0 for ICMP-ECHO events, absolute time in seconds (unix time) otherwise

5. Source IP address: Dotted Internet form

Relay mode: Relay interface IP on RELAY-CLNT, INADDR_ANY on RELAY-SRVR.
Server mode: Client IP.

6. Destination IP address: Dotted Internet form

Relay mode: Client IP on RELAY-CLNT, Server IP on RELAY-SRVR.
Server mode: Server IP.

7. Client Identifier: Hex representation (0-9, A-F)

Relay mode: MAC address
Server mode: BOOTP - MAC address; DHCP - client id

8. Vendor Class identifier (white space converted to periods (.)).

Relay mode: Always "N/A"
Server mode: Vendor class ID tokenized by converting white space characters to periods (.)

9. MAC address: Hex representation (0-9, A-F)

Relay mode: MAC address

Server mode: MAC address

The format of this record is subject to change between releases.

Transactions are logged to the console if daemon is in debug mode (-d).

Logging transactions impact daemon performance.

It is suggested that you periodically rotate the DHCP transaction log file to keep it from growing until it fills the filesystem. This can be done in a fashion similar to that used for the general system message log `/var/adm/messages` and is best accomplished using the facilities provided by [logadm\(1M\)](#).

-n Disable automatic duplicate IP address detection. When this option is specified, the DHCP server does not attempt to verify that an IP address it is about to offer a client is not in use. By default, the DHCP server pings an IP address before offering it to a DHCP/BOOTP client, to verify that the address is not in use by another machine.

-o *DHCP_offer_time* Specifies the number of seconds the DHCP server should cache the offers it has extended to discovering DHCP clients. The default setting is 10 seconds. On slow network media, this value can be increased to compensate for slow network performance. This option affects only DHCP server mode.

-r *IP_address | hostname, . . .* This option enables BOOTP relay agent mode. The option argument specifies a comma-separated list of IP addresses or hostnames of DHCP or BOOTP servers to which the relay agent is to forward BOOTP requests. When the daemon is started in this mode, any DHCP tables are ignored, and the daemon simply acts as a BOOTP relay agent.

A BOOTP relay agent listens to UDP port 68, and forwards BOOTP request packets received on this port to the destinations specified on the command line. It supports the BROADCAST flag described in RFC 1542. A BOOTP relay

agent can run on any machine that has knowledge of local routers, and thus does not have to be an Internet gateway machine.

Note that the proper entries must be made to the `netmasks` database so that the DHCP server being served by the BOOTP relay agents can identify the subnet mask of the foreign BOOTP/DHCP client's network. See [netmasks\(4\)](#) for the format and use of this database.

`-t dhcptab_rescan_interval`

Specifies the interval in minutes that the DHCP server should use to schedule the automatic rereading of the `dhcptab` information. Typically, you would use this option if the changes to the `dhcptab` are relatively frequent. Once the contents of the `dhcptab` have stabilized, you can turn off this option to avoid needless reinitialization of the server.

`-v`

Verbose mode. The daemon displays more messages than in the default mode. Note that verbose mode can reduce daemon efficiency due to the time taken to display messages. Messages are displayed to the current TTY if the debugging option is used; otherwise, messages are logged to the `syslogd` facility. This option can be used in both DHCP/BOOTP server mode and BOOTP relay agent mode.

示例

示例 1 Starting a DHCP Server in BOOTP Compatibility Mode

The following command starts a DHCP server in BOOTP compatibility mode, permitting the server to automatically allocate permanent IP addresses to BOOTP clients which are not registered in the server's table; limits the server's attention to incoming datagrams on network devices `le2` and `tr0`; drops BOOTP packets whose hop count exceeds 2; configures the DHCP server to cache extended DHCP offers for 15 seconds; and schedules `dhcptab` rescans to occur every 10 minutes:

```
# in.dhcpd -i le2,tr0 -h 2 -o 15 -t 10 -b automatic
```

示例 2 Starting the Daemon in BOOTP Relay Agent Mode

The following command starts the daemon in BOOTP relay agent mode, registering the hosts `bladerunner` and `10.0.0.5` as relay destinations, with debugging and verbose modes enabled, and drops BOOTP packets whose hop count exceeds 5:

```
# in.dhcpd -d -v -h 5 -r bladerunner,10.0.0.5
```

文件

- `/etc/inet/dhcpsvc.conf`
- `/etc/init/hosts`
- `/usr/lib/inet/dhcp/nsu/rfc2136.so.1`

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/dhcp
Interface Stability	Obsolete

另请参见

[svcs\(1\)](#), [cron\(1M\)](#), [dhcpgmr\(1M\)](#), [dhtadm\(1M\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [logadm\(1M\)](#), [pntadm\(1M\)](#), [svcadm\(1M\)](#), [syslogd\(1M\)](#), [syslog\(3C\)](#), [dhcpsvc.conf\(4\)](#), [dhcp_network\(4\)](#), [dhcptab\(4\)](#), [ethers\(4\)](#), [hosts\(4\)](#), [netmasks\(4\)](#), [nsswitch.conf\(4\)](#), [attributes\(5\)](#), [dhcp\(5\)](#), [smf\(5\)](#)

《Oracle Solaris 管理：IP 服务》

Alexander, S., and R. Droms, DHCP Options and BOOTP Vendor Extensions, RFC 2132, Silicon Graphics, Inc., Bucknell University, March 1997.

Droms, R., Interoperation Between DHCP and BOOTP, RFC 1534, Bucknell University, October 1993.

Droms, R., Dynamic Host Configuration Protocol, RFC 2131, Bucknell University, March 1997.

Wimer, W., Clarifications and Extensions for the Bootstrap Protocol, RFC 1542, Carnegie Mellon University, October 1993.

附注

The `in.dhcpd` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/dhcp-server
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名 in.discardd – UDP or TCP discard protocol service

用法概要 in.discardd

FMRI svc:/internet/discard:default

描述 FMRI stands for Fault Management Resource Identifier. It is used to identify resources managed by the Fault Manager. See [fmd\(1M\)](#) and [smf\(5\)](#).

The in.discardd service provides the server-side of the discard protocol. This protocol is used for debugging and bandwidth measurement and is available on both TCP and UDP transports through port 9.

The in.discardd service is an [inetd\(1M\)](#) [smf\(5\)](#) delegated service. The in.discardd detects which transport is requested by examining the socket it is passed by the inetd daemon.

The discard service simply throws away any data it receives from the client.

TCP-based service Once a connection is established, the in.discardd discards any data received. No response is generated. The connection remains open until the client terminates it.

UDP-based service The in.discardd listens for UDP datagrams. When a datagram is received, the server discards it. No response is sent.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/legacy-network-services
Interface Stability	Committed

另请参见 [inetd\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

RFC 863

引用名 in.echod – UDP or TCP echo protocol service daemon

用法概要 in.echod

FMRI `svc:/internet/echo:default`

描述 FMRI stands for Fault Management Resource Identifier. It is used to identify resources managed by the Fault Manager. See [fmd\(1M\)](#) and [smf\(5\)](#).

The `in.echod` service provides the server-side of the echo protocol. This protocol is used for debugging and bandwidth measurement and is available on both TCP and UDP transports, through port 7.

The `in.echod` service is an [inetd\(1M\)](#) [smf\(5\)](#) delegated service. The `in.echod` detects which transport is requested by examining the socket it is passed by the `inetd` daemon.

TCP-based service Once a connection is established, the `in.echod` echoes any data received from the client back to the client. The server echoes data until the client program terminates the connection.

UDP-based service The `in.echod` listens for UDP datagrams. When a datagram is received, the server creates a UDP datagram containing the data it received and sends it to the client.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/legacy-network-services
Interface Stability	Committed

另请参见 [inetd\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

RFC 862

引用名 inetadm – 观察或配置 inetd 控制的服务

用法概要

```
inetadm
inetadm -?
inetadm -p
inetadm -l {FMRI | pattern}
inetadm -e {FMRI | pattern}
inetadm -d {FMRI | pattern}
inetadm -m {FMRI | pattern}... {name=value}...
inetadm -M {name=value}...
```

描述 inetadm 实用程序针对 inetd 管理的 SMF 服务提供以下功能：

- 提供所安装的所有此类服务的列表。
- 列出服务的属性和值。
- 允许启用和禁用服务。
- 允许修改服务的属性值以及 inetd 提供的缺省值。

有关 SMF 服务的说明，请参见 [smf\(5\)](#)。

如果没有参数，inetadm 会列出由 [inetd\(1M\)](#) 控制的所有服务，包括它们当前的运行状态以及它们是否已启用等属性。

选项

对于接受一个或多个 FMRI 操作数的选项（有关 FMRI 的说明，请参见 [smf\(5\)](#)），如果操作数指定的是一个服务（而不是一个服务实例）并且该服务只有一个实例，inetadm 会对该实例进行操作。

如果提供了一个包含多个实例的服务名称，或提供了一个与多个实例匹配的模式，则会显示一条警告消息并忽略该操作数。

对于那些接受 *name=value* 参数的选项，可在 [inetd\(1M\)](#) 手册页中查看每个可能的名称及允许值的说明。

所支持的选项如下：

-?	显示用法消息。
-p	以 <i>name=value</i> 对的形式列出 inetd 提供的所有缺省 inet 服务属性值。如果该值为布尔型的，则会将其列出为 TRUE 或 FALSE。
-l {FMRI pattern}...	以 <i>name=value</i> 对的形式列出指定服务实例的所有属性。此外，如果属性值是从 inetd 提供的缺省值继承而来的，则会通过标记来标识 <i>name= value</i> 对（缺省）。当属性没有指定的服务实例缺省值时，将会发生属性继承。

<code>-e {FMRI pattern}...</code>	启用指定的服务实例。
<code>-d {FMRI pattern}...</code>	禁用指定的服务实例。
<code>-m {FMRI pattern}...{name=value}...</code>	更改所标识的服务实例的指定属性的值。应使用以空格分隔的 <i>name=value</i> 对的形式来指定属性。要删除特定于实例的某个值并接受属性的缺省值，只需指定属性而不指定其值，例如， <i>name=</i> 。
<code>-M {name=value}...</code>	更改指定的 <i>inetd</i> 缺省属性的值。应使用以空格分隔的 <i>name=value</i> 对的形式来指定属性。

示例

示例1 显示服务的属性

以下命令显示 *spray* 服务的属性。

```
# inetadm -l network/rpc/spray:default
SCOPE      NAME=VALUE
           name="sprayd"
           endpoint_type="tli"
           proto="datagram_v"
           isrpc=TRUE
           rpc_low_version=1
           rpc_high_version=1
           wait=TRUE
           exec="/usr/lib/netsvc/spray/rpc.sprayd"
           user="root"
default    bind_addr=""
default    bind_fail_max=-1
default    bind_fail_interval=-1
default    max_con_rate=-1
default    max_copies=-1
default    con_rate_offline=-1
default    failrate_cnt=40
default    failrate_interval=60
default    inherit_env=TRUE
default    tcp_trace=FALSE
default    tcp_wrappers=FALSE
default    connection_backlog=10
```

示例2 显示缺省属性

以下命令显示缺省属性。

```
# inetadm -p
NAME=VALUE
bind_addr=""
bind_fail_max=-1
bind_fail_interval=-1
```

示例2 显示缺省属性 (续)

```

max_con_rate=-1
max_copies=-1
con_rate_offline=-1
failrate_cnt=40
failrate_interval=60
inherit_env=TRUE
tcp_trace=FALSE
tcp_wrappers=FALSE
default connection_backlog=10

```

示例3 更改服务的属性值

以下命令将 spray 服务的 rpc_high_version 更改为 3，将 tcp_trace 更改为 TRUE。

```

# inetadm -m network/rpc/spray:default \
    rpc_high_version=3 tcp_trace=TRUE
# inetadm -l network/rpc/spray:default
SCOPE    NAME=VALUE
         name="sprayd"
         endpoint_type="tli"
         proto="datagram_v"
         isrpc=TRUE
         rpc_low_version=1
         rpc_high_version=3
         wait=TRUE
         exec="/usr/lib/netsvc/spray/rpc.sprayd"
         user="root"
default bind_addr=""
default bind_fail_max=-1
default bind_fail_interval=-1
default max_con_rate=-1
default max_copies=-1
default con_rate_offline=-1
default failrate_cnt=40
default failrate_interval=60
default inherit_env=TRUE
         tcp_trace=TRUE
default tcp_wrappers=FALSE
default connection_backlog=10

```

退出状态

将返回下列退出值：

- 0 操作已成功完成。
- 1 发生了致命错误。伴随产生的错误消息将提供进一步的信息。
- 2 提供的参数无效，例如模糊的服务 FMRI 或模式。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）

另请参见 [inetd\(1M\)](#)、[svccfg\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

引用名

inetconv – 将 `inetd.conf` 项转换为 smf 服务清单，将其导入 smf 系统信息库

用法概要

```
inetconv -?
inetconv [-f] [-n] [-i srcfile] [-o destdir]
inetconv -e [-n] [-i srcfile]
```

描述

`inetconv` 实用程序将包含 `inetd.conf(4)` 记录的文件转换成 `smf(5)` 服务清单，然后将这些清单导入到 smf 系统信息库。`inetd.conf` 文件转换完成后，更改 inet 服务某些方面的唯一方法就是使用 `inetadm(1M)` 实用程序。

输入文件中的服务行与生成的清单具有一对一的关系。缺省情况下，清单使用以下模板命名：

```
<svcname>.<proto>.xml
```

`<svcname>` 标记替换为服务的名称，`<proto>` 标记替换为服务的协议。资源行的服务名称或协议中的斜杠 (/) 替换为下划线 (_)。

服务行记录为已转换服务的属性。

转换过程中，如果发现服务行格式不对或是用于内部 `inetd` 服务，则不会生成清单，并跳过该服务行。

转换过程不会改动输入文件。

选项

所支持的选项如下：

- ? 显示用法消息。
- e 启用输入文件中所列的 smf 服务。
- f 如果指定该选项，则目标目录中发现与要生成的服务清单具有相同名称的服务清单时，`inetconv` 将会覆盖该清单。否则会生成一则错误消息，且不会执行此服务的转换。
- i *srcfile* 允许指定一个备用输入文件 *srcfile*。如果未指定此选项，`inetd.conf(4)` 文件会用作输入。
- n 转换过程中禁用生成清单的自动导入。然后，如果希望将生成的清单导入 `smf(5)` 系统信息库，可以通过 `svccfg(1M)` 实用程序进行此操作。
如果指定 `-e` 选项，`-n` 选项仅显示将要启用的 smf 服务。
- o 允许为生成的清单指定备用目标目录 *destdir*。如果未指定此选项，清单将会放在 `/lib/svc/manifest/network/rpc` 中，如果服务是 RPC 服务，则会放在 `/lib/svc/manifest/network` 中。

示例

示例1 从 `inetd.conf` 生成 `smf` 清单

以下命令从 `inetd.conf(4)` 生成 `smf(5)` 清单，并将它们放在 `/var/tmp` 中，覆盖任何之前存在的相同名称的清单，然后将它们导入 `smf` 系统信息库。

```
# inetconv -f -o /var/tmp
100232/10 -> /var/tmp/100232_10-rpc_udp.xml
Importing 100232_10-rpc_udp.xml ...Done
telnet -> /var/tmp/telnet-tcp6.xml
Importing telnet-tcp6.xml ...Done
```

示例2 从备用输入文件生成清单

以下命令指定一个不同的输入文件，不会将得到的清单装入 `smf` 系统信息库。

```
# inetconv -n -i /export/test/inet.svcs -o /var/tmp
100232/10 -> /var/tmp/100232_10-rpc_udp.xml
telnet -> /var/tmp/telnet-tcp6.xml
```

退出状态

将返回下列退出值：

- 0 操作成功完成（没有错误）。
- 1 指定选项无效。
- 2 一个或多个服务行格式不对，因此未为其生成清单。
- 3 导入一个或多个生成的清单时发生错误。
- 4 发生了系统错误。

文件

`/lib/svc/manifest/network/{rpc}/<svcname>-<proto>.xml`
缺省输出清单文件名。

属性

有关以下属性的说明，请参见 `attributes(5)`：

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）

另请参见

`inetadm(1M)`、`inetd(1M)`、`svccfg(1M)`、`inetd.conf(4)`、`attributes(5)`、`smf(5)`

引用名 `inetd` – Solaris 管理工具委托的面向 `inet` 服务的重启程序

用法概要 `inetd` [*configuration-file*] `start` | `stop` | `refresh`

`svc:/network/inetd:default`

描述 `inetd` 是服务管理工具 (Service Management Facility, SMF) 委托的面向 Internet 服务的重启程序。其基本职责是管理服务状态以便响应管理请求、系统故障和服务故障；并在适当时侦听服务的网络请求。

服务不再通过编辑 `inetd` 配置文件 `inetd.conf(4)` 进行管理，而应使用 `inetconv(1M)` 将配置文件内容转换为 SMF 格式服务，然后使用 `inetadm(1M)` 和 `svcadm(1M)` 管理这些服务。`inetconv` 在对服务进行转换之后，对 `inetd` 配置文件中的旧数据所做的任何更改都不会生效。但是，当 `inetd` 发现配置文件进行了更改时，会通知管理员。有关详细信息，请参见“`inetd` 方法”一节下有关 `start` 方法的说明。

另请注意，当前 `inetd` 无法在 SMF 外部运行。这表示无法从命令行运行 `inetd`，而以前的 `inetd` 支持从命令行运行。如果您尝试执行此操作，则会向标准错误发送一条消息，该消息显示以前的 `inetd` 所支持的选项到 `inetd` 的 SMF 版本的映射。

`inetd` 代表处于 `online` 或 `degraded` 状态的所有服务侦听连接。当用户启用服务并且 `inetd` 尝试代表服务进行侦听时，该服务将进入上述状态之一。如果其他服务器（不管是单独的 Internet 服务还是第三方 Internet 服务）已在侦听同一端口，侦听尝试可能会失败。如果发生此种情况，`inetd` 会记录此条件，并继续尝试按配置的次数在配置的时间间隔绑定到此端口。有关更多详细信息，请参见下面“服务属性”下的属性 `bind_fail_max`。

启动 `inetd` 时，会读取由 `inetd` 管理的所有 SMF 服务的配置。刷新 `inetd` 时，会重新读取这些配置，在响应 SMF 请求或当 `inetd` 收到 `SIGHUP` 信号时会出现此种情况。有关配置刷新的行为，请参见“`inetd` 方法”下有关 `refresh` 方法的说明。

使用 `inetadm(1M)` 或 `svccfg(1M)` 实用程序可以对 SMF 系统信息库中的 Internet 服务进行配置更改。相较于 `svccfg` 而言，`inetadm` 的优势在于它提供了 Internet/RPC 服务上下文。

服务状态 `inetd` 为其管理的每个服务实现一台状态机，作为其服务管理职责的一部分。此计算机中的状态由 `smf(5)` 状态集组成。这些状态的语义如下所示：

uninitialized (未初始化)

`inetd` 尚未处理此服务。

online (联机)

服务正在处理新的网络请求，并且可能具有活动的现有连接。

degraded (降级)

服务可以侦听和处理为此服务指定的部分（而非所有）协议的请求，耗尽了其侦听重试次数，从而进入此状态。现有网络连接可能处于活动状态。

offline (脱机)

连接可能处于活动状态，但未在处理新请求。这是瞬时状态。服务可能因下列任何原因而处于 **offline** (脱机) 状态：

- 不符合服务的相关性。如果符合服务相关性，将重新评估此服务的状态。
- 服务已超出其配置的连接速率限制 `max_con_rate`。服务的连接脱机计时器 `con_rate_offline` 过期时，将重新评估此服务的状态。
- 服务已达到其允许的活动连接数 `max_copies`。活动连接数降至 `max_copies` 以下时，将重新评估此服务的状态。
- `inetd` 无法代表服务侦听其所有协议。如上所述，`inetd` 最多按配置的最大次数在配置的时间间隔重试。侦听尝试成功或达到重试限制时，将重新评估此服务的状态。

disabled (禁用)

管理员已禁用服务，该服务不接受新连接，并且没有任何连接处于活动状态。需要管理员介入才能退出此状态。

maintenance (维护)

服务处于此状态的原因包括：服务错误运行且需要管理员注意，或者管理员已请求此状态。

导致错误运行的事件包括：`inetd` 在超出服务的绑定重试限制之前无法代表任何内容侦听服务的任何协议，非启动方法返回非成功返回值，以及服务超出其故障速率。

请求 **maintenance** 状态的目的是为了对服务执行维护，例如，应用修补级别。在此状态下，不会处理新请求，但现有连接可能处于活动状态。需要管理员介入才能退出此状态。

使用 `inetadm(1M)` 获取受管服务的当前状态。

服务方法

作为状态转换的一部分，`inetd` 将执行服务提供的某个方法集（如果提供）。支持的方法集包括：

inetd_start

执行此方法可处理 **online** 或 **degraded** 服务的请求。由于没有单独的状态来区分服务和活动连接，因此此方法不会作为状态转换的一部分执行。

inetd_offline

在服务从 **online** 或 **degraded** 状态转至 **offline** 状态时执行。对于在执行时执行其自己的侦听的 **wait** 类型的服务，此方法应当会导致此服务停止侦听。在禁用 **online/degraded** 服务的情况下，将在 **disable** 方法之前执行此方法。对于 **wait** 类型的服务，必须实现此方法。

inetd_online

在服务从 **offline** 状态切换到 **online** 状态时执行。此方法允许服务作者在服务开始处理请求之前执行某些准备工作。

inetd_disable

在服务从 **offline** 状态切换到 **disabled** 状态时执行。此方法应当会导致服务的任何活动连接终止。

inetd_refresh

在同时满足下列两个条件时执行：

- 通过框架或 SIGHUP 刷新 **inetd**，或者请求开始刷新服务，并且
- 服务当前处于 **online** 状态，并且没有需要将服务置于 **offline** 状态并重新联机的任何配置更改。

唯一强制性方法是 **inetd_start** 方法。在没有任何其他方法的情况下，**inetd** 不会运行任何方法，但其行为就像成功运行了某个方法一样。

服务属性

SMF 受管服务的配置存储在 SMF 系统信息库中。此配置包括服务的基本配置，服务的每个方法的配置以及缺省配置适用于所有 **inetd** 受管服务。

有关查看和修改服务配置及缺省值的详细信息，请参见 **inetadm(1M)**。

服务的基本配置存储在服务中名为 **inetd** 的属性组内。构成基本配置的这些属性如下所示：

bind_addr

服务应绑定到的网络接口的地址。导致服务接受任何网络接口上的连接的空字符串值。

bind_fail_interval

失败绑定尝试与重置之间的时间间隔（以秒为单位）。值 **0** 和 **-1** 指定是否不会尝试重试，并且处理第一次故障的方式是否与处理超出 **bind_fail_max** 的方式相同。

bind_fail_max

inetd 在放弃前重试绑定到服务的关联端口的最大次数。值 **-1** 指定不会强制任何重试限制。如果在达到任何强制限制前未绑定到服务协议，该服务将进入 **maintenance** 状态；否则，如果未绑定到所有协议，该服务将进入 **degraded** 状态。

con_rate_offline

服务超出其配置的最大连接速率 **max_con_rate** 时将保持脱机的时间（以秒为单位）。值 **0** 和 **-1** 指定是否禁用连接速率限制。

connection_backlog

后备队列大小。表示可在服务器的侦听端点中排队的传入客户机请求数目的限制。

endpoint_type

服务使用的套接字的类型或表示基于 TLI 的服务的值 **tli**。有效套接字类型值包括：**stream**、**dgram**、**raw**、**seqpacket**。

failrate_cnt

服务的故障速率限制的计数部分。故障速率限制适用于 **wait** 类型的服务，当在给定时间内启动服务的 **count** 实例时，将达到此限制。超出此速率将导致服务转换到 **maintenance** 状态。这与前面的 **inetd** 的行为不同，它会继续无限期地在每 10 分钟

重试一次。`failrate_cnt` 在使用服务请求之前检查错误运行的服务器的计数，这些服务器将继续重新启动，从而耗费系统资源。故障速率等效于以前的 `inetd` 的 `-r` 选项。值 `0` 和 `-1` 指定是否禁用此功能。

`failrate_interval`

服务的故障速率的时间部分（以秒为单位）。值 `0` 和 `-1` 指定是否禁用此故障速率限制功能。

`inherit_env`

如果为 `true`，则将 `inetd` 的环境传递到服务的启动方法。不管此设置如何，`inetd` 都会在启动方法环境中设置变量 `SMF_FMRI`、`SMF_METHOD` 和 `SMF_RESTARTER`，以及在方法上下文中设置的任何环境变量。[smf_method\(5\)](#) 中介绍了这些变量。

`isrpc`

如果为 `true`，则为 RPC 服务。

`max_con_rate`

`nowait` 类型的服务允许的最大连接速率（以每秒连接数为单位）。值 `0` 和 `-1` 指定是否禁用连接速率限制。

`max_copies`

可并行运行的 `nowait` 服务的最大副本数。值 `0` 和 `-1` 指定是否禁用此副本限制。

`name`

可设置为以下值之一：

- `getservbyname(3SOCKET)` 已知的服务名称。
- 如果 `isrpc` 设置为 `true`，则为 `getrpcbyname(3NSL)` 已知的服务名称。
- 如果 `isrpc` 设置为 `true`，则为有效 RPC 程序编号。

`proto`

对于基于套接字的服务，此属性是服务支持的协议的列表。有效协议包括：`tcp`、`tcp6`、`tcp6only`、`udp`、`udp6` 和 `udp6only`。对于 TLI 服务，这是服务支持的 `getnetconfigent(3NSL)` 所识别的 `netid` 的列表以及值 `tcp6only` 和 `udp6only`。在此列表中，RPC/TLI 服务还支持 `nettype`，并且 `inetd` 首先尝试将列表成员解释为这些服务类型的 `nettype`。值 `tcp6only` 和 `udp6only` 是 `inetd` 中的新值，这些值请求 `inetd` 仅侦听并传递真正的 IPv6 请求（而非 IPv4 映射请求）。请参见下面的“为基于套接字的服务配置协议”。

`rpc_low_version`

支持的最低 RPC 版本。此参数在 `isrpc` 设置为 `true` 时是必需的。

`rpc_high_version`

支持的最高 RPC 版本。此参数在 `isrpc` 设置为 `true` 时是必需的。

`tcp_keepalive`

如果为 `true`，则允许在连接的套接字上定期传输消息。如果连接方无法响应这些消息，则将连接视为已断开。这仅适用于 `endpoint_type` 设置为 `streams` 且 `wait` 设置为 `false` 的服务。

tcp_trace

如果为 **true**，并且是 **nowait** 类型的服务，**inetd** 则会使用 **syslog(3C)** 工具针对每个传入连接记录客户机的 IP 地址和 TCP 端口号以及服务名称。**inetd** 使用 **syslog** 工具 **code** 守护进程以及 **notice** 优先级级别。有关 **syslog** 代码和严重级别的说明，请参见 **syslog.conf(4)**。此日志独立于 TCP 包装器工具执行的记录。

tcp_trace 等效于以前的 **inetd** 的 **-t** 选项（以及 **/etc/default/inetd** 属性 **ENABLE_CONNECTION_LOGGING**）。

tcp_wrappers

如果为 **true**，则启用 TCP 包装器访问控件。这仅适用于 **endpoint_type** 设置为 **streams** 且 **wait** 设置为 **false** 的服务。**syslog** 工具 **code** 守护进程用于记录允许的连接（使用 **notice** 严重级别）和拒绝的通信（使用 **warning** 严重级别）。有关 **syslog** 代码和严重级别的说明，请参见 **syslog.conf(4)**。TCP 包装器工具的接口稳定性及其配置文件容易发生变化。由于 TCP 包装器工具不受 Oracle 控制，因此版本内部经常发生不兼容情况。请参见 **attributes(5)**。

有关配置 TCP 包装器的更多信息，请参见 **tcpd(1M)** 和 **hosts_access(4)** 手册页，这些手册页作为 Solaris 操作系统的一部分在 **/usr/sfw/man** 中提供。这些手册页不包含在标准 Solaris 手册页中，可在 **/usr/man** 上获取。

tcp_wrappers 等效于以前的 **inetd** 的 **/etc/default/inetd** 属性 **ENABLE_TCPWRAPPERS**。

wait

如果为 **true**，则为 **wait** 类型的服务，否则为 **nowait** 类型的服务。**wait** 类型的服务具有下列特征：

- 其 **inetd_start** 方法将在执行时接管服务的绑定端点上的侦听职责。
- 执行此方法后，**inetd** 将等待此方法退出，然后才会恢复侦听职责。

数据报服务器必须配置为类型 **wait**，因为它们始终通过原始数据报端点进行调用，此端点将参与绑定到指定服务的服务的传送。它们没有单独的“侦听”和“接受”套接字。面向连接的服务（如 TCP 流服务）可设计为类型 **wait** 或 **nowait**。

服务的许多基本属性都是可选的。如果缺少这些属性，则会从 **inetd** 服务中的 **defaults** 属性组内提供的缺省值集获取这些属性的值。下面列出了这些属性及其种子值。请注意，这些值可通过 **inetadm(1M)** 进行配置。

```
bind_fail_interval  -1
bind_fail_max       -1
con_rate_offline    -1
connection_backlog  10
failrate_count      40
failrate_time       60
inherit_env         true
max_con_rate        -1
max_copies          -1
tcp_keepalive       false
```



```
tcp_trace      false
tcp_wrappers   false
```

为服务指定的每个方法的配置都存储在 SMF 系统信息库中与相应方法同名的属性组内。这些方法允许的属性集包括为 [svc.startd\(1M\)](#) 管理的服务指定的属性集。（有关更多详细信息，请参见 [svc.startd\(1M\)](#)。）此外，对于 `inetd_start` 方法，可以设置 `arg0` 属性。

`arg0` 属性允许与 `inetd` 服务一起使用外部包装器程序。具体而言，它允许将服务的启动方法的第一个参数 `argv[0]` 设置为除服务器程序路径以外的内容。

如果希望使用外部包装器程序，并将参数传递给服务的守护进程，应将这些参数作为包装器程序的参数合并到 `exec` 属性中。例如：

```
exec='/path/to/wrapper/prog service_daemon_args'
arg0='/path/to/service/daemon'
```

除 [smf_method\(5\)](#) 中提到的特殊方法标记以外，`inetd` 还对 `wait` 类型的服务支持 `:kill_process` 标记。这会导致与提供了 `:kill` 标记的行为相同的行为，但仅向 `wait` 类型的服务的 `start` 方法的父进程发送 `kill` 信号，而不是向其包含的进程合同的所有成员发送此信号（请参见 [process\(4\)](#)）。

为基于套接字的服务配置协议

当为基于套接字的服务配置 `inetd` 时，根据此服务支持的内容，可以选择在上面的 `proto` 属性下介绍的备用项。下面是要使用的 `proto` 值的准则：

- 对于仅支持 IPv4 的服务：则为 `tcp` 和 `udp`
- 对于仅支持 IPv6 的服务：则为 `tcp6only` 和 `udp6only`
- 对于同时支持 IPv4 和 IPv6 的服务：
 - 已过时且不推荐使用的值：`tcp6` 和 `udp6`
 - 建议值：使用两个仅在原始字段存在差异的单独项。其中一个项具有 `tcp` 值，另一个项具有 `tcp6only` 值；或者其中一个项具有 `udp`，另一个项具有 `udp6only` 值。

有关同时支持 IPv4 和 IPv6 的服务的配置示例，请参见“示例”。

inetd 方法

`inetd` 提供下面列出的方法，以供主重启程序 [svc.startd\(1M\)](#) 使用。

start

导致 `inetd` 开始提供服务。这会导致 `inetd` 开始处理其受管服务的 `smf` 请求以及处于 `online` 或 `degraded` 状态的服务的网络请求。

此外，`inetd` 还会检查自上次执行 [inetconv\(1M\)](#) 转换以来其监视的 [inetd.conf\(4\)](#) 格式的配置文件是否已发生更改。如果已发生更改，则会在 `syslog` 中记录一条消息，通知管理员重新运行 `inetconv` 以便实现所做的更改。

stop

导致 `inetd` 停止提供服务。此时，`inetd` 将未处于 `maintenance` 或 `disabled` 状态的各个服务转换为 `offline` 状态，同时在此过程中运行任何适当的方法。

refresh
导致针对每个受管服务执行刷新，并检查 `inetd.conf(4)` 格式配置文件中的更改（如在 `start` 方法中一样）。刷新服务时，其行为取决于服务的当前状态：

- 如果服务处于 `maintenance` 或 `disabled` 状态，则不会执行任何操作，因为将在此服务退出相应状态时读取和使用配置；
- 如果服务处于 `offline` 状态，则会读取配置并立即使用任何更改。
- 如果服务处于 `online` 或 `degraded` 状态，并且配置已发生更改且必须执行重新绑定才能与配置相符，服务将使用绑定的新配置转换到 `offline` 状态并重新联机；
- 如果服务处于 `online` 状态，并且不需要重新绑定，则将运行此服务的 `inetd_refresh` 方法（如果提供），以便允许 `onlinewait` 类型的服务使用任何其他更改。

选项 不支持任何选项。

操作数 *configuration-file*
为旧服务文件 (`inetd.conf(4)`) 指定一个备用位置。

start|stop|refresh
指定应当运行的 `inetd` 方法。

示例 **示例 1** 配置同时支持 IPv4 和 IPv6 的服务
下面的命令演示同时支持 IPv4 和 IPv6 的现有服务以及如何向这些服务分配 `proto` 属性。

```
example# svcs -a | grep mysvc
online      15:48:29 svc:/network/mysvc:dgram4
online      15:48:29 svc:/network/mysvc:dgram6
online      15:51:47 svc:/network/mysvc:stream4
online      15:52:10 svc:/network/mysvc:stream6

# inetadm -M network/rpc/mysvc:dgram4 proto=udp
# inetadm -M network/rpc/mysvc:dgram6 proto=udp6only
# inetadm -M network/rpc/mysvc:stream4 proto=tcp
# inetadm -M network/rpc/mysvc:stream6 proto=tcp6only
```

有关这些命令的说明，请参见 `svcs(1)` 和 `inetadm(1M)`。

属性 有关以下属性的说明，请参见 `attributes(5)`：

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）

另请参见

[fmd\(1M\)](#)、[inetadm\(1M\)](#)、[inetconv\(1M\)](#)、[svcadm\(1M\)](#)、[svccfg\(1M\)](#)、[svcs\(1\)](#)、[svc.startd\(1M\)](#)

附注

inetd 守护进程与 Solaris 9 及早期 Solaris 操作系统版本中的同名守护进程执行相同的功能，但两者的实现方式有很大区别。在当前的 Solaris 发行版中，inetd 是 Solaris 管理工具（请参见 [smf\(5\)](#)）的一部分，并且将仅在此工具中运行。

/etc/default/inetd 文件已过时。属性 `ENABLE_CONNECTION_LOGGING` 和 `ENABLE_TCP_WRAPPERS` 表示的功能现在分别作为 `tcp_trace` 和 `tcp_wrappers` 属性使用。上文中的“服务属性”下介绍了这些属性。

引用名	in.fingerd, fingerd – remote user information server
用法概要	/usr/sbin/in.fingerd [-s]
描述	fingerd implements the server side of the Name/Finger protocol, specified in RFC 742. The Name/Finger protocol provides a remote interface to programs which display information on system status and individual users. The protocol imposes little structure on the format of the exchange between client and server. The client provides a single command line to the finger server which returns a printable reply. fingerd waits for connections on TCP port 79. Once connected, it reads a single command line terminated by RETURN-LINEFEED and passes the arguments to finger(1), prepended with -s. fingerd closes its connections as soon as the output is finished.
选项	fingerd supports the following option. -s Enable secure mode. Deny forwarding of queries to other remote hosts.
文件	/var/adm/utmpx User and accounting information. /etc/passwd System password file. /var/adm/lastlog Last login times. \$HOME/.plan User's plans. \$HOME/.project User's projects.
用法	fingerd and in.fingerd are IPv6-enabled. See ip6(7P) .
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-servers

另请参见 [finger\(1\)](#), [svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [svcadm\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#), [ip6\(7P\)](#)

Harrenstien, Ken, RFC 742, NAME/FINGER, Network Information Center, SRI International, Menlo Park, Calif., December 1977.

附注 Connecting directly to the server from a TIP or an equally narrow-minded TELNET-protocol user program can result in meaningless attempts at option negotiation being sent to the server, which foul up the command line interpretation. fingerd should be taught to filter out IAC's and perhaps even respond negatively (IAC does not) to all option commands received.

The in.fingerd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/network/finger:default

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名	infocmp – compare or print out terminfo descriptions
用法概要	<pre> /usr/bin/infocmp [-d] [-c] [-n] [-I] [-L] [-C] [-r] [-u] [-s d i l c] [-v] [-V] [-1] [-w width] [-A directory] [-B directory] [termname]...</pre>
描述	infocmp compares a binary terminfo entry with other terminfo entries, rewrites a terminfo description to take advantage of the use= terminfo field, or prints out a terminfo description from the binary file (term) in a variety of formats. It displays boolean fields first, then numeric fields, followed by the string fields. If no options are specified and zero, or one <i>termname</i> is specified, the -I option is assumed. If more than one <i>termname</i> is specified, the -d option is assumed.
选项	The -d , -c , and -n options can be used for comparisons. infocmp compares the terminfo description of the first terminal <i>termname</i> with each of the descriptions given by the entries for the other terminal's <i>termname</i>. If a capability is defined for only one of the terminals, the value returned will depend on the type of the capability: F for boolean variables, -1 for integer variables, and NULL for string variables. -d Produce a list of each capability that is different between two entries. This option is useful to show the difference between two entries, created by different people, for the same or similar terminals. -c Produce a list of each capability that is common between two entries. Capabilities that are not set are ignored. This option can be used as a quick check to see if the -u option is worth using. -n Produce a list of each capability that is in neither entry. If no <i>termname</i> is given, the environment variable TERM will be used for both of the <i>termnames</i>. This can be used as a quick check to see if anything was left out of a description. The -I , -L , and -C options will produce a source listing for each terminal named. -I Use the terminfo names. -L Use the long C variable name listed in < term.h >. -C Use the termcap names. The source produced by the -C option may be used directly as a termcap entry, but not all of the parameterized strings may be changed to the termcap format. infocmp will attempt to convert most of the parameterized information, but anything not converted will be plainly marked in the output and commented out. These should be edited by hand. -r When using -C , put out all capabilities in termcap form. If no <i>termname</i> is given, the environment variable TERM will be used for the terminal name. All padding information for strings will be collected together and placed at the beginning of the string where termcap expects it. Mandatory padding (padding information with a trailing '/') will become optional.

All termcap variables no longer supported by terminfo, but are derivable from other terminfo variables, will be displayed. Not all terminfo capabilities will be translated; only those variables which were part of termcap will normally be displayed. Specifying the `-r` option will take off this restriction, allowing all capabilities to be displayed in termcap form.

Note that because padding is collected to the beginning of the capability, not all capabilities are displayed. Mandatory padding is not supported. Because termcap strings are not as flexible, it is not always possible to convert a terminfo string capability into an equivalent termcap format. A subsequent conversion of the termcap file back into terminfo format will not necessarily reproduce the original terminfo source.

Some common terminfo parameter sequences, their termcap equivalents, and some terminal types which commonly have such sequences, are:

terminfo	termcap	Representative Terminals
%p1%c	%. adm	
%p1%d	%d hp, ANSI standard, vt100	
%p1%'x'%'%+%c	%+x concept	
%i	%i ANSI standard, vt100	
%p1?%'x'%'%>%t%p1%'y'%'%+%;	%>xy concept	
%p2 is printed before %p1	%r hp	

`-u` Produce a terminfo source description of the first terminal *termname* which is relative to the sum of the descriptions given by the entries for the other terminals' *termnames*. It does this by analyzing the differences between the first *termname* and the other *termnames* and producing a description with `use=` fields for the other terminals. In this manner, it is possible to retrofit generic terminfo entries into a terminal's description. Or, if two similar terminals exist, but were coded at different times, or by different people so that each description is a full description, using `infocmp` will show what can be done to change one description to be relative to the other.

A capability is displayed with an at-sign (@) if it no longer exists in the first *termname*, but one of the other *termname* entries contains a value for it. A capability's value is displayed if the value in the first *termname* is not found in any of the other *termname* entries, or if the first of the other *termname* entries that has this capability gives a different value for that capability.

The order of the other *termname* entries is significant. Since the terminfo compiler `tic` does a left-to-right scan of the capabilities, specifying two `use=` entries that contain differing entries for the same capabilities will produce different results, depending on the order in which the entries are given. `infocmp` will flag any such inconsistencies between the other *termname* entries as they are found.

Alternatively, specifying a capability *after* a `use=` entry that contains, it will cause the second specification to be ignored. Using `infocmp` to recreate a description can be a useful check to make sure that everything was specified correctly in the original source description.

Another error that does not cause incorrect compiled files, but will slow down the compilation time, is specifying superfluous use= fields. `infocmp` will flag any superfluous use= fields.

- s Sorts the fields within each type according to the argument below:
 - d Leave fields in the order that they are stored in the terminfo database.
 - i Sort by terminfo name.
 - l Sort by the long C variable name.
 - c Sort by the termcap name.
- If the -s option is not given, the fields are sorted alphabetically by the terminfo name within each type, except in the case of the -C or the -L options, which cause the sorting to be done by the termcap name or the long C variable name, respectively.
- v Print out tracing information on standard error as the program runs.
- V Print out the version of the program in use on standard error and exit.
- l Print the fields one to a line. Otherwise, the fields are printed several to a line to a maximum width of 60 characters.
- wwidth Changes the output to *width* characters.

The location of the compiled terminfo database is taken from the environment variable `TERMINFO`. If the variable is not defined, or the terminal is not found in that location, the system terminfo database, usually in `/usr/share/lib/terminfo`, is used. The options -A and -B may be used to override this location.

- A *directory* Set `TERMINFO` for the first *termname*.
- B *directory* Set `TERMINFO` for the other *termnames*. With this, it is possible to compare descriptions for a terminal with the same name located in two different databases. This is useful for comparing descriptions for the same terminal created by different people.

文件 `/usr/share/lib/terminfo/??/*` Compiled terminal description database.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [captainfo\(1M\)](#), [tic\(1M\)](#), [curses\(3CURSES\)](#), [terminfo\(4\)](#), [attributes\(5\)](#)

引用名	in.iked – daemon for the Internet Key Exchange (IKE)
用法概要	<pre>/usr/lib/inet/in.iked [-d] [-f filename] [-p level] /usr/lib/inet/in.iked -c [-f filename]</pre>
描述	in.iked performs automated key management for IPsec using the Internet Key Exchange (IKE) protocol. in.iked implements the following: ■ IKE authentication with either pre-shared keys, DSS signatures, RSA signatures, or RSA encryption. ■ Diffie-Hellman key derivation using either 768, 1024, 1536, 2048, 3072, or 4096-bit public key moduli, or 256, 384, or 521-bit elliptic curve moduli. ■ Authentication protection with cipher choices of AES, DES, Blowfish, or 3DES, and hash choices of either HMAC-MD5 or HMAC-SHA-1. Encryption in in.iked is limited to the IKE authentication and key exchange. See ipsecesp(7P) for information regarding IPsec protection choices. in.iked is managed by the following smf(5) service: <pre>svc:/network/ipsec/ike</pre> This service is delivered disabled because the configuration file needs to be created before the service can be enabled. See ike.config(4) for the format of this file. See “Service Management Facility” for information on managing the smf(5) service. in.iked listens for incoming IKE requests from the network and for requests for outbound traffic using the PF_KEY socket. See pf_key(7P). in.iked has two support programs that are used for IKE administration and diagnosis: ikeadm(1M) and ikecert(1M). The ikeadm(1M) command can read the <code>/etc/inet/ike/config</code> file as a rule, then pass the configuration information to the running in.iked daemon using a doors interface. <pre>example# ikeadm read rule /etc/inet/ike/config</pre> Refreshing the ike smf(5) service provided to manage the in.iked daemon sends a SIGHUP signal to the in.iked daemon, which will (re)read <code>/etc/inet/ike/config</code> and reload the certificate database. The preceding two commands have the same effect, that is, to update the running IKE daemon with the latest configuration. See “Service Management Facility” for more details on managing the in.iked daemon.

When Trusted Extensions are enabled (see [labeld\(1M\)](#)), `in.iked` can be used in the global zone to negotiate labeled security associations. On labeled systems using `in.iked`, UDP ports 500 and 4500 must be configured as multi-level ports for the global zone in the `tnzonecfg` file (see `tnzonecfg(4)`, part of the Solaris Trusted Extensions Reference Manual). See [ike.config\(4\)](#) for more information on configuring `in.iked` to be label-aware.

Service Management Facility

The IKE daemon (`in.iked`) is managed by the service management facility, [smf\(5\)](#). The following group of services manage the components of IPsec:

```
svc:/network/ipsec/ipsecalg      (See ipsecalg(1M))
svc:/network/ipsec/policy       (See ipsecconf(1M))
svc:/network/ipsec/manual-key   (See ipseckey(1M))
svc:/network/ipsec/ike          (see ike.config(4))
```

The `manual-key` and `ike` services are delivered disabled because the system administrator must create configuration files for each service, as described in the respective man pages listed above.

The correct administrative procedure is to create the configuration file for each service, then enable each service using [svcadm\(1M\)](#).

The `ike` service has a dependency on the `ipsecalg` and `policy` services. These services should be enabled before the `ike` service. Failure to do so results in the `ike` service entering maintenance mode.

If the configuration needs to be changed, edit the configuration file then refresh the service, as follows:

```
example# svcadm refresh ike
```

The following properties are defined for the `ike` service:

config/admin_privilege

Defines the level that [ikeadm\(1M\)](#) invocations can change or observe the running `in.iked`. The acceptable values for this property are the same as those for the `-p` option. See the description of `-p` in `OPTIONS`.

config/config_file

Defines the configuration file to use. The default value is `/etc/inet/ike/config`. See [ike.config\(4\)](#) for the format of this file. This property has the same effect as the `-f` flag. See the description of `-f` in `OPTIONS`.

config/debug_level

Defines the amount of debug output that is written to the `debug_logfile` file, described below. The default value for this is `op` or `operator`. This property controls the recording of information on events such as re-reading the configuration file. Acceptable value for `debug_level` are listed in the [ikeadm\(1M\)](#) man page. The value `all` is equivalent to the `-d` flag. See the description of `-d` in `OPTIONS`.

config/debug_logfile

Defines where debug output should be written. The messages written here are from debug code within in.iked. Startup error messages are recorded by the [smf\(5\)](#) framework and recorded in a service-specific log file. Use any of the following commands to examine the logfile property:

```
example# svcs -l ike
example# svcprop ike
example# svccfg -s ike listprop
```

The values for these log file properties might be different, in which case both files should be inspected for errors.

config/ignore_errors

A boolean value that controls in.iked's behavior should the configuration file have syntax errors. The default value is `false`, which causes in.iked to enter maintenance mode if the configuration is invalid.

Setting this value to `true` causes the IKE service to stay online, but correct operation requires the administrator to configure the running daemon with [ikeadm\(1M\)](#). This option is provided for compatibility with previous releases.

These properties can be modified using [svccfg\(1M\)](#) by users who have been assigned the following authorization:

```
solaris.smf.value.ipsec
```

PKCS#11 token objects can be unlocked or locked by using `ikeadm token login` and `ikeadm token logout`, respectively. Availability of private keying material stored on these PKCS#11 token objects can be observed with: `ikeadm dump certcache`. The following authorizations allow users to log into and out of PKCS#11 token objects:

```
solaris.network.ipsec.ike.token.login
solaris.network.ipsec.ike.token.logout
```

See [auths\(1\)](#), [ikeadm\(1M\)](#), [user_attr\(4\)](#), [rbac\(5\)](#).

The service needs to be refreshed using [svcadm\(1M\)](#) before a new property value is effective. General, non-modifiable properties can be viewed with the [svcprop\(1\)](#) command.

```
# svccfg -s ipsec/ike setprop config/config_file = \
/new/config_file
# svcadm refresh ike
```

Administrative actions on this service, such as enabling, disabling, refreshing, and requesting restart can be performed using [svcadm\(1M\)](#). A user who has been assigned the authorization shown below can perform these actions:

```
solaris.smf.manage.ipsec
```

The service's status can be queried using the [svcs\(1\)](#) command.

The `in.iked` daemon is designed to be run under `smf(5)` management. While the `in.iked` command can be run from the command line, this is discouraged. If the `in.iked` command is to be run from the command line, the `ike smf(5)` service should be disabled first. See `svcadm(1M)`.

选项

The following options are supported:

- c Check the syntax of a configuration file.
- d Use debug mode. The process stays attached to the controlling terminal and produces large amounts of debugging output. This option is deprecated. See “Service Management Facility” for more details.
- f *filename* Use *filename* instead of `/etc/inet/ike/config`. See `ike.config(4)` for the format of this file. This option is deprecated. See “Service Management Facility” for more details.
- p *level* Specify privilege level (*level*). This option sets how much `ikeadm(1M)` invocations can change or observe about the running `in.iked`.

Valid *levels* are:

- 0 Base level
- 1 Access to preshared key info
- 2 Access to keying material

If -p is not specified, *level* defaults to 0.

This option is deprecated. See “Service Management Facility” for more details.

安全

This program has sensitive private keying information in its image. Care should be taken with any core dumps or system dumps of a running `in.iked` daemon, as these files contain sensitive keying information. Use the `coreadm(1M)` command to limit any corefiles produced by `in.iked`.

文件

<code>/etc/inet/ike/config</code>	Default configuration file.
<code>/etc/inet/secret/ike.privatekeys/*</code>	Private keys. A private key <i>must</i> have a matching public-key certificate with the same filename in <code>/etc/inet/ike/publickeys/</code> .
<code>/etc/inet/ike/publickeys/*</code>	Public-key certificates. The names are only important with regard to matching private key names.
<code>/etc/inet/ike/crls/*</code>	Public key certificate revocation lists.
<code>/etc/inet/secret/ike.preshared</code>	IKE pre-shared secrets for Phase I authentication.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [svcs\(1\)](#), [coreadm\(1M\)](#), [ikeadm\(1M\)](#), [ikecert\(1M\)](#), [labeld\(1M\)](#), [svccfg\(1M\)](#), [svcadm\(1M\)](#), [ike.config\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [ipsecesp\(7P\)](#), [pf_key\(7P\)](#)

See [tnzonecfg\(4\)](#), part of the Solaris Trusted Extensions Reference Manual.

Harkins, Dan and Carrel, Dave. RFC 2409, Internet Key Exchange (IKE). Network Working Group. November 1998.

Maughan, Douglas, Schertler, M., Schneider, M., Turner, J. RFC 2408, Internet Security Association and Key Management Protocol (ISAKMP). Network Working Group. November 1998.

Piper, Derrell, RFC 2407, The Internet IP Security Domain of Interpretation for ISAKMP. Network Working Group. November 1998.

Fu, D.; Solinos, J., RFC 4753, ECP Groups for IKE and IKEv2. Network Working Group. January 2007.

Lepinski, M.; Kent, S., RFC 5114, Additional Diffie-Hellman Groups for Use with IETF Standards. Network Working Group. January 2008.

引用名 init – 进程控制初始化

用法概要 /usr/sbin/init [0123456abcQqSs]

描述 init 是缺省的最初用户进程。（正如 [kernel\(1M\)](#) 中所述，引导期间提供给内核的选项可能会导致调用备用原始用户进程）。init 启动服务管理工具的核心组件 [svc.configd\(1M\)](#) 和 [svc.startd\(1M\)](#)，如果这些组件失败则会重新启动它们。如下所述，为了实现向后兼容性，init 还可以根据 /etc/inittab 启动或重新启动一般进程。

下面给出的运行级别和系统引导描述仅为兼容性目的提供，除此以外服务管理工具 [smf\(5\)](#) 已将其废弃。

init 失败 如果 init 由于系统关闭以外的任何原因而退出，它将会以进程-ID 1 重新启动。

定义的运行级别 在任意给定时间，系统都处于八个可能的运行级别。运行级别是一个软件配置，其下只存在一组选定的进程。init 为这些运行级别中的每个级别产生的进程都在 /etc/inittab 中进行定义。init 可以处于八个运行级别 0–6 和 S 或 s（S 和 s 相同）中的一个级别。特权用户运行 /usr/sbin/init 时，运行级别会改变。

init 和系统引导 引导系统时，会调用 init 并进行以下操作。首先，它会读取 `svc:/system/environment:init` 服务的属性。这些属性中包含与语言环境相关的环境的值，例如 LANG 或 LC_CTYPE。init 随后会在 /etc/inittab 中查找 initdefault 项（请参见 [inittab\(4\)](#)）。如果 initdefault 项：

存在

 仅当尚未为 [svc.startd\(1M\)](#) 指定选项/里程碑属性时，init 通常会使用该项中指定的运行级别作为要进入的初始运行级别相同。

不存在

 服务管理工具 [smf\(5\)](#) 检查其在 [svc.startd\(1M\)](#) 中的指定配置，然后进入选项/里程碑属性指定的里程碑。

/etc/inittab 中的 initdefault 项对应于以下运行级别：

S 或 s

 init 进入单用户状态。在该状态下，将打开系统控制台设备 (/dev/console) 以读取和写入，并调用命令 /usr/sbin/su（请参见 [su\(1M\)](#)）。使用 init 更改系统的运行级别。请注意，如果终止 shell（使用文件结束符），则 /etc/inittab 不存在时，init 仅重新初始化到单用户状态。

0–6

 init 进入相应的运行级别。运行级别 0、5 和 6 是为关闭系统而保留的状态。运行级别 2、3 和 4 都可用作多用户操作状态。

如果这是开机后 init 第一次进入了除单用户状态之外的运行级别，init 会首先扫描 /etc/inittab 以查找 boot 和 bootwait 项（请参见 [inittab\(4\)](#)）。如果进入的运行级别与这些项的运行级别匹配，会在对 /etc/inittab 进行任何其他处理之前执行这些项。这样，可以在系统允许用户之前进行操作系统的任何特殊初始化，如挂载文件系统。然后，init 扫描 /etc/inittab，并执行该运行级别要处理的所有其他项。

要在 `/etc/inittab` 中产生每个进程，`init` 读取每个项，对于应该重新产生的每个项，它会派生一个子进程。派生了 `/etc/inittab` 指定的所有进程后，`init` 等会待其一个后代进程的终止、`powerfail` 信号或来自其他 `init` 进程的要求更改系统运行级别的信号。出现以上一个条件时，`init` 会重新检查 `/etc/inittab`。

添加 inittab 项

可随时向 `/etc/inittab` 添加新项，但是，`init` 仍会等待出现以上三个条件之一后才重新检查 `/etc/inittab`。要解决此问题，`init Q` 或 `init q` 命令可以唤醒 `init` 以立即重新检查 `/etc/inittab`。

`init` 在引导时间工作时，只要系统从单用户状态改为其他运行状态，`init` 都会将控制台的 `ioctl(2)` 状态设置为文件 `/etc/ioctl.syscon` 中保存的那些模式。进入单用户状态时，`init` 都会向此文件写入。

运行级别更改

提出运行级别更改请求时，`init` 或某个指令会向未在目标运行级别中定义的所有进程发送警告信号 (`SIGTERM`)。在 `init` 或其指令通过发送中止信号 (`SIGKILL`) 来强行终止这些进程之前，会观察一个五秒钟的最小时间间隔。此外，`init` 还会通知 `svc.startd(1M)`，说明将更改运行级别。`svc.startd(1M)` 随后禁止系统使用与运行级别更改相对应的里程碑所依赖的服务集合。

`init` 收到信号告知其产生的进程已经终止时，它会在 `/var/adm/utmpx` 和 `/var/adm/wtmpx` 中记录其终止的事实和原因，如果有的话（请参见 `who(1)`）。产生的进程的历史保存在 `/var/adm/wtmpx` 中。

如果 `init` 收到 `powerfail` 信号 (`SIGPWR`)，它会扫描 `/etc/inittab`，查找 `powerfail` 和 `powerwait` 类型的特殊项。将在进一步处理发生之前调用这些项（如果运行级别允许）。这样，`init` 可以在操作系统电源关闭期间执行各种清理和记录功能。

设置环境变量

对于 `svc:/system/environment:init` 服务的属性列表中的环境变量（例如时区和字符格式化等项），您可以为其设置缺省值。

安全

`init` 使用 `pam(3PAM)` 来进行会话管理。通过 `/etc/pam.conf` 列出的 PAM 配置策略可以指定要用于 `init` 的会话管理模块。下面是部分 `pam.conf` 文件，包含针对使用 UNIX 会话管理模块的 `init` 的项。

```
init    session    required    pam_unix_session.so.1
```

如果没有用于 `init` 服务的项，将使用用于“其他”服务的项。

选项

0

进入固件。

1

将系统置于系统管理员模式。挂载所有本地文件系统。只剩下一小部分必需的内核进程在运行。该模式用于安装可选实用程序包等管理任务。所有文件均可访问，没有用户在系统上登录。

该请求对应于 `smf(5)` 将系统里程碑限制为 `svc:/milestone/single-user:default` 的请求。

2

将系统置于多用户模式。将产生所有多用户环境终端进程和守护进程。该状态通常称为多用户状态。

该请求对应于 **smf(5)** 将系统里程碑限制为 `svc:/milestone/multi-user:default` 的请求。

3

通过使本地资源可在网络中使用来扩展多用户模式。

该请求对应于 **smf(5)** 将系统里程碑限制为 `svc:/milestone/multi-user-server:default` 的请求。

4

可以定义为备用多用户环境配置。它对于系统操作不是必需的，通常不使用。

5

关闭机器，从而安全拔下电源。如果可能，拔下机器上的电源。

6

停止操作系统，重新引导至 `/etc/inittab` 中的 `initdefault` 项定义的状态。

缺省情况下，启用服务

`svc:/system/boot-config:default.config/fastreboot_default` 属性设置为 `true` 时，`init 6` 会根据系统的具体功能，跳过某些固件初始化和测试步骤。

a,b,c

只处理设置了 `a`、`b` 或 `c` 运行级别的那些 `/etc/inittab` 项。这些是伪状态，可以定义它们来运行某些命令，但是不会改变当前的运行级别。

Q,q

重新检查 `/etc/inittab`。

S,s

进入单用户模式。这是唯一不需要存在正确格式化的 `/etc/inittab` 文件的运行级别。如果该文件不存在，缺省情况下，`init` 可以进入的唯一合法运行级别是单用户模式。在单用户模式下时，将会挂载基本系统操作所需的文件系统。系统降为单用户模式时，这些文件系统会保持挂载（即使是由远程文件服务器提供），且任何其他文件系统也会保留挂载。向下转换为单用户模式过程中，`init` 或 `init.d` 脚本启动的、应只在多用户模式下运行的所有进程都会被中止。此外，还将中止具有 `utmpx` 项的任何进程。此最后一个条件确保将中止 `SAC` 启动的所有端口监控器，并且将中止这些端口监控器启动的所有服务，包括 `ttymon` 登录服务。

该请求对应于 **smf(5)** 将系统里程碑限制为 `svc:/milestone/single-user:default` 的请求。

文件

`/dev/console`

系统控制台设备。

`/etc/default/init`

该文件已过时，在将来的发行版中可能会被删除。`init` 进程会读取

`svc:/system/environment:init` 服务的属性，而不是从该文件获取值。为了对之前

通过编辑 `/etc/default/init` 所做的更改进行更改，具有系统管理员或系统配置权限配置文件的管理员可以设置 `init` 服务实例的相应属性并刷新该实例。

该只读文件包含环境变量及其缺省值。这些变量是：

TZ

始终设置为 `localtime`。要设置系统时区，管理员必须设置 `timezone:default` SMF 服务中的 `timezone/localtime` 属性。

CMASK

`init` 使用的以及每个进程从 `init` 进程继承的掩码（请参见 [umask\(1\)](#)）。如果没有设置，`init` 使用从内核继承的掩码。请注意，`init` 在创建文件前始终尝试使用 `022 umask`，而不管 `CMASK` 的设置

LC_CTYPE

字符特征化信息

LC_MESSAGES

消息转换

LC_MONETARY

货币格式化信息

LC_NUMERIC

数字格式化信息

LC_TIME

时间格式化信息

LC_ALL

如果设置，所有其他 `LC_*` 环境变量都会采用此值。

LANG

如果未设置 `LC_ALL`，也未设置任何特定 `LC_*`，`LANG` 值会用于该特定环境变量。

`/etc/inittab`

通过 `init` 控制进程分发。

`/etc/ioctl.syscon`

控制台的 `ioctl` 状态，由 `init` 保存。

`/etc/svc/volatile/init.state`

从故障恢复所需的 `init` 状态。

`/var/adm/utmpx`

用户访问和管理信息。

`/var/adm/wtmpx`

有关用户访问和管理信息的历史记录。

`/var/run/initpipe`

用于内部通信的命名管道。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见 [login\(1\)](#)、[sh\(1\)](#)、[stty\(1\)](#)、[who\(1\)](#)、[kernel\(1M\)](#)、[shutdown\(1M\)](#)、[su\(1M\)](#)、[svc.configd\(1M\)](#)、[svc.](#)

诊断 如果 `init` 发现它在两分钟内从 `/etc/inittab` 重新产生某个项超过十次，它会认定该项中的命令字符串中有错误，并在系统控制台生成错误消息。然后，它拒绝重新产生此项，直到过去了五分钟时间或者它收到来自用户产生的 `init` 命令的信号。这样可以防止 `init` 在以下情况时用光系统资源：有人在 `inittab` 文件中犯了打字错误，或者 `/etc/inittab` 中引用的某个程序已被删除。

附注 `init` 仅可由特权用户运行。

`S` 或 `s` 状态不能在 `/etc/inittab` 中无差别地使用。修改此文件时，最好避免将此状态添加到除 `initdefault` 之外的任何行。

如果未在 `/etc/inittab` 中的 `initdefault` 项中指定缺省状态，会进入状态 `6`。结果，系统会陷入循环，不断地进入固件和重新引导。

如果在引导系统时无法创建 `utmpx` 文件，无论 `/etc/inittab` 中的 `initdefault` 中指定了什么状态，系统都会引导至状态 `s`。如果无法访问 `/var` 文件系统，会发生这种情况。

系统向下转换为 `S` 或 `s` 状态时，会创建 `/etc/nologin` 文件（请参见 [nologin\(4\)](#)）。在后续转换至运行级别 `2` 时，将删除该文件。

`init` 使用 `/var/run/initpipe`（命名管道）进行内部通信。

- 引用名

init.sma – start and stop the snmpd daemon
- 用法概要

/etc/init.d/init.sma start | stop | restart | status
- 描述

The `init.sma` utility is run automatically during installation and each time the system is rebooted. This utility manages the `snmpd`.
- 选项

The following options are supported:

start

Starts the `snmpd` daemon.

stop

Stops the `snmpd` daemon.

restart

Stops then starts the `snmpd` daemon.

status

Reports the `snmpd` daemon's status.
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWsmagt
Interface Stability	Uncommitted

另请参见

[attributes\(5\)](#)

引用名	init.wbem – start and stop the CIM Boot Manager
用法概要	<code>svccfg svc:/application/management/wbem start stop</code>
描述	The <code>init.wbem</code> utility is run automatically during installation, each time the system is rebooted, and by means of the svccfg(1M) command. This method manipulates the CIM Object Manager (CIMOM). <code>init.wbem</code> can be used to start, stop, or retrieve status from the server.
CIM Object Manager	The CIM Object Manager manages CIM objects on a WBEM-enabled system. A CIM object is a computer representation, or model, of a managed resource, such as a printer, disk drive, or CPU. CIM objects are stored internally as Java classes. When a WBEM client application accesses information about a CIM object, the CIM Object Manager contacts either the appropriate provider for that object or the CIM Object Manager Repository. Providers are classes that communicate with managed objects to access data. When a WBEM client application requests data from a managed resource that is not available from the CIM Object Manager Repository, the CIM Object Manager forwards the request to the provider for that managed resource. The provider dynamically retrieves the information. At startup, the CIM Object Manager performs the following functions: ▪ Listens for RMI connections on RMI port 5987 and for XML/HTTP connections on HTTP port 5988. ▪ Sets up a connection to the CIM Object Manager Repository. ▪ Waits for incoming requests. During normal operations, the CIM Object Manager performs the following functions: ▪ Performs security checks to authenticate user login and authorization to access namespaces. ▪ Performs syntactical and semantic checking of CIM data operations to ensure that they comply with the latest CIM Specification. ▪ Routes requests to the appropriate provider or to the CIM Object Manager Repository. ▪ Delivers data from providers and from the CIM Object Manager Repository to WBEM client applications. A WBEM client application contacts the CIM Object Manager to establish a connection when it needs to perform WBEM operations, such as creating a CIM class or updating a CIM instance. When a WBEM client application connects to a CIM Object Manager, it gets a reference to the CIM Object Manager, which it then uses to request services and operations.
System Booting	The <code>init.wbem</code> script is installed in the <code>/etc/init.d</code> directory.

选项 The `init.wbem` function is implemented as a service management facility (SMF) method and is controlled by means of the [svcadm\(1M\)](#) command using the fault management resource identifier (FMRI) `svc:/application/management/wbem`.

The following options are supported through [svcadm\(1M\)](#):

`start` Starts the CIMOM on the local host.

`stop` Stops the CIMOM.

`status` Gets the status of the CIMOM.

The SMF property `options/tcp_listen` is used to control whether CIMOM and the Solaris Management console server will respond to requests from remote systems.

The specification:

```
svc:/application/management/wbem/options/tcp_listen = true
```

...allows remote access and `false` disallows remote access. `false` is the default value.

示例 示例 1 Allowing Access to Remote Systems

The following commands enable CIMOM to allow access from remote systems.

```
# svccfg -s \  
svc:/application/management/wbem setprop options/tcp_listen = true  
# svcadm refresh svc:/application/management/wbem
```

附注 When the `init.wbem` script is invoked by SMF, it does not run the CIMOM directly. The server process is in Java and is too heavyweight to be run immediately at system boot time. Instead, three lightweight processes listen on three different ports that the CIMOM normally uses. This acts similarly to [inetd\(1M\)](#).

Because Java programs cannot inherit file descriptors as other programs can, there is a small time period from when the first connection is made until the server is fully operational where client connections may be dropped. WBEM clients are immune to this, as they will retry until the server comes online.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWwbco

另请参见 [inetd\(1M\)](#), [mofcomp\(1M\)](#), [svccfg\(1M\)](#), [svcadm\(1M\)](#), [wbemadmin\(1M\)](#), [wbemlogviewer\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

引用名

用法概要

描述

选项

文件

属性

inityp2l – create NIS (YP) to LDAP configuration files

```
/usr/lib/netsvc/yp/inityp2l [-m mapping_file_name]
                             [-c config_file_name]
```

The `inityp2l` utility assists with creation of the `NISLDAPmapping` and `ypserv` files. See [NISLDAPmapping\(4\)](#) and [ypserv\(4\)](#). `inityp2l` examines the NIS maps on a system, and through a dialogue with the user, determines which NIS to (and from) LDAP mappings are required. A `NISLDAPmapping` file is then created based on this information. The utility asks users about their LDAP server configuration and a `ypserv` file is created based on this information.

The `inityp2l` utility handles mappings for standard NIS maps and the `auto.*` series of maps. If requested, it creates default mappings for custom maps, with each map entry represented as a single DIT string. `inityp2l` does not handle full custom mapping, but if requested, `inityp2l` will insert comments into the `NISLDAPmapping` file that indicate where these should be added.

To write to the `NISLDAPmapping` or `ypserv` files is potentially dangerous. `inityp2l` warns the user and asks for confirmation before:

1. it overwrites either file
2. it writes to the default `NISLDAPmapping` file location, if this file did not previously exist. This is important because the existence of a file in this location causes NIS components to work NIS to LDAP (N2L) mode when next restarted, rather than to traditional NIS mode.

`inityp2l` assists with rapid creation of a simple N2L configuration files. It is not a general purpose tool for the management of these files. An advanced user who would like to maintain the files or use custom mappings should examine the output of `inityp2l` and customize it by using a standard text editor.

`inityp2l` supports the following options:

- c Specify the name of the generated `ypserv` file. The default location is described in [文件](#).
- m Specify the name of the generated `NISLDAPmapping` file. The default is described in [文件](#).

<code>/var/yp</code>	The directory to be searched for candidate domains (<code>/var/yp/*</code>) and NIS maps (<code>/var/yp/*/*</code>)
<code>/var/yp/NISLDAPmapping</code>	The default location for the generated <code>NISLDAPmapping</code> file
<code>/etc/default/ypserv</code>	The default location for the generated <code>ypserv</code> file

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/nis
Interface Stability	Obsolete

另请参见 [NISLDAPmapping\(4\)](#), [ypserv\(4\)](#), [attributes\(5\)](#)

引用名	in.lpd – BSD print protocol adaptor				
用法概要	/usr/lib/print/in.lpd				
描述	in.lpd implements the network listening service for the BSD print protocol specified in RFC 1179. The BSD print protocol provides a remote interface for systems to interact with a local spooling system. The protocol defines five standard requests from the client to the server: starting queue processing, transferring print jobs, retrieving terse status, retrieving verbose status, and canceling print jobs. The in.lpd service is managed by the service management facility, smf(5), under the service identifier: <pre>svc:/network/lp</pre> Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M). Responsibility for initiating and restarting this service is delegated to inetd(1M). Use inetadm(1M) to make configuration changes and to view configuration information for this service. The service's status can be queried using the svcs(1) command. in.lpd uses the config/log_from_remote property to allow or disallow remote access. The default value of this property, localhost, disallows remote access. inetd waits for connections on TCP port 515. Upon receipt of a connect request, in.lpd is started to service the connection. Once the request has been filled, in.lpd closes the connection and exits.				
示例	示例 1 Allowing Remote Access The following command allows remote access to in.lpd. <pre># inetadm -m svc:/application/print/rfc1179:default bind_addr=""</pre>				
退出状态	The following exit values are returned: <table><tr><td>0</td><td>Successful completion.</td></tr><tr><td>non-zero</td><td>An error occurred.</td></tr></table>	0	Successful completion.	non-zero	An error occurred.
0	Successful completion.				
non-zero	An error occurred.				
文件	/usr/lib/print/bsd-adaptor/bsd_*.so* Spooler translation modules.				
属性	See attributes(5) for descriptions of the following attributes:				

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	print/lp/print-client-commands

另请参见

[svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [svcadm\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

引用名 in.mpathd – IP multipathing daemon

用法概要 /usr/lib/inet/in.mpathd

描述

The `in.mpathd` daemon performs failure and repair detection for IP interfaces that have been placed into an IPMP group (or optionally, for all IP interfaces on the system). It also controls which IP interfaces in an IPMP group are “active” (being used by the system to send or receive IP data traffic) in a manner that is consistent with the administrator's configured policy. Functioning interfaces that are not active are referred to as *alternate* interfaces.

The `in.mpathd` daemon can detect IP interface failure and repair through three methods: by monitoring the `IFF_RUNNING` flag for each IP interface (link-state based failure detection), by sending and receiving ICMP probes on each IP interface (ICMP probe-based failure detection), or by transitive probing. Failure detection by transitive probing is done by sending and receiving ICMP probes on active interfaces, while determining the liveness of alternate interfaces through their ability to send and receive link-layer (that is, transitive) probes to an active interface.

By default, only link-state based failure detection is enabled. This requires the driver to support link-status notification. ICMP probe-based failure detection must be enabled through the configuration of one or more test addresses, which are described below. Transitive probing can be enabled by modifying the value of the `SMF` property (shown below) to `true`:

```
svc:/network/ipmp/config/transitive-probing
```

See `EXAMPLES` for more information on how to modify the value of the `transitive-probing` property.

Both ICMP and transitive probe-based failure detection methods test the entire IP interface send and receive path. The `ipmpstat(1M)` utility can be used to check which failure detection methods are enabled.

If `transitive-probing` is set to `true`, and no test addresses are configured for a given IPMP group, then transitive probing will be used. If it is set to `false` (default value), then transitive probing will not be used under any circumstance.

If only link-state based failure detection is enabled, then the health of the interface is determined solely from the state of the `IFF_RUNNING` flag. If probes have been enabled, the interface is considered failed if either link-state or probes indicate a failure, and repaired once the failure detection method has indicated the failure has been corrected. Although all interfaces in a group need not be configured with the same failure detection methods, transitive probing will be disabled on any interface of a group that has at least one IP test address.

As mentioned above, to perform ICMP probe-based failure detection, `in.mpathd` requires a test address on each IP interface for the purpose of sending and receiving probes. Each address must be marked `NOFAILOVER` (see `ifconfig(1M)`) and `in.mpathd` will be limited to probing targets on the same subnet. Each address may be configured statically or acquired by means of

DHCP. To find targets, `in.mpathd` first consults the routing table for routes on the same subnet, and uses the specified next-hop. If no routes match, it sends all-hosts ICMP probes and selects a subset of the systems that respond. Thus, for probe-based failure detection to operate, there must be at least one neighbor on each subnet that responds to ICMP echo request probes. The `ipmpstat(1M)` utility can be used to display both the current probe target information and the status of sent probes.

Both IPv4 and IPv6 are supported. If an IP interface is plumbed for IPv4 and an IPv4 test address is configured then `in.mpathd` will start sending ICMPv4 probes over that IP interface. Similarly, if an IP interface is plumbed for IPv6 and an IPv6 test address is configured, then `in.mpathd` will start sending ICMPv6 probes over that IP interface. However, note that `in.mpathd` will ignore IPv6 test addresses that are not link-local. If both IPv4 and IPv6 are plumbed, it is sufficient to configure only one of the two, that is, either an IPv4 test address or an IPv6 test address. If both IPv4 and IPv6 test addresses are configured, `in.mpathd` probes using both ICMPv4 and ICMPv6.

As mentioned above, `in.mpathd` also controls which IP interfaces in an IPMP group are “active” (used by the system to send and receive IP data traffic). Specifically, `in.mpathd` tracks the administrative configuration of each IPMP group and attempts to keep the number of active IP interfaces in each group consistent with that configuration. Therefore, if an active IP interface fails, `in.mpathd` will activate an INACTIVE interface in the group, provided one exists (it will prefer INACTIVE interfaces that are also marked STANDBY). Likewise, if an IP interface repairs and the resulting repair leaves the IPMP group with more active interfaces than the administrative configuration specifies, `in.mpathd` will deactivate one of the interfaces (preferably one marked STANDBY), except when the FAILBACK variable is used, as described below. Similar adjustments will be made by `in.mpathd` when offlining IP interfaces (for instance, in response to `if_mpadm(1M)`).

The `in.mpathd` daemon accesses three variable values in `/etc/default/mpathd`: `FAILURE_DETECTION_TIME`, `FAILBACK` and `TRACK_INTERFACES_ONLY_WITH_GROUPS`.

The `FAILURE_DETECTION_TIME` variable specifies the probe-based failure detection time. The shorter the failure detection time, the more probe traffic. The default value of `FAILURE_DETECTION_TIME` is 10 seconds. This means that IP interface failure will be detected by `in.mpathd` within 10 seconds. The IP interface repair detection time is always twice the value of `FAILURE_DETECTION_TIME`. Note that failures and repairs detected by link-based failure detection are acted on immediately, though `in.mpathd` may ignore link state changes if it suspects that the link state is flapping due to defective hardware; see `DIAGNOSTICS`.

By default, `in.mpathd` limits failure and repair detection to IP interfaces that are configured as part of a named IPMP group. Setting `TRACK_INTERFACES_ONLY_WITH_GROUPS` to `no` enables failure and repair detection on all IP interfaces, even if they are not part of a named IPMP group. IP interfaces that are tracked but not part of a named IPMP group are considered to be part of the “anonymous” IPMP group. In addition to having no name, this IPMP group is special in that its IP interfaces are not equivalent and thus cannot take over for one another in

the event of an IP interface failure. That is, the anonymous IPMP group can only be used for failure and repair detection, and provides no high-availability or load-spreading.

As described above, when `in.mpathd` detects that an IP interface has repaired, it activates it so that it will again be used to send and receive IP data traffic. However, if `FAILBACK` is set to `no`, then the IP interface will only be activated if no other active IP interfaces in the group remain. However, the interface may subsequently be activated if another IP interface in the group fails.

SMF Management The `in.mpathd` daemon service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/ipmp:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#).

示例 示例 1 Enabling Fault Detection by Transitive Probing

The following example shows the sequence of SMF commands used to enable fault detection by transitive probing.

```
# svccfg -s svc:/network/ipmp setprop config/transitive-probing=true
# svcadm refresh svc:/network/ipmp:default
```

文件 `/etc/default/mpathd` Contains default values used by the `in.mpathd` daemon.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [if_mpadm\(1M\)](#), [ifconfig\(1M\)](#), [ipmpstat\(1M\)](#), [svcadm\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#), [icmp\(7P\)](#), [icmp6\(7P\)](#),

《Oracle Solaris 管理：IP 服务》

诊断 IP interface *interface_name* has a hardware address which is not unique in group *group_name*; offlining

描述: For probe-based failure detection, load-spreading, and other code IPMP features to work properly, each IP interface in an IPMP group must have a unique hardware address. If this requirement is not met, `in.mpathd` will automatically offline all but one of the IP interfaces with duplicate hardware addresses.

IP interface *interface_name* now has a unique hardware address in group *group_name*; onlining

描述: The previously-detected duplicate hardware address is now unique, and therefore `in.mpathd` has brought *interface_name* back online.

Test address *address* is not unique in group; disabling probe-based failure detection on *interface_name*

描述: For in.mpathd to perform probe-based failure detection, each test address in the group must be unique.

No test address configured on interface *interface_name* disabling probe-based failure detection on it

描述: For in.mpathd to perform probe-based failure detection on an IP interface, it must be configured with a test address: IPv4, IPv6, or both.

IP *interface_name* in group *group_name* is not plumbed for IPv[4|6], affecting IPv[4|6] connectivity

描述: All IP interfaces in a multipathing group must be homogeneously plumbed. For example, if one IP interface is plumbed for IPv4, then all IP interfaces in the group must be plumbed for IPv4, or IPv4 packets will not be able to be reliably sent and received. The STREAMS modules pushed on all IP interfaces must also be identical.

The link has come up on *interface_name* more than 2 times in the last minute; disabling repair until it stabilizes.

描述: To limit the impact of interfaces with intermittent hardware (such as a bad cable), in.mpathd will not consider an IP interface with a frequently changing link state as repaired until the link state stabilizes.

Invalid failure detection time of *time*, assuming default 10000 ms

描述: An invalid value was encountered for FAILURE_DETECTION_TIME in the /etc/default/mpathd file.

Too small failure detection time of *time*, assuming minimum of 100 ms

描述: The minimum value that can be specified for FAILURE_DETECTION_TIME is currently 100 milliseconds.

Invalid value for FAILBACK *value*

描述: Valid values for the boolean variable FAILBACK are yes or no.

Invalid value for TRACK_INTERFACES_ONLY_WITH_GROUPS *value*

描述: Valid values for the boolean variable TRACK_INTERFACES_ONLY_WITH_GROUPS are yes or no.

Cannot meet requested failure detection time of *time* ms on (inet[6] *interface_name*) new failure detection time for group *group_name* is *time* ms

描述: The round trip time for ICMP probes is higher than necessary to maintain the current failure detection time. The network is probably congested or the probe targets are loaded. in.mpathd automatically increases the failure detection time to whatever it can achieve under these conditions.

Improved failure detection time *time* ms on (inet[6] *interface_name*) for group *group_name*

描述: The round trip time for ICMP probes has now decreased and in.mpathd has lowered the failure detection time correspondingly.

IP interface failure detected on *interface_name*

描述: in.mpathd has detected a failure on *interface_name*, and has set the IFF_FAILED flag on *interface_name*, ensuring that it will not be used for IP data traffic.

IP interface repair detected on *interface_name*

描述: in.mpathd has detected a repair on *interface_name*, and has cleared the IFF_FAILED flag. Depending on the administrative configuration, the *interface_name* may again be used for IP data traffic.

All IP interfaces in group *group* are now unusable

描述: in.mpathd has determined that none of the IP interfaces in *group* can be used for IP data traffic, breaking network connectivity for the group.

At least 1 IP interface (*interface_name*) in group *group* is now usable

描述: in.mpathd has determined that at least one of the IP interfaces in *group* can again be used for IP data traffic, restoring network connectivity for the group.

The link has gone down on *interface_name*

描述: in.mpathd has detected that the IFF_RUNNING flag for *interface_name* has been cleared, indicating that the link has gone down.

The link has come up on *interface_name*

描述: in.mpathd has detected that the IFF_RUNNING flag for *interface_name* has been set, indicating that the link has come up.

引用名	in.ndpd – daemon for IPv6 autoconfiguration
用法概要	<code>/usr/lib/inet/in.ndpd [-adt] [-f <i>config_file</i>]</code>
描述	in.ndpd provides both the host and router autoconfiguration components of Neighbor Discovery for IPv6 and Stateless and Stateful Address Autoconfiguration for IPv6. In particular, in.ndpd implements:

- router discovery;
- prefix discovery;
- parameter discovery;
- invocation of stateful address autoconfiguration;
- stateless address autoconfiguration; and
- privacy extensions for stateless address autoconfiguration.

Other aspects of Neighbor Discovery are implemented by [ip6\(7P\)](#), including:

- address resolution;
- neighbor unreachability detection; and
- redirect.

The duplicate address detection function is implemented by the system kernel.

in.ndpd is managed by the service management facility (SMF), by means of the service identifier:

```
svc:/network/routing/ndp:default
```

If the `/etc/inet/ndpd.conf` file does not exist or does not set the variable `AdvSendAdvertisements` to true for a network interface, then in.ndpd will make the node a host for that interface, that is, sending router solicitation messages and then using router advertisement messages it receives to autoconfigure the node. Note that in.ndpd only autoconfigures the addresses of global or site-local scope from the prefix advertisement.

If `AdvSendAdvertisements` is set to true for an interface, then in.ndpd will perform router functions on that interface, that is, sending router advertisement messages to autoconfigure the attached hosts, but not use any advertisements it receives for autoconfiguration. However, when sending advertisements, in.ndpd will use the advertisements it sends itself to autoconfigure its prefixes.

Stateless autoconfiguration requires no manual configuration of hosts, minimal (if any) configuration of routers, and no additional servers. The stateless mechanism enables a host to generate its own addresses and uses local information as well as non-local information that is advertised by routers to generate the addresses. in.ndpd will plumb logical interfaces for each of these addresses.

Stateful autoconfiguration involves the [dhcpgent\(1M\)](#) daemon and the use of the DHCPv6 protocol. The dhcpgent daemon is responsible for plumbing the logical interfaces for the acquired addresses, maintaining the leases, and handling duplicate addresses. in.ndpd starts

the `dhcpgent` daemon automatically and signals when DHCPv6 should be started. `in.ndpd` also detects when `dhcpgent` configures the logical interfaces, and sets the appropriate prefix length on each according to received Routing Advertisement messages. `in.ndpd` will not stop `dhcpgent`; use [ifconfig\(1M\)](#) to control `dhcpgent` if necessary.

Temporary addresses that are autoconfigured for an interface can also be implemented. A temporary address token is enabled for one or more interfaces on a host. However, unlike standard, autoconfigured IPv6 addresses, a temporary address consists of the site prefix and a randomly generated 64 bit number. This random number becomes the interface ID segment of the IPv6 address. A link-local address is not generated with the temporary address as the interface ID.

If the kernel detects a duplicate temporary address, `in.ndpd` will automatically choose another.

Routers advertise all prefixes that have been assigned on the link. IPv6 hosts use Neighbor Discovery to obtain a subnet prefix from a local router. Hosts automatically create IPv6 addresses by combining the subnet prefix with an interface IDs that is generated from an interface's MAC address. In the absence of routers, a host can generate only link-local addresses. Link-local addresses can only be used for communication with nodes on the same link.

For information on how to enable IPv6 address autoconfiguration, see [《Oracle Solaris 管理：IP 服务》](#)

选项

Supported options and equivalent SMF service properties are listed below. SMF service properties are set using a command of the form:

```
# routeadm -m ndp:default key=value
```

-a

Turn off stateless and stateful address auto configuration. When set, the daemon does not autoconfigure any addresses and does not renumber any addresses. This option does the same thing as the following lines in [ndpd.conf\(4\)](#):

```
ifdefault StatefulAddrConf off
ifdefault StatelessAddrConf off
```

Use of this option is equivalent to setting the `stateless_addr_conf` property to false.

-d

Turn on large amounts of debugging output on `stdout`. When set, the program runs in the foreground and stays attached to the controlling terminal. Use of this option is equivalent to setting the `debug` property to true.

-f *config_file*

Use *config_file* for configuration information instead of the default `/etc/inet/ndpd.conf`. Use of this option is equivalent to setting the `config_file` property to the configuration file to be used.

-t
Turn on tracing (printing) of all sent and received packets to stdout. When set, the program runs in the foreground and stays attached to the controlling terminal. As such, this option cannot be run under the SMF.

文件 /etc/inet/ndpd.conf Configuration file. This file is not necessary on a host, but it is required on a router to enable in.ndpd to advertise autoconfiguration information to the hosts.

属性 See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 dhcpagent(1M), ifconfig(1M), routeadm(1M), svcadm(1M), ndpd.conf(4), attributes(5), icmp6(7P), ip6(7P), attributes(5), smf(5)

《Oracle Solaris 管理：IP 服务》

Narten, T., Nordmark, E., Simpson, W. RFC 2461, Neighbor Discovery for IP Version 6 (IPv6). The Internet Society. December 1998.

Thomson, S., Narten, T. RFC 2462, IPv6 Stateless Address Autoconfiguration. The Internet Society. December 1998.

Narten, T., and Draves, R. RFC 3041, Privacy Extensions for Stateless Address Autoconfiguration in IPv6. The Internet Society. January 2001.

诊断 Receipt of a SIGHUP signal will make in.ndpd restart and reread /etc/inet/ndpd.conf.

附注 The in.ndpd daemon service is managed by the service management facility, smf(5), under the service identifier:
svc:/network/routing/ndp:default

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M) or routeadm(1M).

引用名	in.rarpd, rarpd – DARPA Reverse Address Resolution Protocol server
用法概要	<pre>/usr/sbin/in.rarpd [-d] -a /usr/sbin/in.rarpd [-d] device unit</pre>
描述	<code>in.rarpd</code> starts a daemon that responds to Reverse Address Resolution Protocol (RARP) requests. The daemon forks a copy of itself that runs in background. It must be run as root. RARP is used by machines at boot time to discover their Internet Protocol (IP) address. The booting machine provides its Ethernet address in a RARP request message. Using the <code>ethers</code> and <code>hosts</code> databases, <code>in.rarpd</code> maps this Ethernet address into the corresponding IP address which it returns to the booting machine in an RARP reply message. The booting machine must be listed in both databases for <code>in.rarpd</code> to locate its IP address. <code>in.rarpd</code> issues no reply when it fails to locate an IP address. <code>in.rarpd</code> uses the STREAMS-based Data Link Provider Interface (DLPI) message set to communicate directly with the datalink device driver.
选项	The following options are supported: -a Get the list of available network interfaces from IP using the <code>SIOCGIFADDR</code> ioctl and start a RARP daemon process on each interface returned. -d Print assorted debugging messages while executing.
示例	示例 1 Starting An <code>in.rarpd</code> Daemon For Each Network Interface Name Returned From <code>/dev/ip</code>: The following command starts an <code>in.rarpd</code> for each network interface name returned from <code>/dev/ip</code>: <pre>example# /usr/sbin/in.rarpd -a</pre> 示例 2 Starting An <code>in.rarpd</code> Daemon On The Device <code>/dev/le</code> With The Device Instance Number 0 The following command starts one <code>in.rarpd</code> on the device <code>/dev/le</code> with the device instance number 0. <pre>example# /usr/sbin/in.rarpd le 0</pre>
文件	<pre>/etc/ethers File or other source, as specified by nsswitch.conf(4). /etc/hosts File or other source, as specified by nsswitch.conf(4). /tftpboot /dev/ip /dev/arp</pre>

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/boot/network

另请参见 [svcs\(1\)](#), [boot\(1M\)](#), [ifconfig\(1M\)](#), [svcadm\(1M\)](#), [ethers\(4\)](#), [hosts\(4\)](#), [netconfig\(4\)](#), [nsswitch.conf\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [dlpi\(7P\)](#)

Finlayson, R., Mann, T., Mogul, J., and Theimer, M., RFC 903, A Reverse Address Resolution Protocol, Network Information Center, SRI International, June 1984.

附注 The `in.rarpd` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/network/rarp`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	in.rdisc, rdisc – network router discovery daemon
用法概要	<pre>/usr/sbin/in.rdisc [-a] [-f] [-s] [send-address] [receive-address] /usr/sbin/in.rdisc -r [-p preference] [-T interval] [send-address] [receive-address]</pre>
描述	in.rdisc remains part of the software distribution of the Solaris Operating Environment. It is, however, not used by default. in.routed(1M) includes the functionality provided by in.rdisc. See routeadm(1M) for details of how to specify the IPV4 routing daemon. in.rdisc implements the ICMP router discovery protocol. The first form of the command is used on hosts and the second form is used on routers. in.rdisc can be invoked in either the first form (host mode) or second form (router mode). On a host, in.rdisc populates the network routing tables with default routes. On a router, advertises the router to all the hosts. in.rdisc is managed by the service management facility (SMF), by means of the service identifier: <pre>svc:/network/routing/rdisc:default</pre>
Host (First Form)	On a host, in.rdisc listens on the ALL_HOSTS (224.0.0.1) multicast address for ROUTER_ADVERTISE messages from routers. The received messages are handled by first ignoring those listed router addresses with which the host does not share a network. Among the remaining addresses, the ones with the highest preference are selected as default routers and a default route is entered in the kernel routing table for each one of them. Optionally, in.rdisc can avoid waiting for routers to announce themselves by sending out a few ROUTER_SOLICITATION messages to the ALL_ROUTERS (224.0.0.2) multicast address when it is started. A timer is associated with each router address. The address will no longer be considered for inclusion in the routing tables if the timer expires before a new <i>advertise</i> message is received from the router. The address will also be excluded from consideration if the host receives an <i>advertise</i> message with the preference being maximally negative or with a lifetime of zero.
Router (Second Form)	When in.rdisc is started on a router, it uses the SIOCGIFCONF ioctl(2) to find the interfaces configured into the system and it starts listening on the ALL_ROUTERS multicast address on all the interfaces that support multicast. It sends out <i>advertise</i> messages to the ALL_HOSTS multicast address advertising all its IP addresses. A few initial <i>advertise</i> messages are sent out during the first 30 seconds and after that it will transmit <i>advertise</i> messages approximately every 600 seconds. When in.rdisc receives a <i>solicitation</i> message, it sends an <i>advertise</i> message to the host that sent the <i>solicitation</i> message.

When `in.rdisc` is terminated by a signal, it sends out an *advertise* message with the preference being maximally negative.

选项

Supported options and equivalent SMF service properties are listed below. SMF service properties are set using a command of the form:

- `# routeadm -m rdisc:default key=value`
- a Accept all routers independent of the preference they have in their *advertise* messages. Normally, `in.rdisc` only accepts (and enters in the kernel routing tables) the router or routers with the highest preference. Use of this option is equivalent to setting the `accept_all` property to true.
 - f Run `in.rdisc` forever even if no routers are found. Normally, `in.rdisc` gives up if it has not received any *advertise* message after soliciting three times, in which case it exits with a non-zero exit code. If `-f` is not specified in the first form then `-s` must be specified. For SMF execution, this option is required.
 - r Act as a router, rather than a host. Use of this option is equivalent to setting the `act_as_router` property to true.
 - s Send three *solicitation* messages initially to quickly discover the routers when the system is booted. When `-s` is specified, `in.rdisc` exits with a non-zero exit code if it can not find any routers. This can be overridden with the `-f` option. This option is not compatible with SMF execution and is not supported for the `rdisc` service.
 - p *preference* Set the preference transmitted in the *solicitation* messages. The default is zero. Use of this option is equivalent to setting the `preference` property.
 - T *interval* Set the interval between transmitting the *advertise* messages. The default time is 600 seconds. Use of this option is equivalent to setting the `transmit_interval` property.

The send-address and receive-address daemon options can be set by means of the `send_address` and `receive_address` properties.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/routing

另请参见

[in.routed\(1M\)](#), [routeadm\(1M\)](#), [svcadm\(1M\)](#), [ioctl\(2\)](#), [gateways\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [icmp\(7P\)](#), [inet\(7P\)](#)

Deering, S.E., editor, ICMP Router Discovery Messages, RFC 1256, Network Information Center, SRI International, Menlo Park, California, September 1991.

引用名 in.rexecd, rexecd – remote execution server

用法概要 in . rexecd

描述 in . rexecd is the server for the [rexec\(3SOCKET\)](#) routine. The server provides remote execution facilities with authentication based on user names and passwords. It is invoked automatically as needed by [inetd\(1M\)](#), and then executes the following protocol:

1. The server reads characters from the socket up to a null (`\0`) byte. The resultant string is interpreted as an ASCII number, base 10.
2. If the number received in step 1 is non-zero, it is interpreted as the port number of a secondary stream to be used for the `stderr`. A second connection is then created to the specified port on the client's machine.
3. A null terminated user name of at most 16 characters is retrieved on the initial socket.
4. A null terminated password of at most 16 characters is retrieved on the initial socket.
5. A null terminated command to be passed to a shell is retrieved on the initial socket. The length of the command is limited by the upper bound on the size of the system's argument list.
6. rexecd then validates the user as is done at login time and, if the authentication was successful, changes to the user's home directory, and establishes the user and group protections of the user. If any of these steps fail the connection is aborted and a diagnostic message is returned.
7. A null byte is returned on the connection associated with the `stderr` and the command line is passed to the normal login shell of the user. The shell inherits the network connections established by rexecd.

用法 in . rexecd and rexecd are IPv6-enabled. See [ip6\(7P\)](#).

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-servers

另请参见 [svcs\(1\)](#), [inetd\(1M\)](#), [inetadm\(1M\)](#), [svcadm\(1M\)](#), [rexec\(3SOCKET\)](#), [attributes\(5\)](#), [smf\(5\)](#), [ip6\(7P\)](#)

诊断 All diagnostic messages are returned on the connection associated with the `stderr`, after which any network connections are closed. An error is indicated by a leading byte with a value of 1 (0 is returned in step 7 above upon successful completion of all the steps prior to the command execution).

username too long	The name is longer than 16 characters.
password too long	The password is longer than 16 characters.

command too long	The command line passed exceeds the size of the argument list (as configured into the system).
Login incorrect.	No password file entry for the user name existed.
Password incorrect.	The wrong password was supplied.
No remote directory.	The chdir command to the home directory failed.
Try again.	A fork by the server failed.
/usr/bin/sh: ...	The user's login shell could not be started.

附注

The `in.rexecd` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/rexec:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名

in.ripngd – network routing daemon for IPv6

用法概要

```
/usr/sbin/in.ripngd [-s] [-q] [-t] [-p n] [-P] [-v ]
[logfile]
```

描述

in.ripngd is the IPv6 equivalent of [in.routed\(1M\)](#). It is invoked at boot time to manage the network routing tables. The routing daemon uses the Routing Information Protocol for IPv6.

in.ripngd is managed by the service management facility (SMF), by means of the service identifier:

```
svc:/network/routing/ripng:default
```

In normal operation, in.ripngd listens on the [udp\(7P\)](#) socket port 521 for routing information packets. If the host is an internetwork router, it periodically supplies copies of its routing tables to any directly connected hosts and networks.

When in.ripngd is started, it uses the `SIOCGLIFCONF` [ioctl\(2\)](#) to find those directly connected IPv6 interfaces configured into the system and marked “up”; the software loopback interface is ignored. If multiple interfaces are present, it is assumed the host will forward packets between networks. in.ripngd then multicasts a request packet on each IPv6 interface and enters a loop, listening for request and response packets from other hosts.

When a request packet is received, in.ripngd formulates a reply based on the information maintained in its internal tables. The response packet contains a list of known routes. With each route is a number specifying the number of bits in the prefix. The prefix is the number of bits in the high order part of an address that indicate the subnet or network that the route describes. Each route reported also has a “*hop count*” metric. A count of 16 or greater is considered “infinity.” The metric associated with each route returned provides a metric relative to the sender.

The request packets received by in.ripngd are used to update the routing tables if one of the following conditions is satisfied:

- No routing table entry exists for the destination network or host, and the metric indicates the destination is “reachable”, that is, the *hop count* is not infinite.
- The source host of the packet is the same as the router in the existing routing table entry. That is, updated information is being received from the very internetwork router through which packets for the destination are being routed.
- The existing entry in the routing table has not been updated for a period of time, defined to be 90 seconds, and the route is at least as cost-effective as the current route.
- The new route describes a shorter route to the destination than the one currently stored in the routing tables; this is determined by comparing the metric of the new route against the one stored in the table.

When an update is applied, `in.ripngd` records the change in its internal tables and generates a response packet to all directly connected hosts and networks. To allow possible unstable situations to settle, `in.ripngd` waits a short period of time (no more than 30 seconds) before modifying the kernel's routing tables.

In addition to processing incoming packets, `in.ripngd` also periodically checks the routing table entries. If an entry has not been updated for 3 minutes, the entry's metric is set to infinity and marked for deletion. Deletions are delayed an additional 60 seconds to insure the invalidation is propagated throughout the internet.

Hosts acting as internetwork routers gratuitously supply their routing tables every 30 seconds to all directly connected hosts and networks.

选项

`in.ripngd` supports the options listed below. Listed with the options are the equivalent SMF property values. These are set for the `ripng:default` service with a command of the form:

```
# routeadm -m ripng:default key=value
```

- p *n* Send and receive the routing packets from other routers using the UDP port number *n*. Use of this option is equivalent to setting the `udp_port` property.
- P Do not use poison reverse. Use of this option is equivalent to setting the `poison_reverse` property to false.
- q Do not supply routing information. Use of this option is equivalent to setting the `quiet_mode` property to true.
- s Force `in.ripngd` to supply routing information whether it is acting as an internetwork router or not. Use of this option is equivalent to setting the `supply_routes` property to true.
- t Print all packets sent or received to standard output. `in.ripngd` will not divorce itself from the controlling terminal. Accordingly, interrupts from the keyboard will kill the process. Not supported by the `ripng` service.
- v Print all changes made to the routing tables to standard output with a timestamp. Use of this option is equivalent to setting the `verbose` property to true.

Any other argument supplied to this option is interpreted as the name of the file in which the actions of `in.ripngd`, as specified by this option or by `-t`, should be logged instead of being sent to standard output.

The logfile can be specified for the `ripng` service by means of the `log_file` property.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/routing

另请参见 [in.routed\(1M\)](#), [routeadm\(1M\)](#), [svcadm\(1M\)](#), [ioctl\(2\)](#), [attributes\(5\)](#), [smf\(5\)](#), [udp\(7P\)](#)

G. Malkin, R. Minnear, RFC 2080, RIPng for IPv6, January 1997.

附注 The kernel's routing tables may not correspond to those of `in.ripngd` for short periods of time while processes that utilize existing routes exit; the only remedy for this is to place the routing process in the kernel.

`in.ripngd` currently does not support all of the functionality of [in.routed\(1M\)](#). Future releases may support more if appropriate.

`in.ripngd` initially obtains a routing table by examining the interfaces configured on a machine. It then sends a request on all directly connected networks for more routing information. `in.ripngd` does not recognize or use any routing information already established on the machine prior to startup. With the exception of interface changes, `in.ripngd` does not see any routing table changes that have been done by other programs on the machine, for example, routes added, deleted or flushed by means of the [route\(1M\)](#) command. Therefore, these types of changes should not be done while `in.ripngd` is running. Rather, shut down `in.ripngd`, make the changes required, and then restart `in.ripngd`.

引用名	in.rlogind, rlogind – remote login server
用法概要	<pre>/usr/sbin/in.rlogind [-k5eExXciPp] [-s tos] [-S keytab] [-M realm]</pre>
描述	in.rlogind is the server for the rlogin(1) program. The server provides a remote login facility with authentication based on Kerberos V5 or privileged port numbers. in.rlogind is invoked by inetd(1M) when a remote login connection is established. When Kerberos V5 authentication is required (see option -k below), the authentication sequence is as follows: ▪ Check Kerberos V5 authentication. ▪ Check authorization according to the rules in krb5_auth_rules(5). ▪ Prompt for a password if any checks fail and /etc/pam.conf is configured to do so. In order for Kerberos authentication to work, a host/<FQDN> Kerberos principal must exist for each Fully Qualified Domain Name associated with the in.rlogind server. Each of these host/<FQDN> principals must have a keytab entry in the /etc/krb5/krb5.keytab file on the in.rlogind server. An example principal might be: <pre>host/bigmachine.eng.example.com</pre> See kadmin(1M) or gkadmin(1M) for instructions on adding a principal to a krb5.keytab file. See 《Oracle Solaris 管理：安全服务》 for a discussion of Kerberos authentication. If Kerberos V5 authentication is not enabled, then the authentication procedure follows the standard rlogin protocol: ▪ The server checks the client's source port. If the port is not in the range 512-1023, the server aborts the connection. ▪ The server checks the client's source address. If an entry for the client exists in both /etc/hosts and /etc/hosts.equiv, a user logging in from the client is not prompted for a password. If the address is associated with a host for which no corresponding entry exists in /etc/hosts, the user is prompted for a password, regardless of whether or not an entry for the client is present in /etc/hosts.equiv. See hosts(4) and hosts.equiv(4). Once the source port and address have been checked, in.rlogind allocates a pseudo-terminal and manipulates file descriptors so that the slave half of the pseudo-terminal becomes the stdin, stdout, and stderr for a login process. The login process is an instance of the login(1) program, invoked with the -r. The login process then proceeds with the pam(3PAM) authentication process. See SECURITY below. If automatic authentication fails, it reprompts the user to login. The parent of the login process manipulates the master side of the pseudo-terminal, operating as an intermediary between the login process and the client instance of the rlogin program.

In normal operation, a packet protocol is invoked to provide Ctrl-S and Ctrl-Q type facilities and propagate interrupt signals to the remote programs. The login process propagates the client terminal's baud rate and terminal type, as found in the environment variable, TERM.

选项

The following options are supported:

- 5 Same as -k, for backwards compatibility.
- c Requires Kerberos V5 clients to present a cryptographic checksum of initial connection information like the name of the user that the client is trying to access in the initial authenticator. This checksum provides additional security by preventing an attacker from changing the initial connection information. This option is mutually exclusive with the -i option.
- e Creates an encrypted session.
- E Same as -e, for backwards compatibility.
- i Ignores authenticator checksums if provided. This option ignores authenticator checksums presented by current Kerberos clients to protect initial connection information. Option -i is the opposite of option -c.
- k Allows Kerberos V5 authentication with the .k5login access control file to be trusted. If this authentication system is used by the client and the authorization check is passed, then the user is allowed to log in.
- M *realm* Uses the indicated Kerberos V5 realm. By default, the daemon will determine its realm from the settings in the [krb5.conf\(4\)](#) file.
- p Prompts for authentication only if other authentication checks fail.
- P Prompts for a password in addition to other authentication methods.
- s *tos* Sets the IP TOS option.
- S *keytab* Sets the KRB5 keytab file to use. The `/etc/krb5/krb5.keytab` file is used by default.
- x Same as -e, for backwards compatibility.
- X Same as -e, for backwards compatibility.

用法

`rlogind` and `in.rlogind` are IPv6-enabled. See [ip6\(7P\)](#). IPv6 is not currently supported with Kerberos V5 authentication.

Typically, Kerberized `rlogin` service runs on port 543 (`klogin`) and Kerberized, encrypted `rlogin` service runs on port 2105 (`eklogin`). The corresponding FMRI entries are:

```
svc:/network/login:klogin (rlogin with kerberos)
svc:/network/login:eklogin (rlogin with kerberos and encryption)
```

安全

`in.rlogind` uses [pam\(3PAM\)](#) for authentication, account management, and session management. The PAM configuration policy, listed through `/etc/pam.conf`, specifies the modules to be used for `in.rlogind`. Here is a partial `pam.conf` file with entries for the `rlogin` command using the “rhosts” and UNIX authentication modules, and the UNIX account, session management, and password management modules.

<code>rlogin</code>	<code>auth sufficient</code>	<code>pam_rhosts_auth.so.1</code>
<code>rlogin</code>	<code>auth requisite</code>	<code>pam_authtok_get.so.1</code>
<code>rlogin</code>	<code>auth required</code>	<code>pam_dhkeys.so.1</code>
<code>rlogin</code>	<code>auth required</code>	<code>pam_unix_auth.so.1</code>
<code>rlogin</code>	<code>account required</code>	<code>pam_unix_roles.so.1</code>
<code>rlogin</code>	<code>account required</code>	<code>pam_unix_projects.so.1</code>
<code>rlogin</code>	<code>account required</code>	<code>pam_unix_account.so.1</code>
<code>rlogin</code>	<code>session required</code>	<code>pam_unix_session.so.1</code>

With this configuration, the server checks the client's source address. If an entry for the client exists in both `/etc/hosts` and `/etc/hosts.equiv`, a user logging in from the client is not prompted for a password. If the address is associated with a host for which no corresponding entry exists in `/etc/hosts`, the user is prompted for a password, regardless of whether or not an entry for the client is present in `/etc/hosts.equiv`. See [hosts\(4\)](#) and [hosts.equiv\(4\)](#).

When running a Kerberized `rlogin` service (with or without the encryption option), the `pam` service name that should be used is “`krlogin`”.

If there are no entries for the `rlogin` service, then the entries for the “other” service will be used. If multiple authentication modules are listed, then the user may be prompted for multiple passwords. Removing the `pam_rhosts_auth.so.1` entry will disable the `/etc/hosts.equiv` and `~/.rhosts` authentication protocol and the user would always be forced to type the password. The *sufficient* flag indicates that authentication through the `pam_rhosts_auth.so.1` module is sufficient to authenticate the user. Only if this authentication fails is the next authentication module used.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-servers

另请参见

[login\(1\)](#), [svcs\(1\)](#), [rlogin\(1\)](#), [gkadmin\(1M\)](#), [in.rshd\(1M\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [kadmin\(1M\)](#), [svcadm\(1M\)](#), [pam\(3PAM\)](#), [hosts\(4\)](#), [hosts.equiv\(4\)](#), [krb5.conf\(4\)](#), [pam.conf\(4\)](#), [attributes\(5\)](#), [environ\(5\)](#), [krb5_auth_rules\(5\)](#), [pam_authtok_check\(5\)](#), [pam_authtok_get\(5\)](#), [pam_authtok_store\(5\)](#), [pam_dhkeys\(5\)](#), [pam_passwd_auth\(5\)](#), [pam_unix_account\(5\)](#), [pam_unix_auth\(5\)](#), [pam_unix_session\(5\)](#), [smf\(5\)](#)

《Oracle Solaris 管理：安全服务》

诊断

All diagnostic messages are returned on the connection associated with the `stderr`, after which any network connections are closed. An error is indicated by a leading byte with a value of 1.

Hostname for your address unknown. No entry in the host name database existed for the client's machine.

Try again. A *fork* by the server failed.

/usr/bin/sh: ... The user's login shell could not be started.

附注

The authentication procedure used here assumes the integrity of each client machine and the connecting medium. This is insecure, but it is useful in an “open” environment.

A facility to allow all data exchanges to be encrypted should be present.

The `in.rlogind` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/login:rlogin (rlogin)
svc:/network/login:klogin (rlogin with kerberos)
svc:/network/login:eklogin (rlogin with kerberos and encryption)
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名	in.routed, routed – network routing daemon
用法概要	<pre>/usr/sbin/in.routed [-AdghmnqsStVz] [-T <i>tracefile</i> [-v]] [-F <i>net</i>[/<i>mask</i>][,<i>metric</i>]] [-P <i>params</i>]</pre>
描述	The daemon <code>in.routed</code>, often referred to as <code>routed</code>, is invoked at boot time to manage the network routing tables. It uses Routing Information Protocol, RIPv1 (RFC 1058), RIPv2 (RFC 2453), and Internet Router Discovery Protocol (RFC 1256) to maintain the kernel routing table. The RIPv1 protocol is based on the reference 4.3BSD daemon. <code>in.routed</code> is managed by means of the service management facility (SMF), using the fault management resource identifier (FMRI): <pre>svc:/network/routing/route:default</pre> The daemon listens on a udp socket for the route service (see services(4)) for Routing Information Protocol packets. It also sends and receives multicast Router Discovery ICMP messages. If the host is a router, <code>in.routed</code> periodically supplies copies of its routing tables to any directly connected hosts and networks. It also advertises or solicits default routes using Router Discovery ICMP messages. When started (or when a network interface is later turned on), <code>in.routed</code> uses an AF_ROUTE address family facility to find those directly connected interfaces configured into the system and marked “up”. It adds necessary routes for the interfaces to the kernel routing table. Soon after being first started, and provided there is at least one interface on which RIP has not been disabled, <code>in.routed</code> deletes all pre-existing non-static routes in the kernel table. Static routes in the kernel table are preserved and included in RIP responses if they have a valid RIP metric (see route(1M)). If more than one interface is present (not counting the loopback interface), it is assumed that the host should forward packets among the connected networks. After transmitting a RIP request and Router Discovery Advertisements or Solicitations on a new interface, the daemon enters a loop, listening for RIP request and response and Router Discovery packets from other hosts. When a request packet is received, <code>in.routed</code> formulates a reply based on the information maintained in its internal tables. The response packet generated contains a list of known routes, each marked with a “hop count” metric (a count of 16 or greater is considered “infinite”). Advertised metrics reflect the metric associated with an interface (see ifconfig(1M)), so setting the metric on an interface is an effective way to steer traffic. Responses do not include routes with a first hop on the requesting network, to implement in part split-horizon. Requests from query programs such as rtquery(1M) are answered with the complete table. The routing table maintained by the daemon includes space for several gateways for each destination to speed recovery from a failing router. RIP response packets received are used to

update the routing tables, provided they are from one of the several currently recognized gateways or advertise a better metric than at least one of the existing gateways.

When an update is applied, `in.routed` records the change in its own tables and updates the kernel routing table if the best route to the destination changes. The change in the kernel routing table is reflected in the next batch of response packets sent. If the next response is not scheduled for a while, a flash update response containing only recently changed routes is sent.

In addition to processing incoming packets, `in.routed` also periodically checks the routing table entries. If an entry has not been updated for 3 minutes, the entry's metric is set to infinity and marked for deletion. Deletions are delayed until the route has been advertised with an infinite metric to insure the invalidation is propagated throughout the local internet. This is a form of poison reverse.

Routes in the kernel table that are added or changed as a result of ICMP Redirect messages are deleted after a while to minimize black-holes. When a TCP connection suffers a timeout, the kernel tells `in.routed`, which deletes all redirected routes through the gateway involved, advances the age of all RIP routes through the gateway to allow an alternate to be chosen, and advances the age of any relevant Router Discovery Protocol default routes.

Hosts acting as internetwork routers gratuitously supply their routing tables every 30 seconds to all directly connected hosts and networks. These RIP responses are sent to the broadcast address on nets that support broadcasting, to the destination address on point-to-point links, and to the router's own address on other networks. If RIPv2 is enabled, multicast packets are sent on interfaces that support multicasting.

If no response is received on a remote interface, if there are errors while sending responses, or if there are more errors than input or output (see [netstat\(1M\)](#)), then the cable or some other part of the interface is assumed to be disconnected or broken, and routes are adjusted appropriately.

The Internet Router Discovery Protocol is handled similarly. When the daemon is supplying RIP routes, it also listens for Router Discovery Solicitations and sends Advertisements. When it is quiet and listening to other RIP routers, it sends Solicitations and listens for Advertisements. If it receives a good Advertisement and it is not multi-homed, it stops listening for broadcast or multicast RIP responses. It tracks several advertising routers to speed recovery when the currently chosen router dies. If all discovered routers disappear, the daemon resumes listening to RIP responses. It continues listening to RIP while using Router Discovery if multi-homed to ensure all interfaces are used.

The Router Discovery standard requires that advertisements have a default “lifetime” of 30 minutes. That means should something happen, a client can be without a good route for 30 minutes. It is a good idea to reduce the default to 45 seconds using `-P rdisc_interval=45` on the command line or `rdisc_interval=45` in the `/etc/gateways` file. See [gateways\(4\)](#).

While using Router Discovery (which happens by default when the system has a single network interface and a Router Discover Advertisement is received), there is a single default route and a variable number of redirected host routes in the kernel table. On a host with more than one network interface, this default route will be via only one of the interfaces. Thus, multi-homed hosts running with `-q` might need the `no_rdisc` argument described below.

To support “legacy” systems that can handle neither RIPv2 nor Router Discovery, you can use the `pm_rdisc` parameter in the `/etc/gateways`. See [gateways\(4\)](#).

By default, neither Router Discovery advertisements nor solicitations are sent over point-to-point links (for example, PPP). The Solaris OE uses a netmask of all ones (255.255.255.255) on point-to-point links.

`in.routed` supports the notion of “distant” passive or active gateways. When the daemon is started, it reads the file `/etc/gateways` to find such distant gateways that cannot be located using only information from a routing socket, to discover if some of the local gateways are passive, and to obtain other parameters. Gateways specified in this manner should be marked passive if they are not expected to exchange routing information, while gateways marked active should be willing to exchange RIP packets. Routes through passive gateways are installed in the kernel's routing tables once upon startup and are not included in transmitted RIP responses.

Distant active gateways are treated like network interfaces. RIP responses are sent to the distant active gateway. If no responses are received, the associated route is deleted from the kernel table and RIP responses are advertised via other interfaces. If the distant gateway resumes sending RIP responses, the associated route is restored.

Distant active gateways can be useful on media that do not support broadcasts or multicasts but otherwise act like classic shared media, such as some ATM networks. One can list all RIP routers reachable on the HIPPI or ATM network in `/etc/gateways` with a series of “host” lines. Note that it is usually desirable to use RIPv2 in such situations to avoid generating lists of inferred host routes.

Gateways marked external are also passive, but are not placed in the kernel routing table, nor are they included in routing updates. The function of external entries is to indicate that another routing process will install such a route if necessary, and that other routes to that destination should not be installed by `in.routed`. Such entries are required only when both routers might learn of routes to the same destination.

选项

Listed below are available options. Any other argument supplied is interpreted as the name of a file in which the actions of `in.routed` should be logged. It is better to use `-T` (described below) instead of appending the name of the trace file to the command. Associated SMF properties for these options are described, and can be set by means of a command of the form:

```
# routeadm -m route:default name=value
```

-A

Do not ignore RIPv2 authentication if we do not care about RIPv2 authentication. This option is required for conformance with RFC 2453. However, it makes no sense and breaks using RIP as a discovery protocol to ignore all RIPv2 packets that carry authentication when this machine does not care about authentication. This option is equivalent to setting the `ignore_auth` property value to false.

-d

Do not run in the background. This option is meant for interactive use and is not usable under the SME.

-F *net* [*/mask*] [*,metric*]

Minimize routes in transmissions via interfaces with addresses that match *net* (network number)/*mask* (netmask), and synthesizes a default route to this machine with the *metric*. The intent is to reduce RIP traffic on slow, point-to-point links, such as PPP links, by replacing many large UDP packets of RIP information with a single, small packet containing a “fake” default route. If *metric* is absent, a value of 14 is assumed to limit the spread of the “fake” default route. This is a dangerous feature that, when used carelessly, can cause routing loops. Notice also that more than one interface can match the specified network number and mask. See also -g. Use of this option is equivalent to setting the `minimize_routes` property.

-g

Used on internetwork routers to offer a route to the “default” destination. It is equivalent to -F 0/0, 1 and is present mostly for historical reasons. A better choice is -P `pm_rdisc` on the command line or `pm_rdisc` in the `/etc/gateways` file. A larger metric will be used with the latter alternatives, reducing the spread of the potentially dangerous default route. The -g (or -P) option is typically used on a gateway to the Internet, or on a gateway that uses another routing protocol whose routes are not reported to other local routers. Note that because a metric of 1 is used, this feature is dangerous. Its use more often creates chaos with a routing loop than solves problems. Use of this option is equivalent to setting the `offer_default_route` property to true.

-h

Causes host or point-to-point routes not to be advertised, provided there is a network route going the same direction. That is a limited kind of aggregation. This option is useful on gateways to LANs that have other gateway machines connected with point-to-point links such as SLIP. Use of this option is equivalent to setting the `advertise_host_routes` property to false.

-m

Cause the machine to advertise a host or point-to-point route to its primary interface. It is useful on multi-homed machines such as NFS servers. This option should not be used except when the cost of the host routes it generates is justified by the popularity of the server. It is effective only when the machine is supplying routing information, because there is more than one interface. The -m option overrides the -q option to the limited extent

of advertising the host route. Use of this option is equivalent to setting the `advertise_host_routes_primary` property to true.

-n

Do not install routes in kernel. By default, routes are installed in the kernel. Use of this option is equivalent to setting the `install_routes` property to false.

-P *params*

Equivalent to adding the parameter line *params* to the `/etc/gateways` file. Can also be set by means of the `parameters` property.

-q

Opposite of the `-s` option. This is the default when only one interface is present. With this explicit option, the daemon is always in “quiet mode” for RIP and does not supply routing information to other computers. Use of this option is equivalent to setting the `quiet_mode` property to true.

-s

Force `in.routed` to supply routing information. This is the default if multiple network interfaces are present on which RIP or Router Discovery have not been disabled, and if global IPv4 forwarding is turned on (by means of `ipadm(1M)`). Use of this option is equivalent to setting the `supply_routes` property to true.

-S

If `in.routed` is not acting as an internetwork router, instead of entering the whole routing table in the kernel, it enters only a default route for each internetwork router. This reduces the memory requirements without losing any routing reliability. This option is provided for compatibility with the previous, RIPv1-only `in.routed`. Use of this option is generally discouraged. Use of this option is equivalent to setting the `default_routes_only` property to true.

-t

Runs in the foreground (as with `-d`) and logs the contents of the packets received (as with `-zz`). This is for compatibility with prior versions of Solaris and has no SMF equivalent.

-T *tracefile*

Increases the debugging level to at least 1 and causes debugging information to be appended to the trace file. Because of security concerns, do not to run `in.routed` routinely with tracing directed to a file. Use of this option is equivalent to setting the `log_file` property to *trace file path*.

-v

Enables debug. Similar to `-z`, except, where `-z` increments `trace_level`, `-v` sets `trace_level` to 1. Also, `-v` requires the `-T` option. Use of this option is equivalent to setting the `debug` property to true.

-V

Displays the version of the daemon.

- Z
Increase the debugging level, which causes more information to be logged on the tracefile specified with - T or stdout. The debugging level can be increased or decreased with the SIGUSR1 or SIGUSR2 signals or with the [rtquery\(1M\)](#) command.

文件 /etc/defaultrouter If this file is present and contains the address of a default router, the system startup script does not run in . routed. See [defaultrouter\(4\)](#).

 /etc/gateways List of distant gateways and general configuration options for in . routed. See [gateways\(4\)](#).

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/routing

另请参见 [ipadm\(1M\)](#), [route\(1M\)](#), [routeadm\(1M\)](#), [rtquery\(1M\)](#), [svcadm\(1M\)](#), [ioctl\(2\)](#), [inet\(3SOCKET\)](#), [defaultrouter\(4\)](#), [gateways\(4\)](#), [attributes\(5\)](#), [icmp\(7P\)](#), [inet\(7P\)](#), [udp\(7P\)](#)

Internet Transport Protocols, XSI S 028112, Xerox System Integration Standard

Routing Information Protocol, v2 (RFC 2453, STD 0056, November 1998)

RIP-v2 MD5 Authentication (RFC 2082, January 1997)

Routing Information Protocol, v1 (RFC 1058, June 1988)

ICMP Router Discovery Messages (RFC 1256, September 1991)

附注 In keeping with its intended design, this daemon deviates from RFC 2453 in two notable ways:

- By default, in . routed does not discard authenticated RIPv2 messages when RIP authentication is not configured. There is little to gain from dropping authenticated packets when RIPv1 listeners will gladly process them. Using the - A option causes in . routed to conform to the RFC in this case.
- Unauthenticated RIP requests are never discarded, even when RIP authentication is configured. Forwarding tables are not secret and can be inferred through other means such as test traffic. RIP is also the most common router-discovery protocol, and hosts need to send queries that will be answered.

in . routed does not always detect unidirectional failures in network interfaces, for example, when the output side fails.

引用名	in.rshd, rshd – remote shell server
用法概要	<pre>in.rshd [-k5eciU] [-s tos] [-S keytab] [-M realm] [-L env_var] host.port</pre>
描述	in.rshd is the server for the rsh(1) program. The server provides remote execution facilities with authentication based on Kerberos V5 or privileged port numbers. in.rshd is invoked by inetd(1M) each time a shell service is requested. When Kerberos V5 authentication is required (this can be set with Kerberos-specific options listed below), the following protocol is initiated: 1. Check Kerberos V5 authentication. 2. Check authorization according to rules in krb5_auth_rules(5). 3. A null byte is returned on the initial socket and the command line is passed to the normal login shell of the user. (The PATH variable is set to /usr/bin.) The shell inherits the network connections established by in.rshd. In order for Kerberos authentication to work, a host/<FQDN> Kerberos principal must exist for each Fully Qualified Domain Name associated with the in.rshd server. Each of these host/<FQDN> principals must have a keytab entry in the /etc/krb5/krb5.keytab file on the in.rshd server. An example principal might be: <pre>host/bigmachine.eng.example.com</pre> See kadmin(1M) or gkadmin(1M) for instructions on adding a principal to a krb5.keytab file. See 《Oracle Solaris 管理：安全服务》 for a discussion of Kerberos authentication. If Kerberos V5 authentication is not enabled, then in.rshd executes the following protocol: 1. The server checks the client's source port. If the port is not in the range 512-1023, the server aborts the connection. The client's host address (in hex) and port number (in decimal) are the arguments passed to in.rshd. 2. The server reads characters from the socket up to a null (\0) byte. The resultant string is interpreted as an ASCII number, base 10. 3. If the number received in step 2 is non-zero, it is interpreted as the port number of a secondary stream to be used for the stderr. A second connection is then created to the specified port on the client's machine. The source port of this second connection is also in the range 512-1023. 4. A null-terminated user name of at most 16 characters is retrieved on the initial socket. This user name is interpreted as the user identity on the <i>client's</i> machine. 5. A null terminated user name of at most 16 characters is retrieved on the initial socket. This user name is interpreted as a user identity to use on the <i>server's</i> machine.

6. A null terminated command to be passed to a shell is retrieved on the initial socket. The length of the command is limited by the upper bound on the size of the system's argument list.
7. `in.rshd` then validates the user according to the following steps. The remote user name is looked up in the password file and a `chdir` is performed to the user's home directory. If the lookup fails, the connection is terminated. If the `chdir` fails, it does a `chdir` to `/` (root). If the user is not the superuser, (user ID 0), and if the `pam_rhosts_auth` PAM module is configured for authentication, the file `/etc/hosts.equiv` is consulted for a list of hosts considered "equivalent". If the client's host name is present in this file, the authentication is considered successful. See the SECURITY section below for a discussion of PAM authentication.

If the lookup fails, or the user is the superuser, then the file `.rhosts` in the home directory of the remote user is checked for the machine name and identity of the user on the client's machine. If this lookup fails, the connection is terminated
8. A null byte is returned on the initial connection and the command line is passed to the normal login shell of the user. The `PATH` variable is set to `/usr/bin`. The shell inherits the network connections established by `in.rshd`.

选项

The following options are supported:

- 5 Same as `-k`, for backwards compatibility
- c Requires Kerberos V5 clients to present a cryptographic checksum of initial connection information like the name of the user that the client is trying to access in the initial authenticator. This checksum provides additional security by preventing an attacker from changing the initial connection information. This option is mutually exclusive with the `-i` option.
- e Requires the client to encrypt the connection.
- i Ignores authenticator checksums if provided. This option ignores authenticator checksums presented by current Kerberos clients to protect initial connection information. Option `-i` is the opposite of option `-c`.
- k Allows Kerberos V5 authentication with the `.k5login` access control file to be trusted. If this authentication system is used by the client and the authorization check is passed, then the user is allowed to log in.
- L *env_var* List of environment variables that need to be saved and passed along.
- M *realm* Uses the indicated Kerberos V5 realm. By default, the daemon will determine its realm from the settings in the `krb5.conf(4)` file.
- s *tos* Sets the IP TOS option.
- S *keytab* Sets the KRB5 keytab file to use. The `/etc/krb5/krb5.keytab` file is used by default.

-U Refuses connections that cannot be mapped to a name through the [getnameinfo\(3SOCKET\)](#) function.

用法

rshd and in.rshd are IPv6-enabled. See [ip6\(7P\)](#). IPv6 is not currently supported with Kerberos V5 authentication.

The Kerberized rshd service runs on port 544 (kshell). The corresponding FMRI entry is:

```
svc:/network/shell:kshe11 (rshd with kerberos (ipv4 only))
```

安全

in.rshd uses [pam\(3PAM\)](#) for authentication, account management, and session management. The PAM configuration policy, listed through /etc/pam.conf, specifies the modules to be used for in.rshd. Here is a partial pam.conf file with entries for the rsh command using rhosts authentication, UNIX account management, and session management module.

rsh	auth	required	pam_rhosts_auth.so.1
rsh	account	required	pam_unix_roles.so.1
rsh	session	required	pam_unix_projects.so.1
rsh	session	required	pam_unix_account.so.1
rsh	session	required	pam_unix_session.so.1

If there are no entries for the rsh service, then the entries for the “other” service are used. To maintain the authentication requirement for in.rshd, the rsh entry must always be configured with the pam_rhosts_auth.so.1 module.

in.rshd can authenticate using Kerberos V5 authentication or [pam\(3PAM\)](#). For Kerberized rsh service, the appropriate PAM service name is krsh.

文件

/etc/hosts.equiv

\$HOME/.k5login File containing Kerberos principals that are allowed access.

/etc/krb5/krb5.conf Kerberos configuration file.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-servers

另请参见

[rsh\(1\)](#), [svcs\(1\)](#), [gkadmin\(1M\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [kadmin\(1M\)](#), [svcadm\(1M\)](#), [pam\(3PAM\)](#), [getnameinfo\(3SOCKET\)](#), [hosts\(4\)](#), [krb5.conf\(4\)](#), [pam.conf\(4\)](#), [attributes\(5\)](#), [environ\(5\)](#), [krb5_auth_rules\(5\)](#), [pam_authtok_check\(5\)](#), [pam_authtok_get\(5\)](#), [pam_authtok_store\(5\)](#), [pam_dhkeys\(5\)](#), [pam_passwd_auth\(5\)](#), [pam_rhosts_auth\(5\)](#), [pam_unix_account\(5\)](#), [pam_unix_auth\(5\)](#), [pam_unix_session\(5\)](#), [smf\(5\)](#), [ip6\(7P\)](#)

《Oracle Solaris 管理：安全服务》

诊断

The following diagnostic messages are returned on the connection associated with `stderr`, after which any network connections are closed. An error is indicated by a leading byte with a value of 1 in step 8 above (0 is returned above upon successful completion of all the steps prior to the command execution).

locuser too long	The name of the user on the client's machine is longer than 16 characters.
remuser too long	The name of the user on the remote machine is longer than 16 characters.
command too long	The command line passed exceeds the size of the argument list (as configured into the system).
Hostname for your address unknown.	No entry in the host name database existed for the client's machine.
Login incorrect.	No password file entry for the user name existed.
Permission denied.	The authentication procedure described above failed.
Can't make pipe.	The pipe needed for the <code>stderr</code> was not created.
Try again.	A <i>fork</i> by the server failed.

附注

The authentication procedure used here assumes the integrity of each client machine and the connecting medium. This is insecure, but it is useful in an “open” environment.

A facility to allow all data exchanges to be encrypted should be present.

The `in.rshd` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/shell:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名 in.rwhod, rwhod – system status server

用法概要 /usr/sbin/in.rwhod [-m *[ttl]*]

描述 in.rwhod is the server which maintains the database used by the [rwho\(1\)](#) and [ruptime\(1\)](#) programs. Its operation is predicated on the ability to broadcast or multicast messages on a network.

in.rwhod operates as both a producer and consumer of status information. As a producer of information it periodically queries the state of the system and constructs status messages which are broadcast or multicast on a network. As a consumer of information, it listens for other in.rwhod servers' status messages, validating them, then recording them in a collection of files located in the directory /var/spool/rwho.

The rwho server transmits and receives messages at the port indicated in the rwho service specification, see [services\(4\)](#). The messages sent and received are defined in /usr/include/protocols/rwhod.h and are of the form:

```
struct outmp {
    char    out_line[8];    /* tty name */
    char    out_name[8];    /* user id */
    long    out_time;       /* time on */
};
struct whod {
    char    wd_vers;
    char    wd_type;
    char    wd_fill[2];
    int     wd_sendtime;
    int     wd_recvtime;
    char    wd_hostname[32];
    int     wd_loadav[3];
    int     wd_boottime;
    struct  whoent {
        struct outmp we_utm;
        int     we_idle;
    } wd_we[1024 / sizeof (struct whoent)];
};
```

All fields are converted to network byte order prior to transmission. The load averages are as calculated by the [w\(1\)](#) program, and represent load averages over the 1, 5, and 15 minute intervals prior to a server's transmission. The host name included is that returned by the [uname\(2\)](#) system call. The array at the end of the message contains information about the users who are logged in to the sending machine. This information includes the contents of the [utmpx\(4\)](#) entry for each non-idle terminal line and a value indicating the time since a character was last received on the terminal line.

Messages received by the rwho server are discarded unless they originated at a rwho server's port. In addition, if the host's name, as specified in the message, contains any unprintable ASCII characters, the message is discarded. Valid messages received by in.rwhod are placed in files named whod.hostname in the directory /var/spool/rwho. These files contain only the most recent message, in the format described above.

Status messages are generated approximately once every 3 minutes.

选项

The following options are supported:

-m [ttl] Use the rwho IP multicast address (224.0.1.3) when transmitting. Receive announcements both on this multicast address and on the IP broadcast address. If ttl is not specified in.rwhod multicasts on all interfaces but with the IP TimeToLive set to 1 (that is, packets are not forwarded by multicast routers.) If ttl is specified in.rwhod only transmits packets on one interface and setting the IP TimeToLive to the specified ttl.

文件

/var/spool/rwho/whod.* information about other machines

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-servers

另请参见

[ruptime\(1\)](#), [rwho\(1\)](#), [w\(1\)](#), [uname\(2\)](#), [services\(4\)](#), [utmpx\(4\)](#), [attributes\(5\)](#)

警告

This service can cause network performance problems when used by several hosts on the network. It is not run at most sites by default. If used, include the -m multicast option.

附注

This service takes up progressively more network bandwidth as the number of hosts on the local net increases. For large networks, the cost becomes prohibitive.

in.rwhod should relay status information between networks. People often interpret the server dying as a machine going down.

引用名

install – install commands

用法概要

```

/usr/sbin/install -c dira [-m mode] [-u user] [-g group]
    [-o] [-s] file

/usr/sbin/install -f dirb [-m mode] [-u user] [-g group]
    [-o] [-s] file

/usr/sbin/install -n dirc [-m mode] [-u user] [-g group]
    [-o] [-s] file

/usr/sbin/install -d | -i [-m mode] [-u user] [-g group]
    [-o] [-s] dirx...

/usr/sbin/install [-m mode] [-u user] [-g group] [-o] [-s] file
    [dirx]...
```

描述

install is most commonly used in “makefiles” (see [make\(1S\)](#)) to install a file in specific locations, or to create directories within a file system. Each file is installed by copying it into the appropriate directory.

install uses no special privileges to copy files from one place to another. The implications of this are:

- You must have permission to read the files to be installed.
- You must have permission to copy into the destination directory.
- You must have permission to change the modes on the final copy of the file if you want to use the -m option.
- You must be super-user if you want to specify the ownership of the installed file with the -u or -g options. If you are not the super-user, the installed file is owned by you, regardless of who owns the original.

Note that if the ROOT environment variable is set, each of the default directory paths are prefixed by its value (for example, \$ROOT/bin and so on).

install prints messages telling the user exactly what files it is replacing or creating and where they are going.

If no options or directories (*dirx*...) are given, install searches a set of default directories (/bin, /usr/bin, /etc, /lib, and /usr/lib, in that order) for a file with the same name as *file*. When the first occurrence is found, install issues a message saying that it is overwriting that file with *file*, and proceeds to do so. If the file is not found, the program states this and exits.

If one or more directories (*dirx*...) are specified after *file*, those directories are searched before the default directories.

This version of install (/usr/sbin/install) is not compatible with the install binaries in many versions of Unix other than Solaris. For a higher degree of compatibility with other Unix versions, use /usr/ucb/install, which is described in the [install\(1B\)](#) man page.

- 选项
- The following options are supported:
- c *dira*

Install *file* in the directory specified by *dira*, if *file* does not yet exist. If it is found, *install* issues a message saying that the file already exists, and exits without overwriting it.
- f *dirb*

Force *file* to be installed in given directory, even if the file already exists. If the file being installed does not already exist, the mode and owner of the new file is set to 755 and *bin* , respectively. If the file already exists, the mode and owner is that of the already existing file.
- n *dirc*

If *file* is not found in any of the searched directories, it is put in the directory specified in *dirc*. The mode and owner of the new file is set to 755 and *bin*, respectively.
- d

Create a directory. Missing parent directories are created as required as in *mkdir -p*. If the directory already exists, the owner, group and mode is set to the values given on the command line.
- i

Ignore default directory list, searching only through the given directories (*dirx ...*).
- m *mode*

The mode of the new file is set to *mode*. Set to 0755 by default.
- u *user*

The owner of the new file is set to *user*. Only available to the super-user. Set to *bin* by default.
- g *group*

The group id of the new file is set to *group*. Only available to the super-user. Set to *bin* by default.
- o

If *file* is found, save the “found” file by copying it to *OLDfile* in the directory in which it was found. This option is useful when installing a frequently used file such as */bin/sh* or */lib/saf/ttymon*, where the existing file cannot be removed.
- s

Suppress printing of messages other than error messages.

用法

See [largefile\(5\)](#) for the description of the behavior of *install* when encountering files greater than or equal to 2 Gbyte (2³¹ bytes).

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[chgrp\(1\)](#), [chmod\(1\)](#), [chown\(1\)](#), [cp\(1\)](#), [install\(1B\)](#), [make\(1S\)](#), [mkdir\(1\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

引用名	installboot – 在磁盘分区中安装引导块
用法概要	<code>installboot [-F zfs ufs hsfs] bootblk raw-disk-device</code>
描述	<code>boot(1M)</code> 程序 <code>ufsboot</code> 是由驻留在磁盘分区的引导区域中的引导块程序从磁盘装入的。该程序特定于文件系统，必须与要引导的磁盘上的文件系统类型匹配。 引导对象依赖于平台，驻留在 <code>/usr/platform/platform-name/lib/fs/file-system</code> 目录中。可使用 <code>uname(1)</code> 的 <code>-i</code> 选项查找平台名称。可使用以下命令查找文件系统类型： <pre>% fstyp raw-disk-device</pre> 请参见 <code>fstyp(1M)</code>。 <code>installboot</code> 实用程序是一个仅适用于 SPARC 的程序。它在 x86 体系结构上不受支持。x86 用户应该使用 <code>installgrub(1M)</code>。
选项	支持以下选项： <code>-F zfs ufs hsfs</code> 指定要安装的引导块的文件系统类型。如果希望指定 <code>zfs</code> 或 <code>hsfs</code>，则此选项是必需的。缺省值为 <code>ufs</code>。
操作数	<code>bootblk</code> 引导块代码的名称。 <code>raw-disk-device</code> 要在其上安装引导块代码的磁盘设备的名称；它必须是可读写的字符设备。SCSI 或 IPI 驱动器的命名约定为 <code>c?t?d?s?</code>，而 IDE 驱动器的命名约定为 <code>c?d?s?</code>。
示例	示例1 安装 UFS 引导块 在运行命令的平台上，要在控制器 1 的目标 0 的分片 0 上安装 <code>ufs</code> 引导块，请使用以下命令： <pre># installboot /usr/platform/'uname -i'/lib/fs/ufs/bootblk \ /dev/rdsk/c1t0d0s0</pre> 示例2 安装 ZFS 引导块 在运行命令的平台上，要在控制器 1 的目标 0 的分片 0 上安装 <code>ZFS</code> 引导块，请使用如下所示的语法： <pre># installboot -F zfs /usr/platform/'uname -i'/lib/fs/zfs/bootblk \ /dev/rdsk/c1t1d0s0</pre>
文件	<code>/usr/platform/platform-name/lib/fs/</code> 引导对象所驻留的目录。
属性	有关以下属性的说明，请参见 <code>attributes(5)</code> ：

引用名	<code>installf</code> – add a file to the software installation database
用法概要	<pre>installf [-c class] [[-M] -R root_path] [-V fs_file] pkginst pathname [ftype [major minor] [mode owner group]] installf [-c class] [[-M] -R root_path] [-V fs_file] pkginst - installf -f [-c class] [[-M] -R root_path] [-V fs_file] pkginst</pre>
描述	<code>installf</code> informs the system that a pathname not listed in the pkgmap(4) file is being created or modified. It should be invoked before any file modifications have occurred. When the second synopsis is used, the pathname descriptions will be read from standard input. These descriptions are the same as would be given in the first synopsis but the information is given in the form of a list. The descriptions should be in the form: <pre>pathname [ftype [major minor] [mode owner group]]</pre> After all files have been appropriately created and/or modified, <code>installf</code> should be invoked with the <code>-f</code> synopsis to indicate that installation is final. Links will be created at this time and, if attribute information for a pathname was not specified during the original invocation of <code>installf</code>, or was not already stored on the system, the current attribute values for the pathname will be stored. Otherwise, <code>installf</code> verifies that attribute values match those given on the command line, making corrections as necessary. In all cases, the current content information is calculated and stored appropriately. Package commands are largefile(5)-aware. They handle files larger than 2 GB in the same way they handle smaller files. In their current implementations, pkgadd(1M), pkgtrans(1) and other package commands can process a datastream of up to 4 GB.
选项	<code>-c class</code> Class to which installed objects should be associated. Default class is none. <code>-f</code> Indicates that installation is complete. This option is used with the final invocation of <code>installf</code> (for all files of a given class). <code>-M</code> Instruct <code>installf</code> not to use the <code>\$root_path/etc/vfstab</code> file for determining the client's mount points. This option assumes the mount points are correct on the server and it behaves consistently with Solaris 2.5 and earlier releases. <code>-R root_path</code> Define the full path name of a directory to use as the <code>root_path</code>. All files, including package system information files, are relocated to a directory tree starting in the specified <code>root_path</code>. The <code>root_path</code> can be specified when installing to a client from a server (for example, <code>/export/root/client1</code>). <code>installf</code> inherits the value of the <code>PKG_INSTALL_ROOT</code> environment variable. (See <code>ENVIRONMENT VARIABLES</code>, below.) If <code>PKG_INSTALL_ROOT</code> is set, such as when the <code>-R</code> option is used with pkgadd(1M) or pkgrm(1M)

注 – The root file system of any non-global zones must not be referenced with the -R option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See [zones\(5\)](#).

	<i>-V fs_file</i>	Specify an alternative <i>fs_file</i> to map the client's file systems. For example, used in situations where the <i>\$root_path/etc/vfstab</i> file is non-existent or unreliable.
操作数	<i>pkginst</i>	Name of package instance with which the pathname should be associated.
	<i>pathname</i>	Pathname that is being created or modified.
	<i>ftype</i>	A one-character field that indicates the file type. Possible file types include: - b block special device - c character special device - d directory - e a file to be edited upon installation or removal - f a standard executable or data file - l linked file - p named pipe - s symbolic link - v volatile file (one whose contents are expected to change) - x an exclusive directory
	<i>major</i>	The major device number. The field is only specified for block or character special devices.
	<i>minor</i>	The minor device number. The field is only specified for block or character special devices.
	<i>mode</i>	The octal mode of the file (for example, 0664). A question mark (?) indicates that the mode will be left unchanged, implying that the file already exists on the target machine. This field is not used for linked or symbolically linked files.
	<i>owner</i>	The owner of the file (for example, bin or root). The field is limited to 14 characters in length. A question mark (?) indicates that the owner will be left unchanged, implying that the file already exists on the target machine. This field is not used for linked or symbolically linked files.
	<i>group</i>	The group to which the file belongs (for example, bin or sys). The field is limited to 14 characters in length. A question mark (?) indicates that the group will be left unchanged, implying that the file already exists on the target

machine. This field is not used for linked or symbolically linked files.

示例

```
示例1 Basic Usage

The following example shows the use of installf, invoked from an optional pre-install or
post-install script:

# create /dev/xt directory
# (needs to be done before drvininstall)
installf $PKGINST /dev/xt d 755 root sys ||
    exit 2
majno='/usr/sbin/drvininstall -m /etc/master.d/xt
    -d $BASEDIR/data/xt.o -v1.0' ||
    exit 2
i=00
while [ $i -lt $limit ]
do
    for j in 0 1 2 3 4 5 6 7
    do
        echo /dev/xt$i$j c $majno 'expr $i ? 8 + $j'
        644 root sys |
        echo /dev/xt$i$j=/dev/xt/$i$j
    done
    i='expr $i + 1'
    [ $i -le 9 ] && i="0$i" #add leading zero
done | installf $PKGINST - || exit 2
# finalized installation, create links
installf -f $PKGINST || exit 2
```

环境变量

installf inherits the value of the following environment variable. This variable is set when [pkgadd\(1M\)](#) or [pkgrm\(1M\)](#) is invoked with the -R option.

PKG_INSTALL_ROOT If present, defines the full path name of a directory to use as the system's PKG_INSTALL_ROOT path. All product and package information files are then looked for in the directory tree, starting with the specified PKG_INSTALL_ROOT path. If not present, the default system path of / is used.

退出状态

- 0 Successful operation.
- >0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[pkginfo\(1\)](#), [pkgmk\(1\)](#), [pkgparam\(1\)](#), [pkgproto\(1\)](#), [pkgtrans\(1\)](#), [pkgadd\(1M\)](#), [pkgask\(1M\)](#), [pkgchk\(1M\)](#), [pkgrm\(1M\)](#), [removef\(1M\)](#), [pkgmap\(4\)](#), [space\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

[《Application Packaging Developer’s Guide》](#)

附注

When *f*type is specified, all applicable fields, as shown below, must be defined:

<i>f</i> type	Required Fields
p, x, d, f, v, or e	mode owner group
c or b	major minor mode owner group

The `installf` command will create directories, named pipes and special devices on the original invocation. Links are created when `installf` is invoked with the `-f` option to indicate installation is complete.

Links should be specified as *path1=path2*. *path1* indicates the destination and *path2* indicates the source file.

Files installed with `installf` will be placed in the class `none`, unless a class is defined with the command. Subsequently, they will be removed when the associated package is deleted. If this file should not be deleted at the same time as the package, be certain to assign it to a class which is ignored at removal time. If special action is required for the file before removal, a class must be defined with the command and an appropriate class action script delivered with the package.

When classes are used, `installf` must be used in one of the following forms:

```
installf -c class1 . . .
installf -f -c class1 . . .
installf -c class2 . . .
installf -f -c class2 . . .
```

引用名	installgrub – 在磁盘分区中安装 GRUB
用法概要	<code>/usr/sbin/installgrub [-fm] stage1 stage2 raw-device</code>
描述	installgrub 命令是一个仅适用于 x86 的程序。GRUB 代表 G^Rand Uⁿified B^ootl^oader。 installgrub 在磁盘分区的引导区域上安装 GRUB 阶段 1 和阶段 2 文件。如果指定了 -m 选项，则 installgrub 将在磁盘的主引导扇区上安装阶段 1 文件。
选项	installgrub 命令接受下列选项： -f 覆盖主引导扇区时抑制交互。 -m 以交互方式在主引导扇区上安装 GRUB <i>stage1</i>。如果 Solaris 安装在扩展分区上，则必须使用此选项。
操作数	installgrub 命令接受下列操作数： <i>stage1</i> GRUB 阶段 1 文件的名称。 <i>stage2</i> GRUB 阶段 2 文件的名称。 <i>raw-device</i> 要在其上安装 GRUB 代码的设备的名称。它必须是可读写的字符设备。对于磁盘设备，请指定 GRUB 菜单文件所在的分片。（对于 Solaris，它是根分片。）
示例	示例1 在硬盘分片上安装 GRUB 以下命令在根分片为 <code>c0d0s0</code> 的系统上安装 GRUB： <pre>example# /usr/sbin/installgrub /boot/grub/stage1 \ /boot/grub/stage2 /dev/rdisk/c0d0s0</pre>
文件	<code>/boot/grub</code> GRUB 文件所驻留的目录。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	system/core-os
接口稳定性	Uncommitted（未确定）

另请参见 [boot\(1M\)](#)、[fdisk\(1M\)](#)、[fmthard\(1M\)](#)、[kernel\(1M\)](#)、[attributes\(5\)](#)

警告 在主引导扇区上安装 GRUB（-m 选项）会覆盖计算机上当前已安装的任何引导管理器。无论哪个 fdisk 分区处于活动状态，系统都始终在 Solaris 分区上引导 GRUB。

引用名 in.stdiscover – Service Tag Discovery Daemon

用法概要 /usr/lib/inet/in.stdiscover

描述 The in.stdiscover daemon allows a mechanism for discovering the location of the Service Tag Listener. By default, the in.stdiscover daemon listens for discovery probes (using a minimal built-in protocol) on UDP port 6481.

The daemon is under control of the service management facility, [smf\(5\)](#), under its inetd framework. It only runs upon demand and exits when no longer in use.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/management/service-tag
Interface Stability	Private

另请参见 [in.stlisten\(1M\)](#), [stclient\(1M\)](#)

引用名 in.stlisten – Service Tag Listener

用法概要 /usr/lib/inet/in.stlisten

描述 The `in.stlisten` daemon allows a mechanism for discovering the location of the Service Tag. By default, the `in.stlisten` daemon listens for discovery probes (using a minimal built-in protocol) on TCP port 6481.

The daemon is under control of the service management facility, [smf\(5\)](#), under its `inetd` framework. It only runs upon demand and exits when no longer in use.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/management/service-tag
Interface Stability	Private

另请参见 [in.stddiscover\(1M\)](#), [stclient\(1M\)](#), [svccfg\(1M\)](#), [attributes\(5\)](#), [environ\(5\)](#), [smf\(5\)](#)

附注 In open networks where the participants may not always be trusted, it is recommended that you deploy this daemon with the `passphrase-encryption` option. In [smf\(5\)](#) environments, the following command can be used to set the *passphrase*:

```
< prepare a text file "passfile" containing the passphrase >
# chown svctag:daemon passfile
# chmod 600 passfile
svccfg -s svc:/network/stlisten \
    setprop servicetag/passphrase=passfile
```

where *passfile* is the path of a file containing the intended *passphrase*.

This *passphrase* can be subsequently cleared as follows:

```
svccfg -s svc:/network/stlisten \
    setprop servicetag/passphrase=""
```

引用名 in.talkd, talkd – server for talk program

用法概要 in.talkd

描述 talkd is a server used by the [talk\(1\)](#) program. It listens at the UDP port indicated in the “talk” service description; see [services\(4\)](#). The actual conversation takes place on a TCP connection that is established by negotiation between the two machines involved.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-servers

另请参见 [svcs\(1\)](#), [talk\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [svcadm\(1M\)](#), [services\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The protocol is architecture dependent.

The in.talkd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/network/talk

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名	in.telnetd, telnetd – DARPA TELNET protocol server
用法概要	<pre>/usr/sbin/in.telnetd [-a authmode] [-EXUh] [-s tos] [-S keytab] [-M realm]</pre>
描述	in.telnetd is a server that supports the DARPA standard TELNET virtual terminal protocol. in.telnetd is normally invoked in the internet server (see inetd(1M)), for requests to connect to the TELNET port as indicated by the <code>/etc/services</code> file (see services(4)). in.telnetd operates by allocating a pseudo-terminal device for a client, then creating a login process which has the slave side of the pseudo-terminal as its standard input, output, and error. in.telnetd manipulates the master side of the pseudo-terminal, implementing the TELNET protocol and passing characters between the remote client and the login process. When a TELNET session starts up, in.telnetd sends TELNET options to the client side indicating a willingness to do <i>remote</i> echo of characters, and to <i>suppress go ahead</i>. The pseudo-terminal allocated to the client is configured to operate in “cooked” mode, and with XTABS, ICRNL and ONLCR enabled. See termio(7I). in.telnetd is willing to do: <i>echo</i>, <i>binary</i>, <i>suppress go ahead</i>, and <i>timing mark</i>. in.telnetd is willing to have the remote client do: <i>binary</i>, <i>terminal type</i>, <i>terminal size</i>, <i>logout option</i>, and <i>suppress go ahead</i>. in.telnetd also allows environment variables to be passed, provided that the client negotiates this during the initial option negotiation. The DISPLAY environment variable may be sent this way, either by the TELNET general environment passing methods, or by means of the XDISPLOC TELNET option. DISPLAY can be passed in the environment option during the same negotiation where XDISPLOC is used. Note that if you use both methods, use the same value for both. Otherwise, the results may be unpredictable. These options are specified in Internet standards RFC 1096, RFC 1408, RFC 1510, RFC 1571, RFC 2941, RFC 2942, RFC 2946, and RFC 1572. The following Informational draft is also supported: RFC 2952. The banner printed by in.telnetd is configurable. The default is (more or less) equivalent to ‘<code>uname -sr</code>’ and will be used if no banner is set in <code>/etc/default/telnetd</code>. To set the banner, add a line of the form <pre>BANNER="..."</pre> to <code>/etc/default/telnetd</code>. Nonempty banner strings are fed to shells for evaluation. The default banner may be obtained by <pre>BANNER="\r\n\r\n\r\n'uname -s' 'uname -r'\r\n\r\n\r\n"</pre> and no banner will be printed if <code>/etc/default/telnetd</code> contains <pre>BANNER=""</pre>

选项

The following options are supported:

- a *authmode* This option may be used for specifying what mode should be used for authentication. There are several valid values for *authmode*:
 - valid Only allows connections when the remote user can provide valid authentication information to identify the remote user, and is allowed access to the specified account without providing a password.
 - user Only allows connections when the remote user can provide valid authentication information to identify the remote user. The `login(1)` command will provide any additional user verification needed if the remote user is not allowed automatic access to the specified account.
 - none This is the default state. Authentication information is not required. If no or insufficient authentication information is provided, then the `login(1)` program provides the necessary user verification.
 - off This disables the authentication code. All user verification happens through the `login(1)` program.
- E Disables encryption support negotiation.
- h Disables displaying host specific information before login has been completed.
- M *realm* Uses the indicated Kerberos V5 realm. By default, the daemon will determine its realm from the settings in the `krb5.conf(4)` file.
- s *tos* Sets the IP TOS option.
- S *keytab* Sets the KRB5 keytab file to use. The `/etc/krb5/krb5.keytab` file is used by default.
- U Refuses connections that cannot be mapped to a name through the `getnameinfo(3SOCKET)` function.
- X Disables Kerberos V5 authentication support negotiation.

用法

`telnetd` and `in.telnetd` are IPv6-enabled. See [ip6\(7P\)](#).

安全

`in.telnetd` can authenticate using Kerberos V5 authentication, [pam\(3PAM\)](#), or both. By default, the telnet server will accept valid Kerberos V5 authentication credentials from a telnet client that supports Kerberos. `in.telnetd` can also support an encrypted session from such a client if the client requests it.

The telnet protocol only uses single DES for session protection—clients request service tickets with single DES session keys. The KDC must know that host service principals that offer the telnet service support single DES, which, in practice, means that such principals must have single DES keys in the KDC database.

In order for Kerberos authentication to work, a host/*<FQDN>* Kerberos principal must exist for each Fully Qualified Domain Name associated with the telnetd server. Each of these host/*<FQDN>* principals must have a keytab entry in the */etc/krb5/krb5.keytab* file on the telnetd server. An example principal might be:

```
host/bigmachine.eng.example.com
```

See [kadmin\(1M\)](#) or [gkadmin\(1M\)](#) for instructions on adding a principal to a *krb5.keytab* file. See 《Oracle Solaris 管理：安全服务》 for a discussion of Kerberos authentication.

in.telnetd uses [pam\(3PAM\)](#) for authentication, account management, session management, and password management. The PAM configuration policy, listed through */etc/pam.conf*, specifies the modules to be used for *in.telnetd*. Here is a partial *pam.conf* file with entries for the telnet command using the UNIX authentication, account management, session management, and password management modules.

```
telnet  auth requisite      pam_authtok_get.so.1
telnet  auth required      pam_dhkeys.so.1
telnet  auth required      pam_unix_auth.so.1

telnet  account requisite  pam_roles.so.1
telnet  account required  pam_projects.so.1
telnet  account required  pam_unix_account.so.1

telnet  session required  pam_unix_session.so.1

telnet  password required  pam_dhkeys.so.1
telnet  password requisite pam_authtok_get.so.1
telnet  password requisite pam_authtok_check.so.1
telnet  password required  pam_authtok_store.so.1
```

If there are no entries for the telnet service, then the entries for the “other” service will be used. If multiple authentication modules are listed, then the user may be prompted for multiple passwords.

For a Kerberized telnet service, the correct PAM service name is *ktelnet*.

- */etc/default/telnetd*

See [attributes\(5\)](#) for descriptions of the following attributes:

文件
属性

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/telnet

另请参见

[login\(1\)](#), [svcs\(1\)](#), [telnet\(1\)](#), [gkadmin\(1M\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [kadmin\(1M\)](#), [svcadm\(1M\)](#), [pam\(3PAM\)](#), [getnameinfo\(3SOCKET\)](#), [issue\(4\)](#), [krb5.conf\(4\)](#), [pam.conf\(4\)](#), [services\(4\)](#), [attributes\(5\)](#), [pam_authtok_check\(5\)](#), [pam_authtok_get\(5\)](#), [pam_authtok_store\(5\)](#), [pam_dhkeys\(5\)](#), [pam_passwd_auth\(5\)](#), [pam_unix_account\(5\)](#), [pam_unix_auth\(5\)](#), [pam_unix_session\(5\)](#), [smf\(5\)](#), [ip6\(7P\)](#), [termio\(7I\)](#)

《Oracle Solaris 管理：安全服务》

Alexander, S. RFC 1572, TELNET Environment Option. Network Information Center, SRI International, Menlo Park, Calif., January 1994.

Borman, Dave. RFC 1408, TELNET Environment Option. Network Information Center, SRI International, Menlo Park, Calif., January 1993.

Borman, Dave. RFC 1571, TELNET Environment Option Interoperability Issues. Network Information Center, SRI International, Menlo Park, Calif., January 1994.

Crispin, Mark. RFC 727, TELNET Logout Option. Network Information Center, SRI International, Menlo Park, Calif., April 1977.

Marcy, G. RFC 1096, TELNET X Display Location Option. Network Information Center, SRI International, Menlo Park, Calif., March 1989.

Postel, Jon, and Joyce Reynolds. RFC 854, TELNET Protocol Specification. Network Information Center, SRI International, Menlo Park, Calif., May 1983.

Waitzman, D. RFC 1073, TELNET Window Size Option. Network Information Center, SRI International, Menlo Park, Calif., October 1988.

Kohl, J., Neuman, C., The Kerberos Network Authentication Service (V5), RFC 1510. September 1993.

Ts'o, T. and J. Altman, Telnet Authentication Option, RFC 2941. September 2000.

Ts'o, T., Telnet Authentication: Kerberos Version 5, RFC 2942. September 2000.

Ts'o, T., Telnet Data Encryption Option, RFC 2946. September 2000.

Ts'o, T., Telnet Encryption: DES 64 bit Cipher Feedback, RFC 2952. September 2000.

附注

- Some TELNET commands are only partially implemented.
- Binary mode has no common interpretation except between similar operating systems.
- The terminal type name received from the remote client is converted to lower case.

The *packet* interface to the pseudo-terminal should be used for more intelligent flushing of input and output queues.

`in.telnetd` never sends TELNET *go ahead* commands.

The `in.telnetd` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/telnet
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名

in.tftpd, tftpd – Internet Trivial File Transfer Protocol server

用法概要

```
in.tftpd [-d] [-T rexmtval] [-s] [homedir]
```

描述

`tftpd` is a server that supports the Internet Trivial File Transfer Protocol (TFTP).

Before responding to a request, the server attempts to change its current directory to *homedir*; the default directory is `/tftpboot`.

The use of `tftp` does not require an account or password on the remote system. Due to the lack of authentication information, `in.tftpd` will allow only publicly readable files to be accessed. Files may be written only if they already exist and are publicly writable. Note that this extends the concept of “public” to include all users on all hosts that can be reached through the network. This may not be appropriate on all systems, and its implications should be considered before enabling this service.

`in.tftpd` runs with the user ID and group ID set to `[GU]ID_NOBODY` under the assumption that no files exist with that owner or group. However, nothing checks this assumption or enforces this restriction.

选项

-d

Debug. When specified it sets the `SO_DEBUG` socket option.

-s

Secure. When specified, the directory change to *homedir* must succeed. The daemon also changes its root directory to *homedir*.

-T *rexmtval*

Specifies the value of the retransmission timeout in seconds. This also affects the maximum session timeout in that the latter is set to five times the retransmission timeout value.

用法

The `in.tftpd` server is IPv6–enabled. See [ip6\(7P\)](#).

`in.tftpd` supports transfers of greater than 32 MB, per RFC 2348.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/tftp

另请参见

[svcs\(1\)](#), [tftp\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [svcadm\(1M\)](#), [netconfig\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [ip6\(7P\)](#)

Malkin, G. and Harkin, A. RFC 2347, TFTP Option Extension. The Internet Society. May 1998

Malkin, G. and Harkin, A. RFC 2348, TFTP Blocksize Option. The Internet Society. May 1998

Malkin, G. and Harkin, A. RFC 2349, TFTP Timeout Interval and Transfer Size Options. The Internet Society. May 1998

Sollins, K.R. RFC 1350, The TFTP Protocol (Revision 2). Network Working Group. July 1992.

附注

The `tftpd` server only acknowledges the transfer size option that is sent with a read request when the octet transfer mode is specified.

The `in.tftpd.1m` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/tftp/udp6:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

Unlike most [smf\(5\)](#) services, a manifest for the `tftp` service is not included in the system. To create one and enable this service, the administrator should:

1. Edit `/etc/inet/inetd.conf` and uncomment the `tftp` entry.
2. Run `/usr/sbin/inetconv`.

After you run `inetconv`, the `svc:/network/tftp/udp6:default` service is created and enabled.

引用名 in.timed – UDP or TCP time protocol service daemon

用法概要 in.timed

FMRI svc:/internet/time:default

描述 FMRI stands for Fault Management Resource Identifier. It is used to identify resources managed by the Fault Manager. See [fmd\(1M\)](#) and [smf\(5\)](#).

The in.timed service provides the server-side of the time protocol. The time server sends to requestors the time in seconds since midnight, January 1, 1900. The time protocol is available on both TCP and UDP transports through port 37.

The in.timed service is an [inetd\(1M\)](#) [smf\(5\)](#) delegated service. The in.timed detects which transport is requested by examining the socket it is passed by the inetd daemon.

- TCP-based service

Once a connection is established, the in.timed sends the time as a 32-bit binary number and closes the connection. Any received data is ignored.
- UDP-based service

The in.timed listens for UDP datagrams. When a datagram is received, the server generates a UDP datagram containing the time as a 32-bit binary number and sends it to the client. Any received data is ignored.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/legacy-network-services
Interface Stability	Committed

另请参见 [inetd\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

RFC 868

- 引用名

intrd – 中断分发守护进程
- 用法概要

intrd
- 描述

intrd 守护进程在引导时启动，用于监视中断和 CPU 之间的分配，如果 intrd 断定当前的分配不平衡并且会影响系统性能，它将生成并实施新的分配。

所有通知都将通过 [syslogd\(1M\)](#) 传送。

因为 intrd 动态监视系统以获得最佳性能，因此，即使在空闲系统上，它也只会占用少量的 CPU 时间。此行为是正常的。
- 属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/kernel/power
接口稳定性	Uncommitted（未确定）

另请参见

[svcs\(1\)](#)、[svcadm\(1M\)](#)、[syslogd\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

附注

中断分发守护进程由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

```
svc:/system/intrd:default
```

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

引用名 intrstat – 报告中断统计信息

用法概要 `/usr/sbin/intrstat [-c cpulist | -C processor_set_id] [-T u | d] [-x opt[=val]] [interval [count]]`

描述 `intrstat` 实用程序收集并显示运行时中断统计信息。输出是一个包含设备名称和 CPU ID 的表，该表的每一行表示一个设备，每一列表示一个 CPU。该表中每个单元格包含给定 CPU 上给定设备的原始中断数和该 CPU 上该设备的中断处理程序所花费的绝对时间的百分比。

设备名称是以 `{name}#{instance}` 形式给定的。该名称是标准化的驱动程序名称，通常对应于实施该驱动程序的模块名称。请参见 `ddi_driver_name(9F)`。许多 Oracle 交付的驱动程序有自己的手册页。请参见 `Intro(7)`。

如果标准输出是一个终端，该表将包含该终端宽度可以容纳的数据列。如果标准输出不是终端，该表最多可包含 4 个数据列。缺省情况下，会为所有 CPU 收集并显示数据。如果单个表容纳不下数据，则会在多个表中输出该数据。可以使用 `-c` 或 `-C` 选项指定要显示其数据的 CPU 集。

缺省情况下，`intrstat` 会每秒显示一次数据并无限期运行。可以分别使用 `interval` 和 `count` 参数来控制这两个行为。请参见“操作数”部分。

因为 `intrstat` 使用动态搜索，所以它只报告在该命令运行时发生中断的设备。不会显示在 `intrstat` 运行时处于静默状态的任何设备。

`intrstat` 会在系统范围内导致少量的性能降低。因此，缺省情况下，只有超级用户可以运行 `intrstat`。《Solaris 动态跟踪指南》中介绍了管理员如何授予其他用户特权以允许其运行 `intrstat`。

选项 所支持的选项如下：

`-c cpulist`
显示由 `cpulist` 指定的 CPU 的数据。

`cpulist` 可以是单个处理器 ID（例如，4）、处理器 ID 范围（例如，4-6），或者以逗号分隔的处理器 ID 或处理器 ID 范围的列表（例如，4,5,6 或 4,6-8）。

`-C processor_set_id`
显示由 `processor_set_id` 指定的处理器集中的 CPU 的数据。

`intrstat` 会修改其输出以始终反映指定的处理器集中的 CPU。如果向该集中添加了一个 CPU，`intrstat` 会修改其输出以包括该添加的 CPU。如果从该集中删除了一个 CPU，`intrstat` 会修改其输出以排除该删除的 CPU。最多可指定一个处理器集。

`-T u | d`
显示时间戳。

指定 `u` 表示时间的内部表示形式的印刷表示形式。请参见 `time(2)`。指定 `d` 表示标准日期格式。请参见 `date(1)`。

`-x opt[=val]`
启用或修改 DTrace 运行时选项或 D 编译器选项。选项列表可在《Solaris Dynamic Tracing Guide》中找到。布尔选项通过指定其名称启用。包含值的选项通过使用等号 (=) 分隔选项名称和值来设置。

操作数 支持下列操作数：

count
表示 intrstat 在退出之前显示其输出的次数。

interval
表示两次显示 intrstat 输出的间隔秒数。

示例 **示例1** 在不指定选项的情况下使用 intrstat

未指定选项时，intrstat 会显示一个包含陷阱类型和 CPU 的表。缺省终端宽度最多可以容纳 4 列。如果超过 4 个 CPU，则会显示多个表。

以下示例是在基于 Intel IA/32 的单处理器手提电脑上运行 intrstat 时的输出：

```
example# intrstat
      device |      cpu0 %tim
-----+-----
      ata#0 |      166  0.4
      ata#1 |         0  0.0
audioi810#0 |         6  0.0
i8042#0    |      281  0.7
iprb#0     |         6  0.0
uhci#1     |         6  0.0
uhci#2     |         6  0.0

      device |      cpu0 %tim
-----+-----
      ata#0 |      161  0.5
      ata#1 |         0  0.0
audioi810#0 |         6  0.0
i8042#0    |     303  0.6
iprb#0     |         6  0.0
uhci#1     |         6  0.0
uhci#2     |         6  0.0
...

```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/dtrace
接口稳定性	请参见下文。

命令行语法为 Committed（已确定）。人类可阅读的输出为 Uncommitted（未确定）。

另请参见

[dtrace\(1M\)](#)、[trapstat\(1M\)](#)、[attributes\(5\)](#)、[Intro\(7\)](#)、[ddi_driver_name\(9F\)](#)

《Solaris 动态跟踪指南》

引用名

in.uucpd, uucpd – UUCP server

用法概要

/usr/sbin/in.uucpd [-n]

描述

in.uucpd is the server for supporting UUCP connections over networks.

in.uucpd is invoked by [inetd\(1M\)](#) when a UUCP connection is established, that is, a connection to the port indicated in the “uucp” service specification, and executes the following protocol. See [services\(4\)](#):

1. The server prompts with login: . The [uucico\(1M\)](#) process at the other end must supply a username.
2. Unless the username refers to an account without a password, the server then prompts with Password: . The uucico process at the other end must supply the password for that account.

If the username is not valid, or is valid but refers to an account that does not have /usr/lib/uucp/uucico as its login shell, or if the password is not the correct password for that account, the connection is dropped. Otherwise, uucico is run, with the user ID, group ID, group set, and home directory for that account, with the environment variables USER and LOGNAME set to the specified username, and with a -u flag specifying the username. Unless the -n flag is specified, entries are made in /var/adm/utmpx, /var/adm/wtmpx, and /var/adm/lastlog for the username. in.uucpd must be invoked by a user with appropriate privilege (usually root) in order to be able to verify that the password is correct.

安全

in.uucpd uses [pam\(3PAM\)](#) for authentication, account management, and session management. The PAM configuration policy, listed through /etc/pam.conf, specifies the modules to be used for in.uucpd. Here is a partial pam.conf file with entries for uucp using the UNIX authentication, account management, and session management module.

uucp	auth requisite	pam_authok_get.so.1
uucp	auth required	pam_dhkeys.so.1
uucp	auth required	pam_unix_auth.so.1
uucp	account requisite	pam_roles.so.1
uucp	account required	pam_projects.so.1
uucp	account required	pam_unix_account.so.1
uucp	session required	pam_unix_session.so.1

If there are no entries for the uucp service, then the entries for the “other” service will be used. If multiple authentication modules are listed, then the peer may be prompted for multiple passwords.

文件

/var/adm/utmpx	accounting
/var/adm/wtmpx	accounting
/var/adm/lastlog	time of last login

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/uucp

另请参见 [svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [svcadm\(1M\)](#), [uucico\(1M\)](#), [pam\(3PAM\)](#), [pam.conf\(4\)](#), [services\(4\)](#), [attributes\(5\)](#), [pam_authtok_check\(5\)](#), [pam_authtok_get\(5\)](#), [pam_authtok_store\(5\)](#), [pam_dhkeys\(5\)](#), [pam_passwd_auth\(5\)](#), [pam_unix_account\(5\)](#), [pam_unix_auth\(5\)](#), [pam_unix_session\(5\)](#), [smf\(5\)](#)

诊断 All diagnostic messages are returned on the connection, after which the connection is closed.

user read	An error occurred while reading the username.
passwd read	An error occurred while reading the password.
Login incorrect.	The username is invalid or refers to an account with a login shell other than /usr/lib/uucp/uucico, or the password is not the correct password for the account.

附注 The in.uucpd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/network/uucp

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名	iostat – 报告 I/O 统计信息																
用法概要	<pre>/usr/bin/iostat [-cCdDeEiImMnpPrstxXYz] [-l n] [-T u d] [disk]... [interval [count]]</pre>																
描述	iostat 实用程序重复报告终端、磁盘和磁带 I/O 活动，以及 CPU 使用情况。输出的第一行针对自引导后的所有时间；随后每一行仅针对前一个时间间隔。 为计算此信息，内核要维护大量计数器。对于每个磁盘，内核都会对读取、写入、位读取和位写入情况进行计数。内核还会取得队列入口点和退出点处的 hi-res 时间戳，使它可以为每个队列跟踪驻留时间以及累积驻留长度。使用这些值，iostat 会得出关于总处理能力、使用情况、队列长度、事务处理速度和服务时间的非常准确的测量值。对于全体终端，内核只对输入和输出字符数进行计数。 内核状态命令执行期间，系统的 state 可以更改。如果相关，状态更改消息会包括在 iostat 输出中，格式为下列之一： <pre><<device added: sd0>> <<device removed: sd0>> <<partition added: sd0,a>> <<partition removed: sd0,a>> <<NFS mounted: nfs1>> <<NFS unmounted: nfs1>> <<multi-path added: ssd4>> <<multi-path removed: ssd4>> <<controller added: c1>> <<controller removed: c1>> <<processors added: 1, 3>> <<processors removed: 1, 3>></pre> 请注意，根据情况，这些状态更改消息中打印的 names 受 -n 和 -m 选项影响。 有关更多常规系统统计信息，请使用 sar(1)、sar(1M) 或 vmstat(1M)。																
输出	iostat 实用程序的输出包括以下信息。 <table><tr><td>device</td><td>磁盘名称</td></tr><tr><td>r/s</td><td>每秒读取次数</td></tr><tr><td>w/s</td><td>每秒写入次数</td></tr><tr><td>kr/s</td><td>每秒读取的千字节数</td></tr><tr><td></td><td>间隔期间的平均 I/O 大小可以由 kr/s 除以 r/s 计算得出。</td></tr><tr><td>kW/s</td><td>每秒写入的千字节数</td></tr><tr><td></td><td>间隔期间的平均 I/O 大小可以由 kW/s 除以 w/s 计算得出。</td></tr><tr><td>wait</td><td>等待服务的平均事务数（队列长度）</td></tr></table>	device	磁盘名称	r/s	每秒读取次数	w/s	每秒写入次数	kr/s	每秒读取的千字节数		间隔期间的平均 I/O 大小可以由 kr/s 除以 r/s 计算得出。	kW/s	每秒写入的千字节数		间隔期间的平均 I/O 大小可以由 kW/s 除以 w/s 计算得出。	wait	等待服务的平均事务数（队列长度）
device	磁盘名称																
r/s	每秒读取次数																
w/s	每秒写入次数																
kr/s	每秒读取的千字节数																
	间隔期间的平均 I/O 大小可以由 kr/s 除以 r/s 计算得出。																
kW/s	每秒写入的千字节数																
	间隔期间的平均 I/O 大小可以由 kW/s 除以 w/s 计算得出。																
wait	等待服务的平均事务数（队列长度）																

这是保留在设备驱动器队列中等待设备接受的 I/O 操作的数量。

activ 正在处理的平均事务数（已从队列删除但尚未完成）

这是被设备接受但仍未处理的 I/O 操作的数量。

svc_t 事务的平均响应时间（以毫秒为单位）

svc_t 输出报告设备的整体响应时间，而不是服务时间。整体时间包括事务在队列中的时间和事务被处理的时间。在队列中花费的时间以 **wsvc_t** 输出列中的 **-x** 选项显示。处理事务所花费的时间是真正的服务时间。服务时间也以 **-x** 选项显示，出现在同一报告的 **asvc_t** 输出列中。

%w 存在等待服务的事务的时间百分比（队列非空）

%b 磁盘忙碌的时间百分比（事务在进行中）

wsvc_t 等待队列中的平均服务时间（以毫秒为单位）

asvc_t 活动事务的平均服务时间（以毫秒为单位）

wt I/O 等待时间不再以 CPU 时间的百分比来计算，该统计将始终返回零。

选项

所支持的选项如下：

- c** 报告系统在用户模式、系统模式下、等待 I/O 和空闲的时间百分比。有关更多信息，请参见“说明”一节。
- C** 同时也选择了 **-x** 选项时，报告通过 *controller id* 聚合的扩展磁盘统计信息。
- d** 对于每个磁盘，报告每秒传输的千字节数、每秒的传输数量以及平均服务时间（以毫秒为单位）。
- D** 对于每个磁盘，报告每秒读取次数、每秒写入次数以及磁盘使用百分比。
- e** 显示设备错误汇总统计信息。显示全部错误、硬错误、软错误以及传输错误。
- M** 显示全部设备错误统计信息。
- i** 在 **-E** 输出中，显示设备 ID，而非序列号。Device Id 是驱动器通过 [ddi_devid_register\(9F\)](#) 注册的唯一标识符。
- I** 报告每个时间间隔内的计数，而非速度（如适用）。
- l n** 将报告中包括的磁盘数量限制为 *n*；**-d** 和 **-D** 的磁盘限制缺省为 4，对 **-x** 无限制。注意：明确请求的磁盘（请参见下面的 *disk*）不受此磁盘限制的约束。
- m** 报告文件系统挂载点。如果还指定了 **-P** 或 **-p** 选项或者其与 **-Xn** 或 **-en** 一起使用时，该选项最有用。仅当输出中实际列出挂载点时，**-m** 选项才有用。该选项仅可与 **-n** 选项共同使用。

- M 以 MB/秒（而非 KB/秒）显示数据吞吐量。
 - n 以描述性格式显示名称。例如，cXtYdZ、rmt/N、server:/export/path。

缺省情况下，通过诸如 ssd23 或 md301 等的实例名来标识磁盘。结合 -n 选项和 -x 选项，会以 cXtYdZsN 格式显示磁盘名称，从而更容易与其物理硬件特性关联。cXtYdZsN 格式在 FibreChannel (FC) 环境中特别有用，其中 FC 全局名称显示在 t 字段中。
 - p 对于每个磁盘，既按设备也按分区报告统计信息。
 - P 对每个磁盘，仅按分区报告统计信息，没有按磁盘统计信息。
 - r 以逗号分隔格式显示数据。
 - s 取消与 state changes 有关的消息。
 - t 报告每秒向终端读取和写入的字符数。
 - T u | d 显示时间戳。

指定 u 表示时间的内部表示形式的印刷表示形式。请参见 time(2)。指定 d 表示标准日期格式。请参见 date(1)。
 - X 对于 scsi_vhci(7D) 控制下的磁盘，除磁盘 lun 统计信息外，还会报告 lun.controller 的统计信息。
 - x 报告扩展磁盘统计信息。缺省情况下，通过诸如 ssd23 或 md301 等的实例名来标识磁盘。结合 x 选项和 -n 选项，会以 cXtYdZsN 格式显示磁盘名称，从而更容易与其物理硬件特性关联。使用 cXtYdZsN 格式在 FibreChannel 环境中特别有帮助，其中 FC 全局名称显示在 t 字段中。

如果不请求输出显示（无 -x、-e 或 -E），则暗含 -x。
 - Y 对于 scsi_vhci(7D) 控制下的磁盘，除磁盘 lun 统计信息外，还会报告 lun.targetport 和 lun.targetport.controller 的统计信息。

在 -n（描述性）模式下，使用路径的 target-port 属性显示 targetport。没有 -n，会使用较短 port-id 显示 targetport。所有具有相同 target-port 属性值的目标端口共享相同 port-id。target-port 到 port-id 关联在重新引导后不会保留。

如果不请求输出显示（无 -x、-e 或 -E），则暗含 -x。
 - z 不会打印底层数据值都是零的行。
- 选项设置 -xcnCXTdz interval 对于确定是否存在磁盘 I/O 问题以及识别问题特别有用。
- 操作数
- 支持下列操作数：
- count 仅显示 count 报告。

disk 明确指定要报告的磁盘；除了任何明确磁盘外，还将报告达到磁盘限制（请见上面的 *-l* 选项）的任何活动磁盘。

interval 每 *interval* 秒报告一次。

示例

示例1 使用 *iostat* 生成用户和系统操作统计信息

以下命令显示通过 *controller id* 为用户 (*us*) 和系统 (*sy*) 操作聚合的扩展设备统计信息的两个报告。因为 *-n* 选项与 *-x* 选项一起使用，以控制器名称标识设备。

```
example% iostat -xcnXTdz 5

Mon Nov 24 14:58:36 2003
      cpu
      us sy wt id
      14 31  0 20

      extended device statistics
      r/s    w/s    kr/s      kw wait  actv wsvc_t asvc_t  %w  %b device
      3.8    29.9   145.8    44.0 0.0    0.2    0.1    6.4   0   5   c0
666.3   814.8 12577.6 17591.1 91.3  82.3   61.6   55.6   0   2   c12
180.0   234.6  4401.1  5712.6  0.0 147.7    0.0  356.3   0  98  d10

Mon Nov 24 14:58:41 2003
      cpu
      us sy wt id
      11 31  0 22

      extended device statistics
      r/s    w/s    kr/s      kw wait  actv wsvc_t asvc_t  %w  %b device
      0.8    41.0     5.2    20.5 0.0    0.2    0.2    4.4   0   6   c0
565.3   581.7 8573.2 10458.9 0.0   26.6    0.0   23.2   0   3   c12
106.5    81.3  3393.2  1948.6 0.0    5.7    0.0   30.1   0  99  d10
```

示例2 使用 *iostat* 生成 TTY 统计信息

以下命令显示关于不同操作模式下五个磁盘的活动的两个报告。因为使用了 *-x* 选项，以实例名称标识磁盘。

```
example% iostat -x tc 5 2

      extended device statistics
device r/s  w/s kr/s  kw/s wait actv svc_t %w  %b
sd0    0.4  0.3 10.4   8.0  0.0  0.0  36.9  0   1
sd1    0.0  0.0  0.3   0.4  0.0  0.0  35.0  0   0
sd6    0.0  0.0  0.0   0.0  0.0  0.0   0.0  0   0
nfs1   0.0  0.0  0.0   0.0  0.0  0.0   0.0  0   0
nfs2   0.0  0.0  0.0   0.1  0.0  0.0  35.6  0   0

      extended device statistics
      tty
      tin tout
      0   10

      cpu
      us sy wt id
      0  0  0 99
```


示例2 使用 iostat 生成 TTY 统计信息 (续)

device	r/s	w/s	kr/s	kw/s	wait	actv	svc_t	%w	%b	tin	tout	us	sy	wt	id
sd0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	0	155	0	0	0	100
sd1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0						
sd6	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0						
nfs1	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0						
nfs2	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0						

示例3 使用 iostat 生成分区和设备统计信息

以下命令生成每个磁盘的分区和设备统计信息。因为 -n 选项与 -x 选项一起使用，以控制器名称标识磁盘。

example% iostat -xnp

extended device statistics														
r/s	w/s	kr/s	kw/s	wait	actv	wsvc_t	asvc_t	%w	%b	device				
0.4	0.3	10.4	7.9	0.0	0.0	0.0	36.9	0	1	c0t0d0				
0.3	0.3	9.0	7.3	0.0	0.0	0.0	37.2	0	1	c0t0d0s0				
0.0	0.0	0.1	0.5	0.0	0.0	0.0	34.0	0	0	c0t0d0s1				
0.0	0.0	0.0	0.1	0.0	0.0	0.6	35.0	0	0	fujixp:/export/home/user3				

示例4 显示从实例名称向描述性名称的转换

以下示例说明了使用 iostat 将特定实例名称转换成描述性名称。

example% iostat -xn sd1

extended device statistics														
r/s	w/s	kr/s	kw/s	wait	actv	wsvc_t	asvc_t	%w	%b	device				
0.0	0.0	0.0	0.0	0.0	0.0	0.0	0.0	0	0	c8t1d0				

示例5 显示特定磁盘的目标端口和控制器活动

以下示例中共有四个控制器，全都连接到同一个目标端口。

iostat -Y ssd22

extended device statistics														
device	r/s	w/s	kr/s	kw/s	wait	actv	svc_t	%w	%b					
ssd22	0.2	0.0	1.5	0.0	0.0	0.0	0.7	0	0					
ssd22.t2	0.2	0.0	1.5	0.0	0.0	0.0	0.0	0	0					
ssd22.t2.fp0	0.0	0.0	0.4	0.0	0.0	0.0	0.0	0	0					
ssd22.t2.fp1	0.0	0.0	0.4	0.0	0.0	0.0	0.0	0	0					
ssd22.t2.fp2	0.0	0.0	0.4	0.0	0.0	0.0	0.0	0	0					
ssd22.t2.fp3	0.0	0.0	0.4	0.0	0.0	0.0	0.0	0	0					

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	请参见下文。

调用为 Evolving（发展中）。用户可读的输出为 Unstable（不稳定）。

另请参见

[date\(1\)](#)、[sar\(1\)](#)、[sar\(1M\)](#)、[mpstat\(1M\)](#)、[vmstat\(1M\)](#)、[time\(2\)](#)、[attributes\(5\)](#)、[scsi_vhci\(7D\)](#)

附注

因为生成百分比数字中的舍入误差，CPU 使用总量可能与 100 略有差别。

I/O (r/s+w/s) 速度低于 0.5/秒时，svc_t 响应时间不是非常显著。此类情况下，无害峰值相当正常。

mpstat 实用程序报告相同的 wt、usr 和 sys 统计信息。有关更多信息，请参见 [mpstat\(1M\)](#)。

在 zone 中执行时，并且如果池工具处于活动状态，iostat(1M) 仅提供该 zone 绑定到的池的处理器集合中那些处理器的信息。

引用名	ipaddrsel – configure IPv6 default address selection
用法概要	<pre> /usr/sbin/ipaddrsel /usr/sbin/ipaddrsel -f file /usr/sbin/ipaddrsel -d </pre>
描述	Use the <code>ipaddrsel</code> utility to configure the IPv6 default address selection policy table. The policy table is a longest-matching-prefix lookup table that is used for IPv6 source address selection and for destination address ordering when resolving names to <code>AF_INET6</code> addresses. For a description of how the policy table is used for source address selection, see inet6(7P). For a description of how the policy table is used for destination address ordering, see getaddrinfo(3SOCKET). The unmodified policy table is valid for all typical IPv6 deployments. Modify the table only if a circumstance exists for which the default behavior of the IPv6 source address selection or destination address ordering mechanism is unsatisfactory. See the 示例 section for examples of such circumstances. You should carefully consider your addressing strategy before you change the table from the provided default. When the <code>ipaddrsel</code> command is issued without any arguments, the address selection policy currently in use is printed. The format of the output is compatible with the format of the configuration file that the <code>-f</code> option accepts. 注 – If the <code>usesrc</code> subcommand to ifconfig(1M) is applied to a particular physical interface, the selection policy specified by <code>usesrc</code> overrides the source address selection policies specified by <code>ipaddrsel</code>. This is true for packets that are locally generated and for applications that do not choose a non-zero source address using bind(3SOCKET).
The Configuration File	The configuration file that the <code>-f</code> option accepts can contain either comment lines or policy entries. Comment lines have a '#' character as the first non-blank character, and they are ignored by the <code>ipaddrsel</code> utility. Policy entry lines have the following format: <pre> prefix/prefix_length precedence label [# comment] </pre> The <i>prefix</i> must be an IPv6 prefix in a format consistent with inet(3SOCKET). The <i>prefix_length</i> is an integer ranging from 0 to 128. The IPv6 source address selection and destination address ordering algorithms determine the precedence or label of an address by doing a longest-prefix-match lookup using the prefixes in this table, much like next-hop determination for a destination is done by doing a longest-prefix-match lookup using an IP routing table. The precedence is a non-negative integer that represents how the destination address ordering mechanism will sort addresses returned from name lookups. In general, addresses with a higher precedence will be in front of addresses with a lower precedence. Other factors, such as destinations with undesirable source addresses can, however, override these precedence values.

The label is a string of at most fifteen characters, not including the NULL terminator. The label allows particular source address prefixes to be used with destination prefixes of the same label. Specifically, for a particular destination address, the IPv6 source address selection algorithm prefers source addresses whose label is equal that of the destination.

The label may be followed by an optional comment.

The file must contain a default policy entry, which is an entry with `: 0/0` as its *prefix* and *prefix_length*. This is to ensure that all possible addresses match a policy.

选项

The `ipaddrsel` utility supports the following options:

- f *file* Replace the address selection policy table with the policy specified in the *file*.
- d Revert the kernel's address selection policy table back to the default table. Invoking `ipaddrsel` in this way only changes the currently running kernel's policy table, and does not alter the configuration file `/etc/inet/ipaddrsel.conf`. To revert the configuration file back to its default settings, use `ipaddrsel -d`, then dump the contents of the table to the configuration file by redirecting the output of `ipaddrsel` to `/etc/inet/ipaddrsel.conf`.

example# `ipaddrsel -d`
example# `ipaddrsel > /etc/inet/ipaddrsel.conf`

示例

示例 1 The Default Policy in `/etc/inet/ipaddrsel.conf`

The following example is the default policy that is located in `/etc/inet/ipaddrsel.conf`:

# Prefix	Precedence Label
::1/128	50 Loopback
::/96	20 IPv4_Compatible
::ffff:0.0.0.0/96	10 IPv4
2002::/16	30 6to4
::/0	40 Default

示例 2 Assigning a Lower Precedence to Link-local and Site-local Addresses

By default, the destination address ordering rules sort addresses of smaller scope before those of larger scope. For example, if a name resolves to a global and a site-local address, the site local address would be ordered before the global address. An administrator can override this ordering rule by assigning a lower precedence to addresses of smaller scope, as the following table demonstrates.

# Prefix	Precedence Label
::1/128	50 Loopback
::/0	40 Default
2002::/16	30 6to4
fec0::/10	27 Site-Local
fe80::/10	23 Link-Local
::/96	20 IPv4_Compatible

示例 2 Assigning a Lower Precedence to Link-local and Site-local Addresses (续)

```
::ffff:0.0.0.0/96          10 IPv4
```

示例 3 Assigning Higher Precedence to IPv4 Destinations

By default, IPv6 addresses are ordered in front of IPv4 addresses in name lookups. ::ffff:0.0.0.0/96 has the lowest precedence in the default table. In the following example, IPv4 addresses are assigned higher precedence and are ordered in front of IPv6 destinations:

# Prefix	Precedence Label
::1/128	50 Loopback
::/0	40 Default
2002::/16	30 6to4
::/96	20 IPv4-Compatible
::ffff:0.0.0.0/96	60 IPv4

示例 4 Ensuring that a Particular Source Address is Used

This example ensures that a particular source address is used only when communicating with destinations in a particular network.

The following policy table assigns a label of 5 to a particular source address on the local system, 2001:1111:1111::1. The table assigns the same label to a network, 2001:2222:2222::/48. The result of this policy is that the 2001:1111:1111::1 source address will only be used when communicating with destinations contained in the 2001:2222:2222::/48 network. For this example, this network is the ClientNet, which could represent a particular client's network.

# Prefix	Precedence Label
::1/128	50 Loopback
2001:1111:1111::1/128	40 ClientNet
2001:2222:2222::/48	40 ClientNet
::/0	40 Default
2002::/16	30 6to4
::/96	20 IPv4-Compatible
::ffff:0.0.0.0/96	10 IPv4

This example assumes that the local system has one physical interface, and that all global prefixes are assigned to that physical interface.

退出状态

ipaddrsel returns the following exit values:

- 0 ipaddrsel successfully completed.
- >0 An error occurred. If a failure is encountered, the kernel's current policy table is unchanged.

文件	/etc/inet/ipaddrsel.conf	The file that contains the IPv6 default address selection policy to be installed at boot time. This file is loaded before any Internet services are started.

属性 See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 `nscd(1M)`, `inet(3SOCKET)`, `getaddrinfo(3SOCKET)`, `ipaddrsel.conf(4)`, `attributes(5)`, `inet6(7P)`

附注 The ipnodes cache kept by `nsd(1M)` contains addresses that are ordered using the destination address ordering algorithm, which is one of the reasons why `ipaddrsel` is called before `nsd` in the boot sequence. If `ipaddrsel` is used to change the address selection policy after `nsd` has started, you should invalidate the `nsd` ipnodes cache invalidated by invoking the following command:

```
example# /usr/sbin/nscd -i ipnodes
```

引用名

ipadm – 配置 Internet 协议网络接口和 TCP/IP 可调参数

用法概要

```

ipadm create-ip [-t] IP-interface
ipadm delete-ip IP-interface

ipadm create-vni [-t] VNI-interface
ipadm delete-vni VNI-interface

ipadm create-ipmp [-t] [-i interface, [...]] IPMP-interface
ipadm delete-ipmp [-f] IPMP-interface
ipadm add-ipmp [-t] -i interface, [...] [-i interface, [...]]
    IPMP-interface
ipadm remove-ipmp [-t] -i interface, [...] [-i interface, [...]]
    IPMP-interface

ipadm show-if [[-p] -o field, [...]] [interface]
ipadm disable-if -t interface
ipadm enable-if -t interface

ipadm set-ifprop [-t] -m protocol -p prop=value, [...] interface
ipadm reset-ifprop [-t] -m protocol -p prop interface
ipadm show-ifprop [[-c] -o field, [...]] [-p prop, ...]
    [-m protocol] [interface]

ipadm create-addr [-t] -T static [-d]
    -a {local|remote}=addr[/prefixlen],... addrobj
ipadm create-addr [-t] -T dhcp [-w seconds | forever]
    [-h hostname] addrobj
ipadm create-addr [-t] -T addrconf [-i {local|remote}=interface-id]
    [-p {stateful|stateless}={yes|no},...] addrobj

ipadm delete-addr [-r] addrobj
ipadm show-addr [[-p] -o field, [...]] [-d] [addrobj | interface/]
ipadm up-addr [-t] addrobj
ipadm down-addr [-t] addrobj
ipadm refresh-addr [-i] addrobj
ipadm disable-addr -t addrobj
ipadm enable-addr -t addrobj

ipadm set-addrprop [-t] -p prop=value, [...] addrobj
ipadm reset-addrprop [-t] -p prop=value, [...] addrobj
ipadm show-addrprop [[-c] -o field, [...]] [-p prop, [...]] [addrobj]

ipadm set-prop [-t] -p prop=value, [...] protocol
ipadm reset-prop [-t] -p prop protocol
ipadm show-prop [[-c] -o field, [...]] [-p prop, [...]] protocol | protocol]
```

描述

ipadm 命令提供了一组可用于执行下列操作的子命令：

管理接口：

- 创建和删除接口类为 ip、ipmp 和 vni 的接口
- 修改接口属性

- 显示接口配置

管理地址：

- 创建和删除地址
- 修改地址属性
- 显示地址配置

管理 TCP/IP 协议属性：

- 修改 TCP/IP 属性
- 显示 TCP/IP 属性

“子命令”部分后面的“操作数”部分中介绍了 ipadm 子命令的各种操作数。

需要的授权和特权 以下子命令需要 `solaris.network.interface.config` 授权和 `PRIV_SYS_IP_CONFIG` 特权。

<code>create-ip</code>	<code>create-addr</code>
<code>delete-ip</code>	<code>up-addr</code>
<code>create-vni</code>	<code>down-addr</code>
<code>delete-vni</code>	<code>refresh-addr</code>
<code>create-ipmp</code>	<code>disable-addr</code>
<code>delete-ipmp</code>	<code>enable-addr</code>
<code>add-ipmp</code>	<code>set-addrprop</code>
<code>remove-ipmp</code>	<code>reset-addrprop</code>
<code>disable-if</code>	<code>set-prop</code>
<code>enable-if</code>	<code>reset-prop</code>
<code>set-ifprop</code>	
<code>reset-ifprop</code>	

除了上面指定的授权和特权以外，ipadm 子命令 `create-ip`、`create-vni`、`create-ipmp` 和 `enable-if` 还需要 `PRIV_NET_RAWACCESS` 特权。

子命令 支持以下子命令：

`create-ip [-t] IP-interface`
创建处理 IPv4 和 IPv6 包的 IP 接口。IPv4 接口的地址将设置为 `0.0.0.0`，IPv6 接口的地址将设置为 `::`。缺省情况下，此子命令导致保存信息，因此，此接口将在下次重新引导时进行实例化。

在创建接口时针对 IPv4 和 IPv6 隐式启用接口。请参见下面的 `disable-if` 和 `enable-if` 子命令，以便禁用或启用接口。

请注意，`lo0` 是称为回送接口的特殊接口。它是一个虚拟 IP 接口，并且不与任何物理硬件关联。它是在系统上创建的第一批 IP 接口之一，并且 IPv4 地址为 `127.0.0.1`，IPv6 地址为 `::1`。

`-t`、`--temporary`
指定此操作为临时操作，因此不能保存。此操作仅影响活动配置。

delete-ip *IP-interface*

从活动配置中删除 IP 接口。在此接口上配置的所有地址都将删除。并且，所有与此接口相关的持久性信息都将从持久性数据存储中删除，为此，不会在重新引导时实例化 *interface*。要从活动配置禁用接口（而非删除此接口），请使用 **disable-if** 子命令。

create-vni [-t] *VNI-interface*

创建处理 IPv4 和 IPv6 包的 VNI 接口。IPv4 接口的地址将设置为 0.0.0.0，IPv6 接口的地址将设置为 ::。缺省情况下，此子命令导致保存信息，因此，此接口将在下次重新引导时进行实例化。

在创建接口时针对 IPv4 和 IPv6 隐式启用接口。请参见下面的 **disable-if** 和 **enable-if** 子命令，以便禁用或启用接口。

请注意，vni 是一个虚拟接口，并且不与任何硬件关联，因此它是一个特殊接口。请参见 [vni\(7d\)](#)。

-t、--temporary

指定此操作为临时操作，因此不能保存。此操作仅影响活动配置。

delete-vni *VNI-interface*

从活动配置中删除 VNI 接口。在此接口上配置的所有地址都将删除。并且，所有与此 IP 接口相关的持久性信息都将从持久性数据存储中删除，为此，不会在重新引导时实例化接口。要从活动配置禁用此接口（而非删除此接口），请使用 **disable-if** 子命令。

create-ipmp [-t] [-i *interface*,[...]] *IPMP-interface*

创建处理 IPv4 和 IPv6 包的 IPMP 接口。IPv4 接口的地址将设置为 0.0.0.0，IPv6 接口的地址将设置为 ::。缺省情况下，此子命令导致保存信息，因此，此接口将在下次重新引导时进行实例化。

在创建接口时针对 IPv4 和 IPv6 隐式启用接口。请参见下面的 **disable-if** 和 **enable-if** 子命令，以便禁用或启用 IPMP 接口。

-t、--temporary

指定此操作为临时操作，因此不能保存。此操作仅影响活动配置。

-i、--interface *interface*,[...]

以逗号分隔的、要作为底层接口添加到 IPMP 接口的接口列表。指定接口必须存在于活动配置中才能成功添加到 IPMP 组，并且不能存在于任何其他 IPMP 组中。允许使用多个 -i 选项。如果已创建 IPMP 接口，但添加给定的底层接口失败，此命令将以部分成功的形式返回。

delete-ipmp [-f] *IPMP-interface*

从活动配置中删除 IPMP 接口。在此接口上配置的所有地址都将删除。除非指定了 -f 选项，否则如果 IPMP 接口具有任何底层接口，此命令将失败。并且，所有与此 IPMP 接口相关的持久性信息都将从持久性数据存储中删除，为此，不会在重新引导时实例化接口。要从活动配置仅禁用此接口（而非删除此接口），请使用 **disable-if** 子命令。

-f、--force

如果 IPMP 接口具有任何底层接口，指定此选项将首先从组中删除所有底层接口，然后再删除此 IPMP 接口。

add-ipmp [-t] -i *interface*,[...] [-i *interface*,[...]...] *IPMP-interface*

向给定 IPMP 接口添加一个或多个底层 IP 接口。

-t、--temporary

指定此操作为临时操作，因此不能保存。此操作仅影响活动配置。

-i、--interface *interface*,[...]

以逗号分隔的、要作为底层接口添加到 IPMP 接口的接口列表。指定接口必须存在于活动配置中才能成功添加到 IPMP 组，并且不能存在于任何其他 IPMP 组中。如果已至少添加一个接口，并且添加其余接口失败，此命令将以部分成功的形式返回。允许使用多个 -i 选项。

remove-ipmp [-t] -i *interface*,[...] [-i *interface*,[...]...] *IPMP-interface*

从 IPMP 接口删除一个或多个底层 IP 接口。

-t、--temporary

指定此操作为临时操作，因此不能保存。此操作仅影响活动配置。

-i、--interface *interface*,[...]

要从 IPMP 接口中删除的底层接口的逗号分隔列表。指定接口必须已成为给定 IPMP 组的底层接口。允许使用多个 -i 选项。如果已至少删除一个接口，并且删除其余接口失败，此命令将以部分成功的形式返回。

show-if [[-p] -o *field*,...] [*interface*]

显示系统中配置的所有网络接口（包括仅位于持久性配置中的网络接口）或指定网络接口的网络接口配置信息。

-o *field*,...、--output *field*,[...]

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。*field* 名称必须是下面列出的某个字段，或者是特殊值 **all** 以便显示所有字段。对于每个网络接口，可显示下列字段：

IFNAME

IP 接口的名称。

CLASS

表示下列项之一：

ip

在底层数据链路上方检测到的接口。

ipmp

在一个或多个底层 IP 接口上方创建的 IPMP 接口。

loopback

回送接口。

vni

虚拟 IP 接口。请参见 [vni\(7d\)](#)。

STATE

表示显示的接口的下列项之一：

ok (正常)

指示已分配接口所需的资源。对于某些接口，此标志还指示链路已启动。

offline (脱机)

接口处于脱机状态，因此不能发送或接收 IP 数据通信。请参见 [if_mpadm\(1M\)](#)。

failed (失败)

表示数据链路已关闭。如果接口包含在 IPMP 组中，则还可能表示此接口已失败（也即，设置了 IFF_FAILED）。失败的接口将不会用于发送或接收 IP 数据通信。如果对 IPMP 组中的某一物理 IP 接口设置了此标志，IP 数据通信将继续经由 IPMP 组中的其他可用物理 IP 接口传送。如果对 IPMP IP 接口设置了此标志，则整个组将失败，无法通过该组中的任何接口发送或接收数据通信。请参见 [in.ndpd\(1M\)](#)。

down (关机)

表示已通过管理方式关闭接口，从而防止通过此接口发送或接收任何 IP 包。

disabled (禁用)

表示已使用 `disable-if` 子命令从活动配置中禁用接口。

ACTIVE

此值为 `yes` 或 `no`，具体取决于系统是否正在使用此 IP 接口传输 IP 数据通信。

CURRENT

对于活动配置中的接口对象，此选项表示下列任何标志。

b

接口支持广播

m

接口支持多播

p

接口是点对点链路

v

接口是虚拟接口（例如，[vni\(7d\)](#), `loopback`），也就是说，物理接口没有底层硬件。

s

IPMP 接口以管理方式标记为备用。请参见 [in.ndpd\(1M\)](#)。

l

接口是 IPMP 接口的底层接口。请参见 [in.ndpd\(1M\)](#)。

- i
底层接口处于非活动状态。请参见 [in.ndpd\(1M\)](#)。
- V
接口是 VRRP 接口
- a
VRRP 接口处于接受模式 (~IFF_NOACCEPT)
- Z
已通过管理方式强制实施接口的 IP 地址的第 3 层保护。
- 4
接口可处理 IPv4 包
- 6
接口可处理 IPv6 包

请注意，b 和 p 互斥。

PERSISTENT

指定在重新引导时实例化或使用 `enable-if` 子命令重新启用接口对象时将应用的配置。此值可以为下列任何项或所有项：s、l、4 和 6（请参见上文）。缺省情况下，不会显示此字段，此字段仅在指定了 `all` 或 `persistent` 以及 `-o` 时显示。

OVER

创建 IPMP 接口所基于的底层接口。此字段不适用于其他接口类。

-p、--parsable

使用稳定的可供计算机解析的格式显示。`-o` 选项需要与此选项一起使用。请参见下面的“可解析的输出格式”。

`disable-if -t interface`

通过从活动配置中删除指定 *interface* 来禁用它。在此接口上配置的所有地址都将禁用。如果在一开始时就以持久方式创建了接口对象，持久性配置保持不变。要重新启用此接口，用户应使用 `enable-if`。

-t、--temporary

指定禁用是临时操作，并且更改仅应用于活动配置。

`enable-if -t interface`

通过从持久性存储读取配置来启用给定接口。将应用所有持久性接口属性（如果有），并且将启用给定接口上的所有持久性地址（如果有）。

-t、--temporary

指定启用是临时操作，并且更改仅应用于活动配置。

set-ifprop [-t] -m protocol -p prop=value[...] interface

将接口属性修改为用户指定的值。如果此属性采用多个值，则应使用逗号作为分隔符来指定这些值。一次只能指定一个属性。使用 **show-ifprop** 子命令可以检索接口上支持的属性以及属性的可能值。一次只能修改一个属性。

-t、**--temporary**

指定更改是临时操作，并且更改仅应用于活动配置。

-m protocol、**--module protocol**

确定应对 IPv4 还是 IPv6 包应用属性。

-p prop=value[...]、**-prop prop=value[...]**

要设置为指定值的属性。

reset-ifprop [-t] -m protocol -p prop interface

将指定接口的属性重置为其缺省值。如果未使用 **-t**，则将删除属性的所有保存的值。一次只能修改一个属性。

-t、**--temporary**

指定重置是临时操作，并且更改仅应用于活动配置。

-m protocol、**--module protocol**

确定重置的属性将影响 IPv4 还是 IPv6 包。

-p prop、**-prop prop**

要设置为指定值的属性。

show-ifprop [[-c] -o field[...]] [-p prop,...] [-m protocol] [interface]

显示创建的所有接口或指定接口的一个或多个属性的当前和持久性值。通过向 **-p** 选项提供以逗号分隔的属性名称，一次可以检索多个相关属性。如果未指定 **-p** 选项，则显示所有可用的接口属性。

-o field[...]、**--output field[...]**

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。*field* 名称必须是下面列出的某个字段，或者是特殊值 **all** 以便显示所有字段。对于每个接口，可显示下列字段：

IFNAME

接口的名称。

PROPERTY

属性的名称。

PROTO

属性所属的协议的名称。当前支持的协议包括 IPv4 和 IPv6。

PERM

属性的读/写权限。显示的值将为 **r**（只读）、**w**（只写）或 **rw**（读写）。

CURRENT

属性的当前值。对于已禁用的接口，由于未设置值，因此此选项显示为 **--**。

PERSISTENT

属性的持久性值。持久性值是将在重新引导时重新应用的值。

DEFAULT

属性的缺省值。如果属性没有缺省值，则显示 --。

POSSIBLE

属性可具有的值的分号分隔列表。如果这些值在某个数字范围中，则可以显示简写形式 *min - max*。如果可能的值未知，则显示 ?，或者如果可能的值无限制，则显示 --。

-c、--parsable

使用稳定的可供计算机解析的格式显示。-o 选项需要与此选项一起使用。请参见下面的“可解析的输出格式”。

-p *prop*,...、--prop=*prop*

要显示的属性的逗号分隔列表。请参见子命令说明后面的接口属性部分。

-m *protocol*、--module *protocol*

显示与给定协议匹配的属性。有效值为 *ipv4* 和 *ipv6*。

有关接口属性的支持列表，请参见下面的“接口属性”。

create-addr [-t] -T static [-d] -a {local | remote}=addr[/prefixlen],... addrobj

在 *addrobj* 指定的接口上创建静态 IPv4 或 IPv6 地址。创建 *addrobj* 时所在的接口必须已存在。创建的静态地址将通过 *addrobj* 标识。

缺省情况下，配置的地址将标记为 *up*，因此可将其用作传出和传入包的源或目标。

所有地址对象均在创建时启用。有关如何禁用或启用地址对象的说明，请参见 *disable-addr* 和 *enable-addr* 子命令。

不能在临时对象中执行持久性操作。也就是说，如果接口是临时创建的，用户不能以持久方式创建地址对象。

如果 *addrobj* 名称中指定的接口为 IPMP 接口，则会在 IPMP 接口中创建一个静态数据地址。如果 *addrobj* 名称中指定的接口为 IPMP 组的底层接口，则会在底层接口中创建一个静态测试地址。

-t、--temporary

指定配置的地址是临时地址，并且更改仅应用于活动配置。

-d、--down

指定应将配置的地址标记为 *down*，也就是说，该地址将不会用作 IP 包的源或目标。

-a {local | remote}=addr[/prefixlen],...**--address {local | remote}=addr[/prefixlen],...**

addr 表示与本地或远程端点对应的文本 IP 地址或主机名（适用于点对点接口）。

如果指定了主机名，则使用 `/etc/hosts` 中的条目以唯一方式获取其数值。如果此文件中未定义数字 IP 地址，则使用在 `nsswitch.conf(4)` 中为 `hosts` 或 `ipnodes` 指定的解析器顺序以唯一方式获取此数值。如果给定主机名存在多个条目，则将生成错误。由于先创建 IP 地址，然后才会在引导进程中将命名服务联机，因此应将使用的所有主机名包含在 `/etc/hosts` 中，这一点非常重要。

如果未在命令行中明确指定 *prefixlen*，则按下面列出的顺序来执行搜索，以便获取地址的网络掩码：

1. 使用在 `nsswitch.conf(4)` 中为 `netmasks` 指定的顺序
2. 使用 RFC 791 中定义的一类网语义解释 IPv4 地址，使用 RFC 4291 中的定义解释 IPv6 地址。

对于点对点接口，除本地端点的地址以外，还必须指定远程端点的地址（例如，`-a local=laddr,remote=raddr`）。如果指定了远程端点的 *prefixlen*，则将返回错误。

请注意，如果接口仅需要本地地址，请使用 `-a` 选项按如下方法直接指定此地址：`-a addr[/prefixlen]`。此地址将自动视为本地地址。

create-addr [-t] -T dhcp [-w seconds | forever] -h hostname addrobj

在 *addrobj* 指定的接口上创建受 DHCP 控制的 IPv4 地址。创建的 IPv4 地址将通过 *addrobj* 标识。如果 *addrobj* 包含底层接口，此命令将创建一个测试地址；如果包含 IPMP 接口，则会创建一个数据地址。

所有地址对象均在创建时启用。有关如何禁用和启用地址对象的说明，请参见 `disable-addr` 和 `enable-addr` 子命令。

不能在临时对象中执行持久性操作。也就是说，如果接口是临时创建的，用户不能以持久方式创建地址对象。

如果 *addrobj* 名称中指定的接口为 IPMP 接口，则会在 IPMP 接口中创建通过 DHCP 获取的地址以作为数据地址。

-h hostname

指定客户机要让 DHCP 服务器将客户机的租用 IPv4 地址映射到的主机名。不保证 DHCP 服务器能够完成主机名请求。

-t, --temporary

指定配置的地址是临时地址，并且更改仅应用于活动配置。

-w seconds | forever, --wait seconds | forever

指定等待操作完成的时间（以秒为单位）。如果未指定等待时间间隔，此操作无法立即完成，缺省情况下，`ipadm` 将等待 120 秒钟以便完成请求的操作。请注意，缺省等待时间将在将来的发行版中发生更改。也可以使用具有明确含义的符号值 `forever`。

```
create-addr [-t] -T addrconf [-i {local | remote}=interface-id] [-p {stateful | stateless}={yes | no},...] addrobj
```

在 *addrobj* 指定的接口上创建自动配置的 IPv6 地址。创建的 IPv6 地址将通过 *addrobj* 标识。

系统使用缺省接口 ID（对于介质类型的以太网，接口 ID 即为此接口的 MAC 地址）生成自动配置的地址。使用 *-i* 选项可覆盖此行为。

缺省情况下：

- IPv6 地址将根据路由器声明的前缀自动配置，如 RFC 4862 中所述，并且...
- 将使用 DHCPv6 服务器提供的 IPv6 地址在指定接口上自动配置 IPv6 地址，如 RFC 3315 中所述。（也就是说，*-p stateful=yes, stateless=yes* 为缺省选项。）

所有地址对象均在创建时启用。有关如何禁用和启用地址对象的说明，请参见 *disable-addr* 和 *enable-addr* 子命令。

不能在临时对象中执行持久性操作。也就是说，如果接口是临时创建的，用户不能以持久方式创建地址对象。

如果 *addrobj* 名称中指定的接口为 IPMP 接口，则会在 IPMP 接口上创建通过 IPv6 自动配置获取的地址以作为数据地址。

-t、*--temporary*

指定配置的地址是临时地址，并且更改仅应用于活动配置。

-i {local | remote}=interface-id、*--interface-id {local | remote}=interface-id*

指定用于生成自动配置的地址的接口 ID。

对于点对点接口，可指定远程端点的接口 ID（例如，*-i local=lid,remote=rid*）。

请注意，如果接口仅需要本地接口 ID，请使用 *-i* 选项按如下方式直接指定此 ID：*-i lid*。此接口 ID 将自动视为本地接口 ID。

-p {stateful | stateless}={yes | no},...

--prop {stateful | stateless}={yes | no},...

指定是否应启用自动配置的有状态和/或无状态方法。

如果指定 *-p stateful=no*，则不会执行基于 DHCPv6 指定的 IPv6 地址的有状态自动配置。

如果指定 *-p stateless=no*，则不会执行基于路由器声明的前缀的无状态自动配置。

如果指定 *-p stateful=no, stateless=no*，则不会执行两种自动配置方法。

通过 *-T addrconf* 选项，将在缺省情况下使用 *-p stateful=yes, stateless=yes*。

delete-addr [-r] *addrobj*

在 *addrobj* 指定的接口中删除 *addrobj* 标识的所有地址。此命令还会从持久性数据存储中删除这些地址；因此，这些地址不会在重新引导时实例化。

如果地址对象是受 DHCP 控制的地址，**delete-addr** 将从系统中删除此地址，而不会通知 DHCP 服务器，并记录当前租用以备后用。

-r、--release

如果 *addrobj* 是受 DHCP 控制的地址，此选项将导致通过通知服务器在接口上放弃受 DHCP 控制的 IP 地址，并放弃当前租用。

show-addr [[-p] -o *field*[,...]] [-d] [*addrobj* | *interface*]

显示给定 *addrobj* 或指定接口上配置的所有地址对象（包括仅位于持久性配置中的地址对象）的地址信息。

-p、--parsable

使用稳定的可供计算机解析的格式显示。**-o** 选项需要与此选项一起使用。请参见下面的“可解析的输出格式”。

-o *field*[,...]、--output *field*[,...]

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。*field* 名称必须是下面列出的某个字段，或者是特殊值 **all** 以便显示所有字段。对于每个接口，可显示下列字段：

ADDROBJ

地址对象的名称。

TYPE

地址对象的类型。此类型将为下列值之一：**from-gz**、**static**、**dhcp** 或 **addrconf**。**static**、**dhcp** 和 **addrconf** 类型对应于 **create-addr** 的 **-T** 选项所指定的地址对象的类型。**from-gz** 类型将仅在非全局区域中显示，并且表示该地址是根据 **allowed-address** 属性配置的，此属性是从全局区域为非全局专用 IP 区域配置的。

STATE

地址对象的状态。此字段仅在指定了 **all** 和 **-o** 时显示。此字段表示下列值之一：

disabled (禁用)

地址不包含在活动配置中（请参见 **disable-addr** 和 **disable-if**）。

down (关机)

地址通过管理方式关闭（请参见 **down-addr**）。

duplicate (重复)

重复地址检测 (duplicate address detection, DAD) 发现地址与其他系统的 IP 地址冲突，在解决冲突之前，无法使用此地址。系统将定期重新运行 DAD，确定是否已解决冲突。或者，也可以使用 **refresh-addr** 立即重新运行 DAD。

inaccessible (无法访问)

由于配置地址所在的 IP 接口已失败，因此无法使用此地址。

ok (正常)

地址已启用且正常运行。系统将接受发往此地址的 IP 包，并且将根据配置的 IP 源地址选择策略使用此地址发出 IP 包。

tentative (暂定)

地址当前正在进行重复地址检测（例如，作为 `up-addr` 或 `refresh-addr` 的一部分）。

CURRENT

对于活动配置中的地址对象，此选项表示下列任何标志。缺省情况下，不会显示此字段，此字段仅在指定了 `all` 或 `current` 以及 `-o` 时显示。

d (deprecated)

除非接口上没有其他可用地址，或者应用程序已显式绑定到此地址，否则无法用作传出包的源地址。

p (private)

路由守护进程未声明的地址。

t (temporary)

在 RFC 3041 中定义的临时 IPv6 地址。

U (up)

标记地址，以便用作传出/传入包的源/目标。

u (unnumbered)

地址与系统中的某个其他链路的本地地址匹配。

PERSISTENT

指定在重新引导时实例化或使用 `enable-addr` 子命令重新启用地址对象时将应用的配置。此值可以为下列任何项或所有项：`U`、`p` 和 `d`（请参见上文）。

ADDR

数字 IPv4 或 IPv6 地址。对于点对点接口，显示两个端点的地址 (`laddr-->raddr`)。对于类型为 `dhcp` 的地址对象，如果此地址对象的状态为 `disabled`，或者如果此地址为 `0.0.0.0`（适用于 IPv4）或 `::`（适用于 IPv6），则会显示一个问号 (?)。

CID-TYPE

dhcpagent(1M) 使用的客户机 ID 的类型（如果使用 DHCP 获取此地址）。对于 IPv4，此选项显示在构造 RFC 4361 客户机 ID 时使用的 DUID 的类型。此类型为 `DUID-LLT`、`DUID-EN`、`DUID-LL`、`other` 或 `default` 之一。缺省输出中不会显示此字段。通过使用 `-d` 或者使用 `cid-type` 或 `all` 以及 `-o`，可以显示此字段。

DUID-LLT

类型 1 RFC 3315 DUID 用于构造 CID-VALUE（例如，`1,1,63463777,0a:0b:0c:0d:0e:0f`）。有关详细信息，请参见 RFC。

DUID-EN

类型 2 RFC 3315 DUID 用于构造 CID-VALUE (例如, 1,1,63463777,0a:0b:0c:0d:0e:0f)。有关详细信息, 请参见 RFC。

DUID-LL

类型 3 RFC 3315 DUID 用于构造 CID-VALUE (例如, 1,1,63463777,0a:0b:0c:0d:0e:0f)。有关详细信息, 请参见 RFC。

other

类型包含在 {0,4-65535} 中的 RFC 3315 DUID 用于派生客户机 ID (例如, 4,0x734633), 或者 CID-VALUE 是不符合 RFC 3315 的原始客户机 ID (例如, Sun,0xab3146)。

default

表示 RFC 3315 DUID 未用于构造客户机 ID, 而是根据 RFC 2132 使用接口的 MAC 地址派生客户机 ID。CID-VALUE 将包含后跟有 MAC 地址十六进制字符串的字符串 0x01。这仅适用于 IPv4。

CID-VALUE

dhcpage(1M) 使用的客户机 ID 的值 (如果使用 DHCP 获取此地址)。使用的格式遵循文件 /etc/default/dhcpage 中的配置参数 CLIENT_ID 的格式。请参见 dhcpage(1M) 中 CLIENT_ID 的说明。如果 CID-TYPE 是缺省值, 则 CID-VALUE 包含根据 RFC 2132 构造的旧 CLIENT-ID。缺省输出中不会显示此字段。通过使用 -d 或者使用 cid-type 或 all 以及 -o 选项, 可以显示此字段。

BEGIN

对于使用 DHCP 获取的地址而言, 这是租用的开始时间 (如果可用)。此时间按 LC_TIME 本地环境变量控制的格式显示。对于并非使用 DHCP 配置的地 址或没有租用的 DHCP 地址, 将显示 -- (双连字符)。缺省输出中不会显示此字段。通过使用 -d 或者使用 cid-type 或 all 以及 -o 选项, 可以显示此字段。

EXPIRE

对于使用 DHCP 获取的地址而言, 这是租用过期的时间 (如果可用)。此时间按 LC_TIME 本地环境变量控制的格式显示。对于并非使用 DHCP 配置的地 址或没有租用的 DHCP 地址, 将显示 -- (双连字符)。缺省输出中不会显示此字段。通过使用 -d 或者使用 cid-type 或 all 以及 -o 选项, 可以显示此字段。

RENEW

对于使用 DHCP 获取的地址而言, 这是上次续订租用的时间。此时间按 LC_TIME 本地环境变量控制的格式显示。对于并非使用 DHCP 配置的地 址或没有租用的 DHCP 地址, 将显示 -- (双连字符)。缺省输出中不会显示此字段。通过使用 -d 或者使用 cid-type 或 all 以及 -o 选项, 可以显示此字段。

-d、--dhcp

显示使用 DHCP 获取的地址的 dhcp 状态字段。显示的字段包括 ADDR_OBJ、STATE、ADDR、CID-TYPE、CID-VALUE、BEGIN、EXPIRE 和 RENEW。此选

项仅显示便于阅读的输出，并且无法与 `-p` 配合使用。

注 - 在某些情况下，地址对象的名称以下条 () 开头，如在 `e1000g1/_a` 或 `e1000g1/_b` 中。对于使用 `ipadm` 编程库的应用程序，自动生成这些名称。您还将看到地址对象名称中含有问号 (?) 的地址。这表示这些地址是在 `ipadm` 库外部创建的，因此 `ipadm` 不知道这些地址。

`down-addr [-t] addrobj`

通过 `addrobj` 标识的地址标记为已关闭，因此无法将其用作传出/传入包的源/目标。如果在调用 `down-addr` 之前已将地址对象标记为关闭，此命令不起作用。如果地址对象的类型为 `addrconf`，此命令将返回错误。

`-t`、`--temporary`

指定配置的地址是临时地址，并且更改仅应用于活动配置。如果地址对象类型为 `dhcp`，此选项则是强制性的。

`up-addr [-t] addrobj`

通过 `addrobj` 标识的地址标记为已启用，因此可以将其用作传出/传入包的源/目标。如果系统因地址对象是重复地址而将地址对象标记为已关闭，或者如果在调用 `up-addr` 之前将此地址标记为已启用，此命令不起作用。如果地址对象的类型为 `addrconf`，此命令将返回错误。

`-t`、`--temporary`

指定配置的地址是临时地址，并且更改仅应用于活动配置。如果地址对象类型为 `dhcp`，此选项则是强制性的。

`refresh-addr [-i] addrobj`

如果 `addrobj` 的类型为 `static`，则会在地址对象标识的地址上重新启动 DAD (Duplicate Address Detection，重复地址检测) (如果需要)。

如果 `addrobj` 的类型为 `dhcp`，DHCP 客户机守护进程将延长在此地址上获取的租用持续时间。

如果 `addrobj` 的类型为 `addrconf`，此命令将返回错误。

`-i`、`--inform`

对于指定 IP 地址，从 DHCP 获取网络配置参数，而不在其中获取租用。当通过 DHCP 以外的机制获取 IP 地址时，此选项非常有用。

`disable-addr -t addrobj`

通过从活动配置中删除地址来禁用它。如果在一开始时就以持久方式创建了地址对象，持久性配置保持不变。要重新启用此 `addrobj`，用户应使用 `enable-addr`。

`-t`、`--temporary`

指定禁用是临时操作，并且更改仅应用于活动配置。

enable-addr -t *addrobj*

通过从持久性存储读取配置来启用给定 *addrobj*。所有持久性地址属性都适用于地址对象。此子命令要求存在启用地址对象所在的接口。如果活动配置中缺少此接口本身，且此接口位于持久性存储中，也就是说，如果禁用此接口，那么用户必须在调用 **enable-addr** 之前运行 **enable-if**。

-t、**--temporary**

指定启用是临时操作，并且更改仅应用于活动配置。

set-addrprop [-t] -p prop=*value*[...] *addrobj*

在指定 *addrobj* 上设置属性的值。如果 *addrobj* 映射到多个地址，属性更改将应用于 *addrobj* 引用的所有地址。一次只能指定一个属性。使用 **show-addrprop** 子命令可以检索 *addrobj* 上支持的属性以及属性的可能值。如果 *addrobj* 的类型为 **addrconf**，此命令将返回错误。

-t、**--temporary**

指定更改是临时操作，并且更改仅应用于活动配置。

-p prop=*value*[...]、--prop prop=*value*[...]

要设置为指定值的属性。

reset-addrprop [-t] -p prop *addrobj*

将给定地址属性重置为其缺省值。如果未使用 **-t**，则将删除属性的所有持久性值。一次只能修改一个属性。如果 *addrobj* 的类型为 **addrconf**，此命令将返回错误。

-t、**--temporary**

指定重置是临时操作，并且更改仅应用于活动配置。

-p prop、--prop prop

要重置的属性。

show-addrprop [[-c] -o field[...]] [-p prop,...] [*addrobj*]

显示配置的所有地址对象或指定 *addrobj* 的一个或多个属性的当前和持久性值。通过向 **-p** 选项提供以逗号分隔的属性名称，一次可以检索多个相关属性。如果未指定 **-p** 选项，则显示所有可用属性。如果 *addrobj* 的类型为 **addrconf**，此命令将返回错误。

-o field[...]、**--output field[...]**

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。*field* 名称必须是下面列出的某个字段，或者是特殊值 **all** 以便显示所有字段。对于每个 *addrobj*，可显示下列字段：

ADDROBJ

地址对象的名称。

PROPERTY

属性的名称。

PERM

属性的读/写权限。显示的值将为 **r**（只读）、**w**（只写）或 **rw**（读写）。

CURRENT

属性的当前值。对于禁用的地址，由于未设置该值，因此该值显示为双连字符 (--)。

PERSISTENT

属性的持久性值。持久性值是将在重新引导时重新应用的值。

DEFAULT

属性的缺省值。如果属性没有缺省值，则显示双连字符 (--)。

POSSIBLE

属性可具有的值的分隔列表。如果这些值在某个数字范围中，则可以显示简写形式 *min - max*。如果可能的值未知，则显示问号 (?)，或者如果可能的值无限制，则显示双连字符 (--)。

-c、--parsable

使用稳定的可供计算机解析的格式显示。**-o** 选项需要与此选项一起使用。请参见下面的“可解析的输出格式”。

-p prop,...、--prop=prop

要显示的属性的逗号分隔列表。请参见子命令说明后面的地址对象属性部分。

set-prop [-t] -p prop[+|-]=value[,...] protocol

将协议属性的值修改为指定的 *value*。如果此属性采用多个值，则应使用逗号作为分隔符来指定这些值。一次只能指定一个属性。缺省情况下，此值为持久性值，并且将在重新引导时重新应用。使用 **show-prop** 子命令可以检索协议上支持的属性以及属性的可能值。

支持下列协议：**ip**、**ipv4**、**ipv6**、**icmp**、**tcp**、**udp** 和 **sctp**。

请注意，对于某些属性，可以全局设置属性值，也可以根据接口设置属性值。使用 **set-ifprop** 子命令可以设置基于接口的值。此种情况下，如果管理员选择自定义属性基于接口的值，以便与全局值区分开来，基于接口的值将覆盖此接口的全局设置。

-t、--temporary

指定属性更改是临时操作，并且更改仅应用于活动配置。

-p prop[+|-]=value[,...]、--prop prop[+|-]=value[,...]

要设置为指定值的属性。此属性还提供下列“限定符”，以便执行添加、删除和赋值操作。

+

向值的当前列表添加给定值。

-

从值的当前列表删除给定值。

=

执行新的赋值，并删除所有当前值。

有关如何使用限定符的详细信息，请参见“示例”部分。

reset-prop [-t] -p *prop* *protocol*

将指定协议的属性重置为此属性的缺省值。如果未使用 **-t**，则将删除属性的所有持久性值。一次只能修改一个属性。

-t、**--temporary**

指定重置是临时操作，并且更改仅应用于活动配置。

-p *prop*、**--prop *prop***

要重置的属性。

show-prop [[-c] -o *field*[...]] [-p *prop*[...] *protocol* | *protocol*]

显示支持的所有协议或指定协议的一个或多个属性的当前和持久性值。通过向 **-p** 选项提供以逗号分隔的属性名称，一次可以检索多个相关属性。如果未指定 **-p** 选项，则显示所有可用属性。

-o *field*[...]、**--output *field*[...]**

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。*field* 名称必须是下面列出的某个字段，或者是特殊值 **all** 以便显示所有字段。对于每个 *protocol*，可显示下列字段：

PROTO

协议的名称。

PROPERTY

属性的名称。

PERM

属性的读/写权限。显示的值将为 **r**（只读）、**w**（只写）或 **rw**（读写）。

CURRENT

属性的当前值。对于禁用的地址，由于未设置该值，因此该值显示为双连字符 (--)。如果值未知，则显示为一个问号 (?)。如果属性的当前值不在列出的 POSSIBLE 值集中，则显示关键字 **custom**。

PERSISTENT

属性的持久性值。持久性值是在重新引导时重新应用的值。

DEFAULT

属性的缺省值。如果属性没有缺省值，则显示双连字符 (--)。

POSSIBLE

要与 **set-prop** 子命令配合使用的属性设置的值的逗号分隔列表。如果这些值在某个数字范围中，则可以显示简写形式 *min - max*。如果可能的值未知，则显示问号 (?)，或者如果可能的值无限制，则显示双连字符 (--)。

-c、**--parsable**

使用稳定的可供计算机解析的格式显示。**-o** 选项需要与此选项一起使用。请参见下面的“可解析的输出格式”。

`-p prop,... \ --prop=prop`

要显示的属性的逗号分隔列表。请参见子命令说明后面的协议属性部分。

有关各个协议的属性的支持列表，请参见下面的“协议属性”。

可解析的输出格式 `ipadm "show"` 子命令具有一个以计算机可解析的格式显示输出的 `-o` 选项。输出格式是一行或多行以冒号 (:) 分隔的字段。显示的字段特定于使用的子命令，并且在给定子命令的 `-o` 选项的条目下列出。输出仅包括通过 `-o` 选项请求的字段，这些字段按请求的顺序排列。请注意，显示给定子命令的所有字段的 `-o all` 选项不能与可解析的输出选项配合使用。

当请求多个字段时，将在输出之前使用反斜杠 (\) 对任何文本冒号字符进行转义。类似地，也会对文本反斜杠字符进行转义 (\\)。通过使用带有环境变量（按 `IFS=:` 设置）的 shell `read(1)` 函数，可以解析此转义格式。请注意，当仅请求一个字段时，不会进行转义。

协议属性 支持下列协议属性：

注– 存在特定于协议的、以 “_”（下划线）开头的协议属性。这些属性可能会发生更改或者被删除，缺省情况下不显示在 `ipadm show-prop` 输出中。有关详细信息，请参见《Oracle Solaris Tunable Parameters Reference Manual》。

cong_default (TCP, SCTP)

指定创建新的连接时由协议使用的缺省拥塞控制算法。在该连接的生命周期内，应用程序可以在以后选择不同的算法。只有已启用的算法才可以设置为缺省值（请参见 `cong_enabled`）。

cong_enabled (TCP, SCTP)

该选项可以用于启用或禁用拥塞控制算法。缺省情况下，系统上安装的所有算法都处于启用状态。已禁用的算法不能设置为缺省值（请参见 `cong_default`）或由应用程序使用。

可以使用 `set-prop` 子命令和修饰符 `+` 与 `-` 来添加或删除算法。

ecn (TCP)

显式拥塞控制（有关详细信息，请参见 RFC 3168）。可能的值与上面的值相同，即：`never`、`passive` 和 `active`。

extra_priv_ports (TCP, SCTP, UDP)

此选项定义 1-1023 范围以外的其他特权端口。尝试绑定此处列出的端口的任何程序都必须作为根运行。这会防止普通用户在特定端口上启动服务器进程。

使用 `set-prop` 子命令以及 `+`、`-` 和 `=` 修饰符可以对这些端口进行添加、删除或赋值操作。有关用法，请参见下文的“示例”部分。

forwarding (IPv4)、forwarding (IPv6)

启用/禁用全局 IPv4 或 IPv6 转发。配置的所有接口都将启动/停止转发包。各接口可使用 `set-ifprop` 覆盖全局选项。

hostmodel (IPv4)、hostmodel (IPv6)

控制多宿主系统中的 IP 包的发送/接收行为。**hostmodel** 的值可设置为与 RFC 1122 的等效终端系统模型定义对应的 **strong** 或 **weak**。此外，还支持第三个值 **src-priority**。在 **src-priority hostmodel** 方案中，只要配置了包的目标 IP 地址并在某个主机接口上将其标记为 **up**，则会在任何接口上接受此包。当传送包时，如果此包中的 IP 目标有多个路由，系统将首选这样的路由：包中的 IP 源地址在外发接口上配置。如果没有此类路由，系统将退而选择“最佳”路由，就像在弱 ES 情况下一样。

max_buf (TCP, SCTP, UDP, ICMP)

发送或接收套接字缓冲区的最大大小。该属性的当前值限制 **recv_buf** 和 **send_buf** 的最大值。

recv_buf (TCP, SCTP, UDP, ICMP)**send_buf (TCP, SCTP, UDP, ICMP)**

修改指定协议的接收或发送缓冲区大小。这些属性的最大值均受 **max_buf** 属性当前值的限制。

sack (TCP)

通过选择性确认 (Selective acknowledgment, SACK)，接收者可以选择性地确认失序数据，选择性确认旨在提高通过有损链路传输数据的性能。有关 SACK 的信息，请参见 RFC 2018。可能的值及其含义：

never

从不接受 SACK 和外发 SACK 信息。

passive

将接受但不外发 SACK。

active

将接受 SACK 并外发 SACK 信息。

smallest_anon_port (TCP, SCTP, UDP)**largest_anon_port (TCP, SCTP, UDP)**

这些选项定义临时端口上的上限和下限。当建立传出网络连接时，将使用临时（表示短暂）端口。请注意，**smallest_anon_port** 的当前值始终应当小于或等于 **largest_anon_port** 的当前值。

smallest_nonpriv_port (TCP, SCTP, UDP)

此选项定义非特权端口的起点。非特权端口范围通常从 1024 开始。尝试绑定非特权端口的任何程序都不必作为根运行。

ttl (IPv4)、hoplimit (IPv6)

指定将为 IPv4 或 IPv6 头的 **ttl/hoplimit** 字段设置的值。可用于防止系统到达其他系统中 *N* 跳以外的位置，其中 *N* 是指定的值。

接口属性

支持下列接口属性：

arp

在接口上启用/禁用对地址解析协议 (Address Resolution Protocol, ARP) 的使用。ARP 用于在网络级别地址和链路级别地址之间进行映射。当前已对 IPv4 地址与 MAC 地址之间的映射实现此选项。可能的值包括 **on** 或 **off**。缺省值为 **on**。

exchange_routes

在此接口上启用/禁用路由信息交换。可能的值包括 **on** 或 **off**。缺省值为 **off**。

group

指定采用此接口作为其底层接口的 IPMP 接口的组名称。如果此接口是 IPMP 类，此选项指定 IPMP 组的名称。此选项仅在 IPMP 接口上才是读写属性。对于其他接口类，此属性为只读。

forwarding

在接口上启用/禁用 IP 转发。启用 IP 转发后，可将 IP 包转发到接口以及从此接口转发 IP 包。可能的值包括 **on** 或 **off**。缺省值为 **off**。

metric

将接口的路由度量设置为 *n*；如果未指定值，则缺省为 **0**。路由度量由路由协议使用。较高的度量会降低相应路由的适用性。度量为到目标网络或主机的额外跳数。

mtu

将接口的最大传输单元设置为 *n*。对于许多类型的网络，MTU 具有上限，例如，以太网的 MTU 上限为 1500。

nud

在点对点物理接口上启用/禁用邻居不可达检测机制。可能的值包括 **on** 或 **off**。缺省值为 **on**。

standby

指定是否将此接口配置为 IPMP 组的备用接口。此属性不适用于 IPMP 接口。

usesrc

指定用于选择源地址的物理或虚拟接口。如果使用关键字 **none**，则将清除以前的所有选择。缺省值为 **none**。

地址属性

支持下面列出的地址属性。请注意，不支持修改 **addrconf** 地址对象的地址属性。

deprecated

此地址不应在新通信中继续用作源地址，但发往此地址的包仍会按预期方式处理。可能的值包括 **on** 或 **off**。缺省值为 **off**。在类型为 **dhcp** 的地址对象上，不支持此属性。

prefixlen

指定地址的 IPv6 前缀或 IPv4 网络掩码所包含的地址的最左侧的连续位数。其余低位定义地址的主机部分。将 **prefixlen** 转换为地址的文本表现形式时，对于用于网络部分的位位置，此地址包含 1，对于主机部分，此地址包含 0。**prefixlen** 必须指定为一个十进制数字。在类型为 **dhcp** 的地址对象上，不支持此属性。

private

指定地址不应由 `in.routed` 路由守护进程声明。可能的值包括 `on` 或 `off`。缺省值为 `off`。

reqhost

客户机要让 DHCP 服务器将客户机的租用 IPv4 映射到的主机名。不保证能够完成主机名请求。

transmit

启用要使用地址对象引用的地址传输的包。当地址处于启用状态时，这是缺省行为。可能的值包括 `on` 或 `off`。缺省值为 `on`。

zone

指定应在其中放置地址对象引用的所有地址的区域。命名区域在处于就绪或运行中状态的内核中必须处于活动状态。当区域停止或重新引导时，接口将被取消激活。必须将此区域配置为共享 IP 区域。`zonecfg(1M)` 用于向专用 IP 区域分配网络接口名称。要修改区域赋值使之在重新引导后保留，请使用 `zonecfg(1M)`。可能的值是在系统上配置的所有区域的列表。缺省值为 `global`。

操作数

每个 `ipadm` 子命令都在下列某个对象中运行：

addrobj

在网络接口上配置的地址通过 *addrobj* 标识。*addrobj* 包括两个部分。第一个部分为配置地址所在的网络接口的名称。第二个部分是用户指定的字符串，此字符串可使用任何字母数字字符，其长度最多包含 32 个字符，并且必须以字母开头。*addrobj* 的这两个部分使用斜杠 (/) 分隔。地址对象始终表示系统中的唯一地址集。

interface

配置网络地址所在的网络接口的名称。一般说来，此名称可使用任何字母数字字符、下划线 (_) 和句点 (.)，但必须以字母字符开头，并且以数字结尾。

protocol

要配置其属性的 TCP/IP Internet 协议族的名称。支持下列协议：`ip`、`ipv4`、`ipv6`、`icmp`、`tcp`、`sctp` 和 `udp`。

示例**示例 1 创建 IPv4 静态地址**

下面的命令在接口 `bge1`（链路名称 `net1`）上创建地址 `10.2.3.4/24`，并将此地址标记为 `up` 以便使用。

```
# ipadm create-ip net1
# ipadm create-addr -T static -a local=10.2.3.4/24 net1/v4static1
```

下面的命令在接口 `net1` 上创建另一个地址 `10.2.3.5/24`，但将此地址标记为 `down`，直到显式标记为 `up` 为止。

```
# ipadm create-addr -T static -d -a 10.2.3.5/24 net1/v4static2
```

请注意，因为未使用本地地址且只有一个地址，所以假定 `10.2.3.5/24` 为本地地址。

下面的命令将以前标记为 `down` 的地址对象 `net1/v4static2` 标记为 `up`。

示例1 创建IPv4静态地址 (续)

```
# ipadm up-addr net1/v4static2
```

如果在地址对象上设置了 DUPLICATE 标志，refresh-addr 将验证此地址是否仍在网络上重复。如果不重复，此地址将标记为 up。

```
# ipadm refresh-addr net1/v4static2
```

下面的命令列出已配置的地址。此命令显示地址 net1/v4static2 不是重复地址。

```
# ipadm show-addr
ADDROBJ      TYPE    STATE    ADDR
lo0/v4       static  ok       127.0.0.1/8
lo0/v6       static  ok       ::/128
net1/v4static1 static  ok       10.2.3.4/24
net1/v4static2 static  ok       10.2.3.10/24
```

示例2 创建受 DHCPv4 控制的地址

下面的命令在接口 bge1（链路名称 net1）上获取 DHCPv4 地址。

```
# ipadm create-ip net1
# ipadm create-addr -T dhcp net1/dhaddr
# ipadm show-addr net1/dhaddr
ADDROBJ      TYPE    STATE    ADDR
net1/dhaddr   dhcp    ok       10.8.48.173/25
```

下面的命令延长 DHCPv4 地址对象 net1/dhaddr 的租用持续时间。

```
# ipadm refresh-addr net1/dhaddr
```

示例3 创建 IPv6 地址

下面的命令序列使用具有缺省接口 ID 的 in.ndpd 在 bge1（链路名称 net1）上自动配置 IPv6 地址。首先配置链路本地地址，然后 in.ndpd 添加无状态和有状态自动配置地址。

```
# ipadm create-ip net1
# ipadm create-addr -T addrconf net1/v6addr
```

下面的命令创建一个 IPv6 静态地址。为了能够配置不是链路本地地址的 IPv6 地址，接口上应当已配置链路本地地址。这是通过使用 -T addrconf 在上一步骤中实现的。

```
# ipadm create-addr -T static -a local=2ff0::f3ad/64 net1/v6static
```

下面的命令更改 IPv6 地址的前缀长度。

```
# ipadm set-addrprop -p prefixlen=80 net1/v6static
```

通过列出地址，可以查看自动配置的所有地址以及更新后的前缀长度：

示例3 创建IPv6地址 (续)

```
# ipadm show-addr
ADDROBJ      TYPE      STATE      ADDR
lo0/v4        static    ok          127.0.0.1/8
lo0/v6        static    ok          ::/128
net1/v6addr    addrconf ok          fe80::203:baff:fe94:2f01/10
net1/v6addr    addrconf ok          2002:a08:39f0:1:203:baff:\
                                     fe94:2f00/64
net1/v6addr    addrconf ok          2001:db8:1:2::402f/128
net1/v6static  static    ok          2ff0::f3ad/80
```

示例4 配置IPv4通道

下面的第一个命令(ipadm)创建通道源地址。然后，dladm命令创建一个通道链路。最后的ipadm命令在通道IP接口上配置IPv4和IPv6地址。

```
# ipadm create-ip net1
# ipadm create-addr -T static -a local=10.2.3.4/24 net1/v4static
# dladm create-iptun -T ipv4 -a local=10.2.3.4,remote=10.2.3.5 tun0
# ipadm create-ip tun0
# ipadm create-addr -T static \
    -a local=173.129.134.1,remote=173.129.134.2 tun0/v4tunaddr
# ipadm create-addr -T static \
    -a local=2ff1::3344,remote=2ff1::3345 tun0/v6tunaddr
# ipadm show-addr
ADDROBJ      TYPE      STATE      ADDR
lo0/v4        static    ok          127.0.0.1/8
lo0/v6        static    ok          ::/128
net1/v4static  static    ok          10.2.3.4/24
tun0/v4tunaddr static    ok          173.129.134.1-->173.129.134.2
tun0/v6tunaddr static    ok          2ff1::3344-->2ff1::3345
```

示例5 查看所有接口

下面的命令用于查看所有接口。

```
# ipadm show-if -o all
IFNAME CLASS  STATE  ACTIVE CURRENT      PERSISTENT OVER
lo0     loopback ok      yes    -m-v-----46 --46 --
net0     ip      ok      yes    bm-----46 --46 --
e1000g0 ip      ok      yes    bm---l----46 -l46 --
e1000g1 ip      ok      yes    bm---l----46 -l46 --
ipmp0   ipmp    down    yes    bm-----46 --46 e1000g0 e1000g1
tun0     ip      failed  no     -mp-----46 --46 --
vni0     vni     disabled no     bm-v----- --46 --
```

示例6 显示接口属性

下面的命令显示指定接口的所有接口属性。

示例6 显示接口属性 (续)

```
# ipadm show-ifprop net0
IFNAME PROPERTY      PROTO PERM CURRENT PERSISTENT DEFAULT POSSIBLE
net0  arp              ipv4  rw   on    --      on    on,off
net0  forwarding       ipv4  rw   off   on      off   on,off
net0  metric            ipv4  rw   2     2       0     --
net0  mtu               ipv4  rw   1500  --      1500  68-1500
net0  exchange_routes  ipv4  rw   off   --      off   on,off
net0  usesrc            ipv4  rw   none  --      none  --
net0  forwarding       ipv6  rw   off   --      off   on,off
net0  metric            ipv6  rw   2     2       0     --
net0  mtu               ipv6  rw   1500  --      1500  1280-1500
net0  nud              ipv6  rw   on    --      on    on,off
net0  exchange_routes  ipv6  rw   off   on      off   on,off
net0  usesrc            ipv6  rw   none  --      none  --
net0  group             ip    rw   grp0  --      --    --
net0  standby          ip    r-   off   --      off   on,off
```

示例7 配置基于接口的属性

下面的命令将接口 net0 的 IPv4 MTU 设置为 900。

```
# ipadm set-ifprop -m ipv4 -p mtu=900 net0
```

下面的命令将接口 net0 的 IPv6 MTU 设置为 1400。

```
# ipadm set-ifprop -m ipv6 -p mtu 1400 net0
```

查看结果：

```
# ipadm show-ifprop -p mtu net0
IFNAME PROPERTY      PROTO PERM CURRENT PERSISTENT DEFAULT POSSIBLE
net0  mtu              ipv4  rw   900    900      1500  68-1500
net0  mtu              ipv6  rw   1400   1400     1500  1280-1500

# ipadm show-ifprop -m ipv6 -p mtu net0
IFNAME PROPERTY      PROTO PERM CURRENT PERSISTENT DEFAULT POSSIBLE
net0  mtu              ipv6  rw   1400   1400     1500  1280-1500
```

示例8 显示支持的属性

下面的命令显示在 TCP 上支持的属性。

```
# ipadm show-prop tcp
PROTO PROPERTY      PERM CURRENT PERSISTENT DEFAULT POSSIBLE
tcp  ecn              rw   active  active    passive  never,passive,
                                     active
tcp  extra_priv_ports  rw   --      1,65535  2049,4045 1-65535
tcp  largest_anon_port  rw   32768   32768    65535     1024-65535
tcp  sack              rw   active  --      active    never,passive,
```

示例8 显示支持的属性 (续)

						active
tcp	recv_buf	rw	29567	--	49152	2048-1073741824
tcp	send_buf	rw	21354	--	49152	4096-1073741824
tcp	max_buf	ro	65536	--	32768	4096-1073741824
tcp	smallest_anon_port	rw	32768	--	32768	1024-65535
tcp	smallest_nonpriv_port	rw	1024	--	1024	1024-32768

示例9 配置全局 IPv4 转发

下面的命令序列配置全局 IPv4 转发，并覆盖接口 net0 的相应设置。

```
# ipadm set-prop -p forwarding=on ipv4
# ipadm set-ifprop -p forwarding=off -m ipv4 net0
# ipadm show-prop -p forwarding ipv4
```

PROTO	PROPERTY	PERM	CURRENT	PERSISTENT	DEFAULT	POSSIBLE
ipv4	forwarding	rw	on	on	off	on,off

```
# show-ifprop -p forwarding -m ipv4 net0
```

IFNAME	PROPERTY	PROTO	PERM	CURRENT	PERSISTENT	DEFAULT	POSSIBLE
net0	forwarding	ipv4	rw	off	off	off	on,off

示例10 在 set-prop 子命令中使用限定符

下面的命令序列使用加号和减号 (+、-) 限定符将 1047、1048 和 1049 添加为 TCP 的额外特权端口。

```
# ipadm set-prop -p extra_priv_ports=1047 tcp
# ipadm set-prop -p extra_priv_ports+=1048 tcp
# ipadm set-prop -p extra_priv_ports+=1049 tcp
# ipadm set-prop -p extra_priv_ports+=1050 tcp
```

下面的命令删除作为额外特权端口的 1048。

```
# ipadm set-prop -p extra_priv_ports-=1048
```

下面的命令显示 TCP 的所有额外特权端口。

```
# ipadm show-prop -p extra_priv_ports tcp
```

PROTO	PROPERTY	PERM	CURRENT	PERSISTENT	DEFAULT	POSSIBLE
ipv4	extra_priv_ports	rw	1047,1049, 1050	1047,1049, 1050	2049,4045	1-65535

示例11 启用和禁用对象

下面的命令序列启用和禁用接口和地址对象，并显示这些操作的结果。

```
# ipadm create-ip net1
# ipadm create-addr -T static -a local=10.2.3.4/24 net1/v4static
# ipadm set-addrprop -p private=yes net1/v4static
# ipadm show-addr net1/v4static
```

示例 11 启用和禁用对象 (续)

```
ADDROBJ      TYPE  STATE  ADDR
net1/v4static static ok      10.2.3.4/24
```

下面的命令禁用地址对象 net1/v4static。

```
# ipadm disable-addr -t net1/v4static
# ipadm show-addr net1/v4static
ADDROBJ      TYPE  STATE  ADDR
net1/v4static static ok      10.2.3.4/24
```

下面的命令禁用接口对象 net1。

```
# ipadm disable-if -t net1
# ipadm show-if net1 -o all
IFNAME      CLASS  STATE  ACTIVE CURRENT      PERSISTENT OVER
net1        ip      disabled no      bm----- --46      --
```

下面的命令从持久性配置启用接口对象。

```
# ipadm enable-if -t net1
# ipadm show-if net1 -o all
IFNAME      CLASS  STATE  ACTIVE CURRENT      PERSISTENT OVER
net1        ip      ok      yes      bm-----46 --46      --

# ipadm show-addr net1/v4static
ADDROBJ      TYPE  STATE  ADDR
net1/v4static static ok      10.2.3.4/24
```

请注意，当启用接口对象时，还将启用在此接口上配置的所有地址对象。

下面的命令在非全局专用 IP 区域中为 net0 接口创建持久性配置，以便在下次重新引导时使用全局区域中通过 allowed-address 资源提供的地址集来配置 net0 接口。

```
# ipadm create-ip net0
```

此外，还可以通过执行下面的命令，使用非全局专用 IP 区域中的可用 allowed-address 值集来配置 net0 接口，而不必进行重新引导：

```
# ipadm disable-if -t net0
# ipadm enable-if -t net0
```

示例 12 创建 IPMP 接口

下面的命令序列创建一个 IPMP 接口，并向其添加底层接口。

```
# ipadm create-ip e1000g0
# ipadm create-ip e1000g1
# ipadm create-ip e1000g2
# ipadm set-ifprop -p standby-on -m ip e1000g2
# ipadm create-ipmp testgroup0
```


示例 12 创建 IPMP 接口 (续)

```
# ipadm add-ipmp -i e1000g0 -i e1000g1 -i e1000g2 testgroup0
# ipadm create-addr -T static -a local=192.168.80.5/24 testgroup0/data1
# ipadm create-addr -T static -a local=192.168.80.6/24 testgroup0/data2
# ipadm show-if
IFNAME      CLASS    STATE    ACTIVE OVER
lo0         loopback ok       yes  --
net0        ip       ok       yes  --
e1000g0     ip       ok       yes  --
e1000g1     ip       ok       yes  --
ipmp0       ipmp     ok       yes  e1000g0 e1000g1
```

下面的命令序列先禁用后启用 IPMP 接口。

```
# ipadm disable-if -t testgroup0
ipadm show-if
IFNAME      CLASS    STATE    ACTIVE OVER
lo0         loopback ok       yes  --
net0        ip       ok       yes  --
e1000g0     ip       disabled no    --
e1000g1     ip       disabled no    --
ipmp0       ipmp     disabled no    e1000g0 e1000g1
# ipadm enable-if -t testgroup0
```

下面的命令序列删除 IPMP 接口的底层接口，然后删除 IPMP 接口。

```
ipadm remove-ipmp -i e1000g0 -i e1000g1 testgroup0
ipadm delete-ipmp testgroup0
```

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）

另请参见

[read\(1\)](#)、[arp\(1M\)](#)、[cfgadm\(1M\)](#)、[dhcagent\(1M\)](#)、[dladm\(1M\)](#)、[if_mpadm\(1M\)](#)、[ifconfig\(1M\)](#)

《Oracle Solaris Tunable Parameters Reference Manual》

由 Postel, J. 编著，RFC 791，《Internet Protocol - DARPA Internet Program Protocol Specification》，南加州大学信息科学研究所，1981 年 9 月。

由 Hinden, R. 和 S. Deering 合著的《IP Version 6 Addressing Architecture》，RFC 4291，2006 年 2 月。

由 Thomson, S.、Narten, T. 和 T. Jinmei 合著的《IPv6 Stateless Address AutoConfiguration》，RFC 4862，2007 年 9 月。

由 Droms, R.、Bound, J.、Volz, B.、Lemon, T.、Perkins, C. 和 M. Carney 合著的《Dynamic Host Configuration Protocol for IPv6 (DHCPv6)》，RFC 3315，2003 年 7 月。

由 Narten, T.、Draves, R. 和 S. Krishnan 合著的《Privacy Extensions for Stateless Address AutoConfiguration in IPv6》，RFC 4941，2007 年 9 月。

由 S. Routhier, Ed. 编著的《Management Information Base for the Internet Protocol (IP)》，RFC 4293，2006 年 4 月。

由 Braden, R. 编著，RFC 1122，《Requirements for Internet Hosts - Communication Layers》，南加州大学信息科学研究所，1989 年 10 月。

引用名	ipf – 更改 IP 包输入和输出的包过滤列表
用法概要	<pre>ipf [-6AdEInoPRrsVyzZ] [-l block pass nomatch] [-T optionlist] [-F i o a s S] -f filename [-f filename...]</pre>
描述	ipf 实用程序属于与 Solaris IP 过滤器功能关联的一套命令。请参见 ipfilter(5)。 ipf 实用程序打开列出的文件名（将连字符 (-) 视为 <code>stdin</code>）并针对要在包过滤器规则集中添加或删除的规则集解析该文件。 如果没有解析问题，则会将 ipf 处理的每个规则添加到内核的内部列表中。规则将被添加到内部列表的末尾，并按提供给 ipf 时的先后顺序排列。 只能通过 <code>/dev/ipauth</code>、<code>/dev/ipl</code> 和 <code>/dev/ipstate</code> 访问 ipf。这些文件的缺省权限要求 ipf 以 <code>root</code> 的身份运行才能执行所有操作。
启用 Solaris IP 过滤器功能	Solaris IP 过滤器随 Solaris 操作系统一起安装。但是，缺省情况下不启用包过滤。可以使用以下过程激活 Solaris IP 过滤器功能。 1. 担任一个具有 "IP Filter Management"（IP 过滤器管理）权限配置文件的角色（请参见 rbac(5)）或成为超级用户。 2. 配置系统和服务的防火墙策略。请参见 svc.ipfd(1M) 和 ipf(4)。 3. （可选）创建网络地址转换 (network address translation, NAT) 配置文件。请参见 ipnat.conf(4)。 4. （可选）创建地址池配置文件。请参见 ippool(4)。 如果要将一组地址作为单个地址池引用，请创建 <code>ippool.conf</code> 文件。如果希望在引导时装入地址池配置文件，请创建一个名为 <code>/etc/ipf/ippool.conf</code> 的文件，在其中放置地址池。如果不希望在引导时装入地址池配置文件，请将 <code>ippool.conf</code> 文件放置在 <code>/etc/ipf</code> 之外的位置，然后手动激活这些规则。 5. 启用 Solaris IP 过滤器，如下所述： <pre># svcadm enable network/ipfilter</pre> 要在临时禁用包过滤功能后重新启用，请重新引导计算机或输入以下命令： <pre># svcadm enable network/ipfilter</pre> ...这本质上是执行以下 ipf 命令： 1. 启用 Solaris IP 过滤器： <pre># ipf -E</pre> 2. 装入 <code>ippools</code>： <pre># ippool -f <ippool configuration file></pre> 请参见 ippool(1M)。 3. （可选的）激活包过滤：

ipf -f <ipf configuration file>

4. （可选的）激活 NAT：

ipnat -f <IPNAT configuration file>

请参见 [ipnat\(1M\)](#)。

注- 如果重新引导系统，IPfilter 配置会自动激活。

选项

所支持的选项如下：

-6

要解析并装入 IPv6 规则，此选项是必需的。IPv6 规则的装入将来可能会有更改。

-A

设置列表以更改活动列表（缺省）。

-d

打开调试模式。使得它在处理每个过滤器规则时生成过滤器规则的十六进制转储。

-D

禁用过滤器（如果已启用）。对可装入的内核版本无效。

-M

启用过滤器（如果已禁用）。对可装入的内核版本无效。

-F i | o | a

指定要刷新的过滤器列表。参数应该为 **i**（输入）、**o**（输出）或 **a**（删除所有过滤器规则）。可以使用单个字母或以相应字母开头的整个单词。此选项可以放在任何其他选项之前或之后，选项在命令行上的顺序确定了用来执行选项的顺序。

-F s | S

要刷新状态表中的条目，可将 **-F** 选项与 **s**（删除关于任何非完全建立的连接的状态信息）或 **S**（删除整个状态表）一起使用。只能指定这两个选项中的一个。完全建立的连接将会在 **ipfstat -s** 输出中显示为 4/4，这两个方面中任何一个有偏差都表示连接不是完全建立的。

-f filename

指定 **ipf** 应当从哪些文件获取输入来修改包过滤器规则列表。

-I

设置列表以更改非活动的列表。

-l pass | block | nomatch

切换包的缺省日志记录。此选项的有效参数为 **pass**、**block** 和 **nomatch**。如果设置了某个选项，则会记录退出了过滤并且与所设置的类别匹配的任何包。这最适用于记录与所装入的任何规则都不匹配的所有包。

-n

防止 **ipf** 进行任何 **ioctl** 调用或执行会更改当前运行的内核的任何操作。

- O 强制在缺省情况下将规则添加到输出列表（或从中删除规则），而不是在（缺省）输入列表中执行这些操作。
- P 将规则添加为验证规则表中的临时条目。
- R 同时禁用“IP 地址到主机名”解析和“端口号到服务名称”解析。
- r 删除匹配的过滤器规则，而不是将它们添加到内核列表中。
- S 使当前活动的过滤器列表成为备用列表。
- T *optionlist* 允许 IPFilter 内核变量的运行时更改。要允许更改，某些变量要求 IPFilter 处于禁用状态 (-D)，另一些变量则不要求。*optionlist* 参数是以逗号分隔的调优命令列表。调优命令是以下项之一：
 - list** 检索内核中所有变量的列表、它们的最大值、最小值和当前值。
 - 单个变量名称 检索它的当前值。
 - 后跟一个赋值的变量名称 设置一个新值。
 示例如下：


```
# Print out all IPFilter kernel tunable parameters
ipf -T list

# Display the current TCP idle timeout and then set it to 3600
ipf -D -T fr_tcpidletimeout,fr_tcpidletimeout=3600 -E

# Display current values for fr_pass and fr_chksrc, then set
# fr_chksrc to 1.
ipf -T fr_pass,fr_chksrc,fr_chksrc=1
```
- v 打开详细模式。显示与规则处理相关的信息。
- V 显示版本信息。这将显示编译为 **ipf** 二进制代码的版本信息并从内核代码中检索该信息（如果正在运行或存在）。如果它存在于内核中，将会显示有关它的当前状态的信息；例如，日志记录是否活动、缺省过滤等。
- y 手动将由 IP 过滤器维护的内核中接口列表与当前接口状态列表重新同步。

- Z
对于输入文件中的每个规则，将它的统计信息重置为零并显示在它们归零前的统计信息。
- Z
使内核中保留的仅用于过滤的全局统计信息归零。这不会影响片段或状态统计信息。

文件

/dev/ipauth
/dev/ipl
/dev/ipstate
指向 IP 过滤器伪设备的链接。

/etc/ipf/ipf.conf
ipf 启动配置文件的位置。请参见 ipf(4)。

/usr/share/ipfilter/examples/
包含大量的 IP 过滤器示例。

属性

有关以下属性的说明，请参见 attributes(5)：

属性类型	属性值
可用性	network/ipfilter
接口稳定性	Committed（已确定）

另请参见

ipfstat(1M)、ipmon(1M)、ipnat(1M)、ippool(1M)、svcadm(1M)、svc.ipfd(1M)、ipf(4)、ipnat.
《Oracle Solaris 管理：IP 服务》

诊断

需要以 root 的身份运行才能实际在内核中影响包过滤列表。

引用名	ipfs – saves and restores information for NAT and state tables
用法概要	<pre>ipfs [-nv] -l ipfs [-nv] -u ipfs [-nv] [-d <i>dirname</i>] -R ipfs [-nv] [-d <i>dirname</i>] -W ipfs [-nNSv] [-f <i>filename</i>] -r ipfs [-nNSv] [-f <i>filename</i>] -w ipfs [-nNSv] -f <i>filename</i> -i <if1>,<if2></pre>
描述	The <code>ipfs</code> utility enables the saving of state information across reboots. Specifically, the utility allows state information created for NAT entries and rules using "keep state" to be locked (modification prevented) and then saved to disk. Then, after a reboot, that information is restored. The result of this state-saving is that connections are not interrupted.
选项	The following options are supported: -d Change the default directory used with -R and -W options for saving state information. -n Do not take any action that would affect information stored in the kernel or on disk. -v Provides a verbose description of <code>ipfs</code> activities. -N Operate on NAT information. -S Operate on filtering state information. -u Unlock state tables in the kernel. -l Lock state tables in the kernel. -r Read information in from the specified file and load it into the kernel. This requires the state tables to have already been locked and does not change the lock once complete. -w Write information out to the specified file and from the kernel. This requires the state tables to have already been locked and does not change the lock once complete. -R Restores all saved state information, if any, from two files, <code>ipstate.ipf</code> and <code>ipnat.ipf</code>, stored in the <code>/var/db/ipf</code> directory. This directory can be changed with the -d option. The state tables are locked at the beginning of this operation and unlocked once complete. -W Saves in-kernel state information, if any, out to two files, <code>ipstate.ipf</code> and <code>ipnat.ipf</code>, stored in the <code>/var/db/ipf</code> directory. This directory can be changed with the -d option. The state tables are locked at the beginning of this operation and unlocked once complete.

- 文件
- /var/db/ipf/ipstate.ipf
 - /var/db/ipf/ipnat.ipf
 - /dev/ipl
 - /dev/ipstate
 - /dev/ipnat

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	network/ipfilter
Interface Stability	Committed

另请参见

[ipf\(1M\)](#), [ipmon\(1M\)](#), [ipnat\(1M\)](#), [attributes\(5\)](#)

诊断

Arguably, the -W and -R operations should set the locking and, rather than undo it, restore it to what it was previously.

Fragment table information is currently not saved.

引用名 ipfstat – 报告包过滤器统计信息和过滤器列表

用法概要

```
ipfstat [-6aACdfghIilnoRstv]
```

```
ipfstat [-C] [-D addrport] [-P protocol] [-S addrport]  
[-T refreshtime]
```

描述

ipfstat 命令属于与 Solaris IP 过滤器功能关联的一套命令。请参见 [ipfilter\(5\)](#)。

ipfstat 命令使用 `_fr_flags`、`_frstats`、`_filterin` 和 `_filterout` 符号来检查 `/dev/kmem`。要运行和工作，它需要能够读取 `/dev/kmem` 和内核本身。

ipfstat 的缺省行为是检索并显示自内核通过过滤器放置包以来随时间推移累积的统计信息。

ipfstat 的职责是将当前收集的内核统计信息作为向流入和流出内核的包应用过滤器（如果存在）后的结果进行显示。当未提供命令行参数时，这是缺省操作。当提供了 `-i` 或 `-o` 时，ipfstat 将检索并显示当前安装的及内核正使用的过滤器规则的相应列表。

ipfstat 使用内核设备文件获取信息。这些文件的缺省权限要求 ipfstat 以 root 身份运行才能执行所有操作。

ipfstat 命令支持 [kstat\(3KSTAT\)](#) 内核工具。由于此项支持功能，可使用 [kstat\(1M\)](#) 作为 ipfstat 的替代命令。例如：

```
# kstat -m ipf
```

使用 ipfstat -t 选项可使 ipfstat 进入状态排名模式。在此模式下，状态表的显示方式与 Unix top 实用程序显示进程表类似。可以使用 `-C`、`-D`、`-P`、`-S` 和 `-T` 命令行选项限制要显示的状态项和指定显示更新的频率。

在状态排名模式下，可使用以下键影响显示的信息：

- d 选择要显示的信息。
- l 重画屏幕。
- q 退出程序。
- s 在不同排序条件之间切换。
- r 颠倒排序条件。

可按协议编号、IP 包数、字节数、状态项的生存时间对状态进行排序。缺省行为是按字节数进行排序。状态按降序排序，但可以使用 `r` 键使它们按升序排序。

无法以交互方式更改源、目标和协议过滤器或刷新频率。这些操作必须通过命令行执行。

屏幕必须具有 80 个列才能正确显示。不过，ipfstat 不会检查屏幕宽度。

只显示与排序和过滤条件匹配的前 $X-5$ 个项（其中 X 是显示器上的行数）。无法查看更多的项。

选项

所支持的选项如下：

- 6 显示过滤器列表和 IPv6 的状态（如果有）。此选项将来可能会更改。
- a 显示记帐过滤器列表并显示为每个规则统计的字节数。
- A 显示包验证统计信息。
- C 仅当与 `-t` 结合使用时有效。还会显示排名靠前的 `closed` 状态。通常情况下，如果 TCP 连接达到了 `CLOSE_WAIT` 协议状态，则不会显示该连接。如果启用了此选项，将会显示所有状态项。
- d 显示数据时生成调试输出。
- D *addrport* 仅当与 `-t` 结合使用时有效。将状态排名显示限制为仅显示其目标 IP 地址和端口与 *addrport* 参数相匹配的状态项。*addrport* 是以 *ipaddress* [*,port*] 格式指定的。*ipaddress* 和 *port* 应该是数字或字符串 `any`（以该顺序指定任何 IP 地址和端口）。如果未指定 `-D` 选项，则缺省设置为 `-D any, any`。
- f 如果存在任何片段状态信息（统计信息）和保留的状态信息（在内核中），将会显示它们。
- g 显示当前配置的组（包括活动的和非活动的）。
- h 显示每个规则被“命中”的次数。需与 `-i` 结合使用。
- i 显示用于内核 IP 处理的输入端的过滤器列表。
- I 在检索 `inactive/active` 过滤器列表详细信息之间切换。需与 `-i` 结合使用。
- l 当与 `-s` 一起使用时，会显示活动状态项的列表（没有统计信息）。
- n 显示每个规则在输出时的规则编号。
- o 显示用于内核 IP 处理的输出端的过滤器列表。
- P *protocol* 仅当与 `-t` 结合使用时有效。将状态排名显示限制为仅显示与特定协议匹配的状态项。该参数可以是协议名称（如 `/etc/protocols` 中定义的）或协议编号。如果未指定此选项，则会指定任何协议的状态项。
- R 同时禁用“IP 地址到主机名”解析和“端口号到服务名称”解析。
- S *addrport* 仅当与 `-t` 结合使用时有效。将状态排名显示限制为仅显示其源 IP 地址和端口与 *addrport* 参数相匹配的状态项。*addrport* 是以 *ipaddress*

[,port] 格式指定的。 *ipaddress* 和 *port* 应该是数字或字符串 *any*（以该顺序指定任何 IP 地址和端口）。如果未指定 *-s* 选项，则缺省设置为 *-S any, any*。

- s 显示包/流状态信息（仅限于统计信息）。
- T *refreshtime* 仅当与 *-t* 结合使用时有效。指定状态**排名**显示应以何频率进行更新。刷新时间是两次更新之间的秒数。可以使用任何正整数。缺省值（最小更新时间）是 1。
- t 以与 Unix 实用程序 *top* 显示进程表类似的方式显示状态表。可采用多种不同种方式对状态进行排序。
- v 打开详细模式。显示更多调试信息。

文件

- /dev/kmem
- /dev/ksyms
- /dev/ipl
- /dev/ipstate

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	network/ipfilter
接口稳定性	Committed（已确定）

另请参见

[ipf\(1M\)](#)、[kstat\(1M\)](#)、[kstat\(3KSTAT\)](#)、[attributes\(5\)](#)、[ipfilter\(5\)](#)
《Oracle Solaris 管理：IP 服务》

引用名 ipmgmtd – IP 网络接口和 TCP/IP 可调参数管理守护进程

用法概要 /lib/inet/ipmgmtd
svc:/network/ip-interface-management:default

描述 ipmgmtd 是处理针对网络 IP 接口和 IP/TCP/UDP/SCTP/ICMP 可调参数的管理事件的系统守护进程。通过服务管理工具 (Service Management Facility, SMF) 服务实例对其进行控制：

svc:/network/ip-interface-management:default

ipmgmtd 守护进程是由 SMF 服务自动启动的，不应该直接调用。它不是管理或编程接口的一部分。可通过 [ipadm\(1M\)](#) 调用用于管理 IP 接口和上述协议可调参数的管理接口。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Private（专用）

另请参见 [ipadm\(1M\)](#)、[attributes\(5\)](#)

引用名	ipmon – monitors /dev/ipl for logged packets
用法概要	<pre>ipmon [-abDFhnpstvxX] [-N device] [[o] [NSI]] [-O [NSI]] [-P pidfile] [-S device] [-f device] [filename]</pre>
描述	The <code>ipmon</code> command is part of a suite of commands associated with the Solaris IP Filter feature. See ipfilter(5). The <code>ipmon</code> command opens <code>/dev/ipl</code> for reading and awaits data to be saved from the packet filter. The binary data read from the device is reprinted in human readable form. However, IP addresses are not mapped back to hostnames, nor are ports mapped back to service names. The output goes to standard output, by default, or a filename, if specified on the command line. Should the <code>-s</code> option be used, output is sent instead to syslogd(1M). Messages sent by means of <code>syslog</code> have the day, month, and year removed from the message, but the time (including microseconds), as recorded in the log, is still included. Messages generated by <code>ipmon</code> consist of whitespace-separated fields. Fields common to all messages are: ■ The date of packet receipt. This is suppressed when the message is sent to <code>syslog</code>. ■ The time of packet receipt. This is in the form <i>HH:MM:SS.F</i>, for hours, minutes, seconds, and fractions of a second (which can be several digits long). ■ The name of the interface on which the packet was processed, for example, <code>ib1</code>. ■ The group and rule number of the rule, for example, <code>@0:17</code>. These can be viewed with <code>ipfstat -in</code> for input rules or <code>ipfstat -o</code> for output rules. See ipfstat(1M). ■ The action: <code>p</code> for passed, <code>b</code> for blocked, <code>s</code> for a short packet, <code>n</code> did not match any rules, or <code>L</code> for a log rule. ■ The addresses. This is actually three fields: the source address and port (separated by a comma), the symbol <code>→</code>, and the destination address and port. For example: <code>209.53.17.22,80 → 198.73.220.17,1722</code>. ■ <code>PR</code> followed by the protocol name or number, for example, <code>PR tcp</code>. ■ <code>len</code> followed by the header length and total length of the packet, for example, <code>len 20 40</code>. If the packet is a TCP packet, there will be an additional field starting with a hyphen followed by letters corresponding to any flags that were set. See ipf.conf(4) for a list of letters and their flags. If the packet is an ICMP packet, there will be two fields at the end, the first always being <code>icmp</code>, the next being the ICMP message and submessage type, separated by a slash. For example, <code>icmp 3/3</code> for a port unreachable message.
选项	The following options are supported: -a Open all of the device logfiles for reading log entries. All entries are displayed to the same output device (<code>stderr</code> or <code>syslog</code>).

- b
For rules which log the body of a packet, generate hex output representing the packet contents after the headers.
- D
Cause ipmon to turn itself into a daemon. Using subshells or backgrounding of ipmon is not required to turn it into an orphan so it can run indefinitely.
- f *device*
Specify an alternative device/file from which to read the log information for normal IP Filter log records.
- F
Flush the current packet log buffer. The number of bytes flushed is displayed, even if the result is zero.
- h
Displays usage information.
- n
IP addresses and port numbers will be mapped, where possible, back into hostnames and service names.
- N *device*
Set the logfile to be opened for reading NAT log records from or to *device*.
- o *letter*
Specify which log files from which to actually read data. N, NAT logfile; S, state logfile; I, normal IP Filter logfile. The -a option is equivalent to using -o NSI.
- O *letter*
Specify which log files you do not wish to read from. This is most commonly used in conjunction with the -a. Letters available as parameters are the same as for -o.
- p
Cause the port number in log messages always to be printed as a number and never attempt to look it up.
- P *pidfile*
Write the PD of the ipmon process to a file. By default this is /var/run/ipmon.pid.
- s
Packet information read in will be sent through syslogd rather than saved to a file. The default facility when compiled and installed is local0. The following levels are used:

LOG_INFO
Packets logged using the log keyword as the action rather than pass or block.

LOG_NOTICE
Packets logged that are also passed.

- LOG_WARNING
Packets logged that are also blocked.
- LOG_ERR
Packets that have been logged and that can be considered “short”.
- S *device*
Set the logfile to be opened for reading state log records from or to *device*.
- t
Read the input file/device in the way performed by [tail\(1\)](#).
- v
Show TCP window, ack, and sequence fields
- x
Show the packet data in hex.
- X
Show the log header record data in hex.

文件

- /dev/ipl
- /dev/ipnat
- /dev/ipstate

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	network/ipfilter
Interface Stability	Committed

另请参见

[ipf\(1M\)](#), [ipfstat\(1M\)](#), [ipnat\(1M\)](#), [attributes\(5\)](#), [ipfilter\(5\)](#)

《Oracle Solaris 管理：IP 服务》

诊断

ipmon expects data that it reads to be consistent with how it should be saved and aborts if it fails an assertion which detects an anomaly in the recorded data.

引用名 ipmpstat – 显示 IPMP 子系统状态

用法概要 ipmpstat [-n] [-o *field[,...]* [-P]] -a|-g|-i|-p|-t

描述 ipmpstat 命令简要显示有关 IPMP 子系统的信息。如下所述，它支持五种不同的输出模式，每种模式都提供 IPMP 子系统的不同视图（地址、组、接口、探测器和目标）。每次调用至多可以指定一种输出模式，并且确保显示的信息前后一致。该命令还提供一种可解析输出格式，脚本可使用该格式来检查 IPMP 子系统的状态。仅需基本特权来调用 ipmpstat，但是探测模式需要所有特权。

选项 所支持的选项如下：

-a
显示 IPMP 数据地址信息（“地址”输出模式）。

-g
显示 IPMP 组信息（“组”输出模式）。

-i
显示 IP 接口信息（“接口”输出模式）。

-n
以数字形式显示 IP 地址，而不尝试将其解析为主机名。此选项可在任何输出模式下使用。

-o *field[,...]*
仅按顺序显示指定的输出字段。字段名称列表不区分大小写，以逗号分隔字段名称。如下所述，支持的字段名称取决于选定的输出模式。特殊字段名称 **all** 可用于显示给定输出模式的所有字段。

-p
显示 IPMP 探测器信息（“探测器”输出模式）。

-t
显示 IPMP 目标信息（“目标”输出模式）。

-P
如下所述，使用机器可解析格式显示。如果指定此选项，则必须使用 -o 选项指定明确的字段列表。

输出模式 ipmpstat 实用程序支持下面列出的输出模式。请注意，这些模式映射至上述部分选项。

地址模式

地址模式会显示系统上所有 IPMP 数据地址的状态。支持以下输出字段：

ADDRESS

与信息相关联的主机名（或 IP 地址）。请注意，由于可能存在重复的停用地址，因此必须将地址与 **GROUP** 组合在一起才能构成唯一的标识。对于给定 IPMP 组，如果存在重复地址，至多会显示一个地址，且始终优先活动地址。

STATE

地址的状态。如果地址是 **IFF_UP**，则状态为 **up**（活动）（请参见 [ifconfig\(1M\)](#)），如果地址不是 **IFF_UP**，则状态为 **down**（不活动）。

GROUP

托管地址的 **IPMP IP** 接口。

INBOUND

将接收此地址的数据包的底层 **IP** 接口。这种情况会因为响应 **IP** 接口故障等外部事件而变化。如果此字段为空，系统将不接受发送至此地址的 **IP** 数据包（例如，由于地址停用或 **IPMP** 组中没有活动的 **IP** 接口）。

OUTBOUND

将使用此源地址发送数据包的底层 **IP** 接口。这种情况会因为响应 **IP** 接口故障等外部事件而变化。如果此字段为空，系统将不会将此地址作为源来发送数据包（例如，由于地址停用或 **IPMP** 组中没有活动的 **IP** 接口）。

如果未指定 **-o**，则将显示所有输出字段。

组模式

组模式会显示系统中所有 **IPMP** 组的状态。支持以下输出字段：

GROUP

与信息相关联的 **IPMP IP** 接口名称。对于匿名组（请参见 [in.mpathd\(1M\)](#)），此字段将为空。

GROUPNAME

IPMP 组名。对于匿名组，此字段将为空。

STATE

组的状态：

- ok**（正常） 组中的所有接口均可用。
- degraded**（降级） 组中的部分（并非全部）接口可用。
- failed**（失败） 组中无任何接口可用。

FDT

基于探测器的故障检测时间。如果禁用基于探测器的故障检测，此字段将为空。

INTERFACES

组中底层 **IP** 接口的列表。此列表分为三个部分：

1. 活动接口会首先列出，并且不用任何方括号或圆括号括起来。活动接口是指系统正用来发送或接收数据流量的那些接口。
2. 随后会列出 **INACTIVE** 接口，并用圆括号括起来。**INACTIVE** 接口是那些正在工作但根据管理策略未在使用的接口。
3. 最后会列出不可用的接口并用方括号括起来。不可用接口是在当前配置中根本无法使用的那些接口（例如 **FAILED** 或 **OFFLINE**）。

如果未指定 `-o`，则将显示所有输出字段。

接口模式

接口模式会显示系统中由 `in.mpathd` 跟踪的所有 IP 接口的状态。支持以下输出字段：

INTERFACE

与信息相关联的 IP 接口名称。

ACTIVE

此值为 `yes` 或 `no`，具体取决于系统是否正在使用此 IP 接口传输 IP 数据通信。

GROUP

与 IP 接口相关联的 IPMP IP 接口。对于匿名组中的 IP 接口（请参见 [in.mpathd\(1M\)](#)），此字段将为空。

FLAGS

关于 IP 接口的各种信息：

- i 由于处于 **INACTIVE**（非活动）状态而不可用。
- s 标记为 **STANDBY**（备用）。
- m 指定为其 IPMP 组发送/接收 IPv4 多播。
- b 指定为其 IPMP 组发送/接收 IPv4 广播。
- M 指定为其 IPMP 组发送/接收 IPv6 多播。
- d 由于处于 **down**（关闭）状态而不可用。
- h 由于存在重复硬件地址时 `in.mpathd` 将状态转为 **OFFLINE**（脱机）而不可用。

LINK

基于链接的故障检测状态：

up（活动）

链接处于活动状态。

down（不活动）

链接处于不活动状态。

unknown（未知）

网络驱动程序未报告链接状态更改。

PROBE

基于探测器的故障检测状态：

ok（正常）

探测器没有检测到问题。

failed（失败）

探测器检测到故障。

unknown (未知)

由于适合的探测器目标未知，因此无法发送探测器。

disabled (禁用)

由于尚未配置唯一的 IP 测试地址，已禁用探测器。

STATE

接口的整体状态：

ok (正常)

接口已联机并且正基于配置的故障检测方法正常工作。

failed (失败)

接口已联机，但链接状态为 **down** (不活动)，或探测器状态为 **failed** (失败)。

offline (脱机)

接口已脱机。

unknown (未知)

接口已联机，但可能在工作，也可能未在工作，原因是配置的故障检测方法处于 **unknown** (未知) 状态。

如果未指定 **-o**，则将显示所有输出字段。

探测器模式

探测器模式会显示有关 **in.mpathd** 正发送的探测器的信息。与其他输出模式不同，此模式会一直运行，直到使用 **Ctrl-C** 明确终止。支持以下输出字段：

TIME

发送探测器的时间，此时间相对于启动 **ipmpstat** 的时间。如果探测器是在启动 **ipmpstat** 之前发送的，此时间将为负值。

PROBE

代表探测器的标识符。此标识符嵌入了指示探测器类型的前缀，后跟探测器的数字标识符。探测器类型的允许值包括：

i ICMP 探测器

t 过渡探测器

从活动接口发送 **ICMP** 探测器；对于由 **in.mpathd** 通过给定活动接口发送的每个 **IP** 探测器，其数字标识符会递增。此数字标识符与 **ICMP** 探测器数据包的 **icmp_seq** 字段相匹配，可供数据包监视工具用来进行更为详细的分析。

当 **IPMP** 组未在组中的任何接口上配置任何 **NOFAILOVER** 测试地址时，将从当前未用于接收数据包的所有接口发送过渡探测器。对于从给定接口发送的每个过渡探测器，其数据标识符会递增。探测器数据包的格式是实现的内部格式。

INTERFACE

发送探测器的 **IP** 接口。

TARGET

将 ICMP 探测器从活动接口发送到的目标的主机名（或 IP 地址），或者将过渡探测器发送到的 IP 接口的名称。

NETRTT

探测器的网络往返时间。这是从 IP 模块发送探测器到 IP 模块接收确认的时间。如果 `in.mpathd` 给出的结果是探测器已丢失，则此字段将为空。

RTT

探测器的总往返时间。这是从 `in.mpathd` 开始执行发送探测器的代码到处理完 `ack` 的时间。如果 `in.mpathd` 给出的结果是探测器已丢失，则此字段将为空。如果总往返时间中的峰值在网络往返时间中未出现，则表示本地系统自身已过载。

RTTAVG

通过 **INTERFACE** 到 **TARGET** 的平均往返时间。这样有助于识别较慢的目标。如果没有足够的数据来计算平均值，则此字段将为空。

RTTDEV

通过 **INTERFACE** 到 **TARGET** 的往返时间的标准偏差。这样有助于识别较为紧凑的目标。如果没有足够的数据来计算标准偏差，则此字段将为空。

如果未指定 `-o`，则会显示除 **RTTAVG** 和 **RTTDEV** 之外的所有字段。

目标模式

目标模式会显示 IPMP 探测器目标信息。支持以下输出字段：

INTERFACE

与信息相关联的 IP 接口名称。

MODE

探测器目标搜索模式：

routes	通过路由表发现的探测器目标。
multicast	通过多播 ICMP 探测器发现的探测器目标。
disabled	禁用所有基于探测器的故障检测。
transitive	故障检测通过过渡探测来实现，其中 IP 接口的运行状况通过探测组中的其他活动接口来确定。

TESTADDR

外出探测器中使用的源地址。正用于数据流量的活动接口和已明确配置了 **NOFAILOVER** 测试地址的接口将包含用于发送和接收 ICMP 探测器的主机名（或 IP 地址）。组中的其他所有接口将会显示发送探测器的接口的名称。请注意，如果活动 IP 接口同时配置了 IPv4 和 IPv6 测试地址，则会为每个测试地址分别显示探测器目标信息。

TARGETS

ICMP 探测器的探测器目标主机名（或 IP 地址）的空格分隔列表，或者过渡探测器的目标接口的空格分隔列表。IP 目标将按照启动顺序列出，如果找不到任何探测器目标，此字段将为空。

如果未指定 -o，则将显示所有输出字段。

输出格式

缺省情况下，ipmpstat 会为其输出模式使用用户友好型表格格式，其中每行包含有关给定对象的一个或多个信息字段，而此对象也由这些字段中的一个或多个字段来唯一标识。在这种格式中，标识这些字段的标题显示在表上方（以及满屏信息后面），这些字段以空格分隔，空字段使用 --（双连字符）表示，还将使用其他视觉辅助。如果某个字段的值无法确定，其值将显示为 "?"，并且会在标准错误中输出诊断消息。

机器可解析格式也使用表格格式，但此格式设计用于有效执行编程解析。具体来说，机器可解析格式与用户友好型格式的不同体现在以下方面：

- 未显示标题。
- 具有空值的字段不会生成输出，而不是显示 --。
- 字段以单个冒号 (:) 分隔，而不是不定数量的空格。
- 如果请求了多个字段，并且字段值中出现了字符 : 或反斜杠 (\)，则通过在字段前面添加 \ 前缀对其进行转义。

示例

示例1 获取特定接口的故障检测时间

以下代码使用机器可解析输出格式创建 ksh 函数，此函数可输出给定 IPMP IP 接口的故障检测时间：

```
getfdt() {
    ipmpstat -gP -o group,fdt | while IFS=: read group fdt; do
        [[ "$group" = "$1" ]] && { echo "$fdt"; return; }
    done
}
```

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

/usr/sbin/ipmpstat :

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）
机器可解析格式	Committed（已确定）
用户友好型格式	Not-an-Interface（不是接口）

/usr/sbin/ipmpstat 不是 Committed（已确定）接口。

另请参见

[if_mpadm\(1M\)](#)、[ifconfig\(1M\)](#)、[in.mpathd\(1M\)](#)、[attributes\(5\)](#)

引用名

用法概要

描述

选项

文件

ipnat – user interface to the NAT subsystem

ipnat [-CdFhlnRrsv] -f *filename*

The ipnat utility opens a specified file (treating - as stdin) and parses it for a set of rules that are to be added or removed from the IP NAT.

If there are no parsing problems, each rule processed by ipnat is added to the kernel's internal lists. Rules are appended to the internal lists, matching the order in which they appear when given to ipnat.

ipnat's use is restricted through access to /dev/ipauth, /dev/ipl, and /dev/ipstate. The default permissions of these files require ipnat to be run as root for all operations.

ipnat's use is restricted through access to /dev/ipnat. The default permissions of /dev/ipnat require ipnat to be run as root for all operations.

The following options are supported:

-C	Delete all entries in the current NAT rule listing (NAT rules).
-d	Turn debug mode on. Causes a hex dump of filter rules to be generated as it processes each one.
-F	Delete all active entries in the current NAT translation table (currently active NAT mappings).
-f <i>filename</i>	Parse specified file for rules to be added or removed from the IP NAT. <i>filename</i> can be stdin.
-h	Print number of hits for each MAP/Redirect filter.
-l	Show the list of current NAT table entry mappings.
-n	Prevents ipf from doing anything, such as making ioctl calls, which might alter the currently running kernel.
-R	Disable both IP address-to-hostname resolution and port number-to-service name resolution.
-r	Remove matching NAT rules rather than add them to the internal lists.
-s	Retrieve and display NAT statistics.
-v	Turn verbose mode on. Displays information relating to rule processing and active rules/table entries.

/dev/ipnat	Link to IP Filter pseudo device.
/dev/kmem	Special file that provides access to virtual address space.
/etc/ipf/ipnat.conf	Location of ipnat startup configuration file.

/usr/share/ipfilter/examples/ Contains numerous IP Filter examples.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	network/ipfilter
Interface Stability	Committed

另请参见 [ipf\(1M\)](#), [ipfstat\(1M\)](#), [ipnat\(4\)](#), [attributes\(5\)](#)

引用名	ippool – user interface to the IP Filter pools
用法概要	<pre>ippool -a [-dnv] [-m poolname] [-o role] -i ipaddr [/netmask] ippool -A [-dnv] [-m poolname] [-o role] [-S seed] [-t type] ippool -f file [-dnuv] ippool -F [-dv] [-o role] [-t type] ippool -h [-dv] [-m poolname] [-t type] ippool -l [-dv] [-m poolname] [-t type] ippool -r [-dnv] [-m poolname] [-o role] -i ipaddr [/netmask] ippool -R [-dnv] [-m poolname] [-o role] [-t type] ippool -s [-dtv] [-M core] [-N namelist]</pre>
描述	The ippool utility is used to manage information stored in the IP pools subsystem of IP Filter software. Configuration file information can be parsed and loaded into the kernel and currently configured pools can be removed, changed, or inspected. ippool's use is restricted through access to /dev/ippool. The default permissions of /dev/ippool require ippool to be run as root for all operations. The command line options used are divided into two sections: the global options and the instance-specific options. ippool's use is restricted through access to /dev/ipauth, /dev/ipl, and /dev/ipstate. The default permissions of these files require ippool to be run as root for all operations.
选项	ippool supports the option categories described below.
Global Options	The following global options are supported: -d Toggle debugging of processing the configuration file. -n Prevents ippool from doing anything, such as making ioctl calls, that would alter the currently running kernel. -v Turn verbose mode on.
Instance-Specific Options	The following instance-specific options are supported: -a Add a new data node to an existing pool in the kernel. -A Add a new (empty) pool to the kernel. -f file Read in IP pool configuration information from file and load it into the kernel. -F Flush loaded pools from the kernel.

- h Display a list of pools of the type: hash loaded in the kernel.
- l Display a list of pools of the type: tree loaded in the kernel.
- r Remove an existing data node from a pool in the kernel.
- R Remove an existing pool from within the kernel.
- s Display IP pool statistical information.

Other Options

The following, additional options are supported:

- i *ipaddr[/netmask]* Sets the IP address for the operation being undertaken with an all-one's mask or, optionally, a specific netmask, given in either dotted-quad notation or as a single integer.
- m *poolname* Sets the pool name for the current operation.
- M *core* Specify an alternative path to /dev/kmem from which to retrieve statistical information.
- N *namelist* Specify an alternative path to lookup symbol name information when retrieving statistical information.
- o *role* Sets the role with which this pool is to be used. Currently only ipf, auth, and count are accepted as arguments to this option.
- S *seed* Sets the hashing seed to the number specified. For use with hash-type pools only.
- t *type* Sets the type of pool being defined. Must be one of pool, hash, or group-map.
- u When parsing a configuration file, rather than load new pool data into the kernel, unload it.

文件

- /dev/ippool Link to IP Filter pseudo device.
- /dev/kmem Special file that provides access to virtual address space.
- /etc/ipf/ippool.conf Location of ippool startup configuration file.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	network/ipfilter
Interface Stability	Committed

另请参见

[ipf\(1M\)](#), [ipfstat\(1M\)](#), [ippool\(4\)](#), [attributes\(5\)](#)

引用名	ipqosconf – configure the IPQoS facility
用法概要	<pre>/usr/sbin/ipqosconf /usr/sbin/ipqosconf -a <i>conf_file</i> [-vs] /usr/sbin/ipqosconf -c /usr/sbin/ipqosconf -f /usr/sbin/ipqosconf -l /usr/sbin/ipqosconf -L</pre>
描述	The <code>ipqosconf</code> utility configures the Quality of Service facility of the Internet Protocol (IP). Only superusers can use this command. Without arguments, <code>ipqosconf</code> displays the actual IPQoS configuration. Configuration is not preserved across reboot. You must apply the configuration every time that the machine reboots. To apply the configuration early in the boot phase, you can populate the <code>/etc/inet/ipqosinit.conf</code> file, which is then read from the <code>svc:/network/initial:default</code> service.
选项	The following options are supported: -a <i>conf_file</i> Apply the configuration in <i>conf_file</i>. If the <i>conf_file</i> is <code>-</code>, <code>ipqosconf</code> reads from standard input. -c Populate the boot file with the current configuration. -f Flush the configuration. -l List the current applied configuration. -L List the current configuration in verbose mode. In addition to the information that the <code>-l</code> option provides, the <code>-L</code> option provides filters and classes configured through other means than the <code>ipqosconf</code> command. This option also provides the full set of filters that were created by <code>ipqosconf</code> by representing a multi-homed host in a configuration file -s Log messages to <code>syslog</code> during an <code>-a</code> operation. -v Toggle verbose mode during an <code>-a</code> operation. The <code>-v</code> option causes all messages to go to the console in addition to their normal destination. Messages intended to go to <code>syslog</code>, because the <code>-s</code> flag is set or because it is a log message, still go to <code>syslog</code> as well as the console.

Configuration File The configuration file is composed of a format version and a succession of configuration (action) blocks. There are different configuration blocks for each type of action that is being configured.

Format Version The first line of the configuration file specifies the format version contained in the configuration file.

The following entry specifies the format version:

```
fmt_version x.x
```

where *x.x* is the format version. 1.0 is the only supported version.

Configuration Blocks Following the format version, are a succession of configuration (action) blocks that are different for each type of action being configured. A configuration block always has the following structure:

```
action {
    name action_name
    module module_name
    params_clause | ""
    cf_clauses
}

action_name      ::= string
module_name      ::= ipgpc | dlcsmk | dscpmk | flowacct | tswtclmt |
                  tokenmt

params_clause    ::= params {
                    parameters
                    params_stats | ""
                  }

parameters       ::= prm_name_value parameters | ""

prm_name_value   ::= param_name param_value
```

Modules The *param_name* and the types of *param_value* are specific to a given module.

```
params_stats     ::= global_stats boolean

cf_clauses       ::= class_clause cf_clauses |
                  filter_clause cf_clauses | ""

class_clause     ::= class {
                    name class_name
                    next_action next_action_name
                    class_stats | ""
                  }
```

```

class_name      ::= string
next_action_name ::= string
class_stats     ::= enable_stats boolean
boolean         ::= TRUE | FALSE

```

```

filter_clause   ::= filter {
                        name filter_name
                        class class_name
                        parameters
                    }

```

```

filter_name     ::= string

```

There must be exactly one configuration block belonging to module `ipgpc`. The action must be named `ipgpc.classify`. All other actions should be reachable from `ipgpc` by way of parameters of type `action` or the `next_action` of a class.

The set of types that are used for parameters of the different modules are:

```

action          ::= string
protocol        ::= 1..255
port            ::= 1..65535
uint8           ::= 0..255
uint32          ::= 0..4294967296
int32           ::= -2147483648..2147483648
address         ::= <see the description section>
ifname          ::= <interface name recognized by SIOGLIFINDEX ioctl>
enum            ::= string | { string_list }
boolean         ::= TRUE | FALSE
integer_array   ::= { range_value_list }
map_index       ::= uint32
address         ::= ip_address | ip_node_name
user            ::= uid | username
uid             ::= 0..65535
username        ::= string
string_list     ::= string sl_entrys
sl_entrys       ::= ',' string sl_entrys | ""
range_value_list ::= range_value_entry range_value_entrys
range_value_entry ::= range ':' integer_array_value
range           ::= uint32 '-' uint32
integer_array_value ::= string | integer_array_number
integer_array_number ::= uint8 | uint32
range_value_entrys ::= ';' range_value_entry range_value_entrys | ""
ip_node_name    ::= string
ip_address      ::= v4_address | v6_address
v4_address      ::= v4_ip_address / v4_cidr_mask |
v4_ip_address

```

```

v4_cidr_mask      ::= 1-32
v6_address        ::= v6_ip_address / v6_cidr_mask |
v6_ip_address
v6_cidr_mask      ::= 1-128

```

METER module tokenmt configuration syntax:

```

red_action_name    action
yellow_action_name action
green_action_name  action
committed_rate     uint32
committed_burst    uint32
peak_rate          uint32
<if present this signifies that this will be a two rate meter, not
    a single rate meter>
peak_burst         uint32
<this is the 'peak' burst size for a two rate meter, but
    the 'excess' burst size for a single rate meter>
color_aware        boolean
color_map           integer_array
global_stats       boolean

```

METER module tswtclmt configuration syntax:

```

red_action_name    action
yellow_action_name action
green_action_name  action
committed_rate     uint32
peak_rate          uint32
window             uint32
global_stats       boolean

```

MARKER module dscpmk configuration syntax:

```

next_action        action
dscp_map           int_array
dscp_detailed_stats boolean
global_stats       boolean

```

MARKER module dlcosmk configuration syntax:

```

next_action        action
cos                map_index
global_stats       boolean

```

CLASSIFIER module ipgpc configuration syntax:

```

user              user
projid            int32
if_name           ifname

```

direction	enum { LOCAL_IN, LOCAL_OUT, FWD_IN, FWD_OUT}
protocol	protocol
dsfield	uint8
dsfield_mask	uint8
saddr	address
daddr	address
sport	port
dport	port
priority	uint32
precedence	uint32
ip_version	enum { V4, V6 }
global_stats	boolean

ACCOUNTING module flowacct configuration syntax:

next_action	action
timer	uint32
timeout	uint32
max_limit	uint32

Types	<i>action</i>	A string of characters with a matching action definition. The character string can be up to twenty three characters in length. To allow for spaces the string needs to be enclosed in quotes and cannot span lines. Two special actions are pre-defined and can not have an explicit action definition. The two pre-defined actions are continue and drop. continue causes the packet that is passed to it to continue normal processing. drop causes the packet that is passed to it to be dropped.
	<i>address</i>	A machine name or address recognized by getipnodebyname(3SOCKET) . If a machine name is specified, and ip_version has been defined, the query is done using that address family. If a machine name is not specified and ip_version has not been defined, the query is done using the AI_DEFAULT flag to getipnodebyname()(. AF_INET6 . .). CIDR address masks following an IP address are allowed. Specify the CIDR address masks as 1-32 (for v4) or 1-128 (for v6). CIDR addresses are disallowed for node names.
	<i>enum</i>	Either one of the supported values or comma delimited list of support values, enclosed in curly braces.
	<i>ifname</i>	A non-NULL, existing interface name recognized by the SIOGLIFINDEX socket ioctl.
	<i>integer_array</i>	A comma delimited set of <i>range/value</i> pairs, enclosed in curly braces.

Specify *range* in the format *x-y*, where *x* and *y* are integers that denote the range of array indexes to which the value applies. The minimum value for both *x* and *y* is 0. The maximum value for *x* is particular to the parameter. Any array indexes not referred to in the set of ranges are left at their previous value.

map_index A non-negative integer used as an index into any maps associated with a parameter of this type.

The maximum value of this type is dictated by the number of entries in the associated maps. The index starts at 0.

port Either a service name recognized by `getservbyname(3SOCKET)` or an integer 1-65535.

protocol Either a protocol name recognized by `getprotobyname(3SOCKET)` or an integer 1-255.

string A character string. Enclose *string* in quotes. *string* cannot span multiple lines.

user Either a valid user ID or username for the system that is being configured.

Parameters

The configuration file can contain the following parameters

color_aware A value of TRUE or FALSE, indicating whether or not the configured action takes account of the previous packet coloring when classifying.

color_map An integer array that defines which values of the dscp field correspond with which colors for when the *color_aware* parameter is set to TRUE.

committed_burst The committed burst size in bits.

committed_rate The committed rate in bits per second.

cos The value used to determine the underlying driver level priority applied to the packet which is defined in 802.1D.

daddr The destination address of the datagram.

direction The value used to build a filter matching only part of the traffic.

This parameter is of type enum with valid values of LOCAL_IN (local bound traffic), LOCAL_OUT (local sourced traffic), FWD_IN (forwarded traffic entering the system), and FWD_OUT (forwarded traffic exiting the system).

dport The destination port of the datagram.

dscp_detailed_stats	A value of TRUE or FALSE that determines whether detailed statistics are switched on for this dscp action. Specify TRUE to switch on or FALSE to switch off.
dscp_map	The <i>integer_array</i> that supplies the values that IP packets with a given dscp value have their dscp re-marked with. The existing value is used to index into the array where the new value is taken from. The array is of size 64, meaning valid indexes are 0-63 and valid values are also 0-63.
dsfield	The DS field of the IP datagram header. This is an 8-bit value, with each bit position corresponding with the same one in the header; this enables matches to be done on the CU bits. If you specify this parameter, you must also specify the dsfield_mask parameter.
dsfield_mask	The mask applied to the dsfield parameter to determine the bits against which to match. This is an 8-bit value, with each bit position corresponding with the same one in the dsfield parameter.
global_stats	A value of TRUE or FALSE to enable or disable the statistic collection for this action.
green_action_name	The action to be executed for packets that are deemed to be green.
if_name	The name of an interface recognized by the SIOGLIFINDEX ioctl. This parameter is of type ifname.
ip_version	This parameter is of type enum and has valid values of V4 and V6. If it is set to V4 only then only ipv4addresses are requested for a specified hostname. If it is set to V6, only ipv6 addresses are returned if there are any, otherwise v4 mapped v6 addresses are returned. If both V4 and V6 are specified, or if ip_version is not specified, then both ipv4 and ipv6 addresses are requested for a specified hostname.
max_limit	The maximum number of flow entries present at one time in the flowacct actions in the memory resident table.
next_action	The action to be executed when the current action is complete. This value can be either the name of an action defined in the configuration file, or one of the two special action types: drop and continue.
peak_burst	The peak burst size, for a two rate meter, or excess burst size, for a single rate meter, in bits.
peak_rate	The peak rate in bits per second.

precedence	An integer that is used to order filters. If there are two matching filters that have the same priority value, the one with the lower precedence value is the one matched. This parameter should be used because the order of the filters in a configuration file has no influence on their relative precedence.
priority	An integer that represents the relative priority of a filter. If there are two matching filters, the one with the higher priority value is the one matched. Multiple filters can have the same priority.
projid	The project ID of the process sending the data. This value is always -1 for received traffic.
protocol	The Upper Layer Protocol against which this entry is matched.
red_action_name	The action to be executed for packets that are determined to be red.
saddr	The source address of the datagram.
sport	The source port of the datagram.
timeout	The timeout in milliseconds after which flows are written to the accounting file.
timer	The period in milliseconds at which timed-out flows are checked for.
user	The user ID or username of the process sending the data. This value is always -1 for received traffic.
window	The window size in ms.
yellow_action_name	The action to be executed for packets that are determined to be yellow.

安全 None.

示例 示例 1 Sending All Traffic From eng to the AF 1 Class of Service

This example sends all traffic from eng to the AF 1 class of service. It is documented in four separate steps:

The following step creates a tokenmt action with three outcomes:

```
#meter for class 1.
action {
    name AF_CL1
    module tokenmt
    params{
        committed_rate 64
        committed_burst 75
        peak_burst 150
    }
}
```

示例 1 Sending All Traffic From eng to the AF 1 Class of Service (续)

```

        global_stats TRUE
        red_action_name drop
        yellow_action_name markAF12
        green_action_name markAF11
    }
}

```

The following step creates two dscpmk actions:

```

#class 1, low drop precedence.
action {
    name markAF11
    module dscpmk
    params{
        dscp_map {0-63:28}
        dscp_detailed_stats TRUE
        global_stats TRUE
        next_action acct1
    }
}

#class 1, medium drop precedence.
action {
    name markAF12
    module dscpmk
    params {
        dscp_map {0-63:30}
        dscp_detailed_stats TRUE
        global_stats TRUE
        next_action acct1
    }
}

```

The following step creates an accounting action:

```

#billing for transmitted class 1 traffic.
action {
    name acct1
    module flowacct
    params {
        timer 10
        timeout 30
        global_stats TRUE
    }
    max_limit 1024
    next_action continue
}

```

示例1 Sending All Traffic From eng to the AF 1 Class of Service (续)

The following step creates an ipgpc action:

```
#traffic from eng sent, traffic from ebay dropped.
action {
    name ipgpc.classify
    module ipgpc
    class {
        name from_eng
        enable_stats TRUE
        next_action AF_CL1
    }
    class {
        name from_ebay
        enable_stats TRUE
        next_action drop
    }

    filter {
        name from_eng
        saddr eng-subnet
        class from_eng
    }
    filter {
        name from_ebay
        saddr ebay-subnet
        class from_ebay
    }
}
```

文件

/etc/inet/ipqosinit.conf	Contains the IPQoS configuration loaded at boot time. If this file exists, it is read from the network/initial:default service.
/etc/inet/ipqosconf.1.sample	Sample configuration file for an application server
/etc/inet/ipqosconf.2.sample	Sample configuration file that meters the traffic for a specified application
/etc/inet/ipqosconf.3.sample	Sample configuration file that marks the ethernet headers of web traffic with a given user priority

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/ipqos

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

另请参见

[syslog\(3C\)](#), [getipnodebyname\(3SOCKET\)](#), [getprotobyname\(3SOCKET\)](#), [getservbyname\(3SOCKET\)](#), [attributes\(5\)](#), [dlcosmk\(7ipp\)](#), [dscpmk\(7ipp\)](#), [flowacct\(7ipp\)](#), [ipgpc\(7ipp\)](#), [ipqos\(7ipp\)](#), [tokenmt\(7ipp\)](#), [tswtclmt\(7ipp\)](#)

诊断

ipqosconf sends messages to syslog of facility user, severity notice when any changes are made to the IPQoS configuration.

Errors that occur during an ipqosconf operation send an error message to the console by default. For the application of a new configuration if the -s option is set then these messages are sent to syslog as facility user, severity error instead. If the -v option is present during an application then all error and change notification messages are sent to the console as well as their default destination.

引用名	ipsecalgs – configure the IPsec protocols and algorithms table				
用法概要	<pre>ipsecalgs ipsecalgs -l ipsecalgs -s ipsecalgs -a [-P protocol-number -p protocol-name] -k keylen-list [-i inc] [-K default-keylen] -b blocklen-list -n alg-names -N alg-number -m mech-name [-I initialization-vector_length] [-M MAC-length] [-S length-of-salt] [-F flags] [-f] [-s] ipsecalgs -P protocol-number -p protocol-name [-e exec-mode] [-f] [-s] ipsecalgs -r -p protocol-name [] -n alg-name [-s] ipsecalgs -r -p protocol-name [] -N alg-number [-s] ipsecalgs -R -P protocol-number [-s] ipsecalgs -R -p protocol-name [-s] ipsecalgs -e exec-mode -P protocol-number [-s] ipsecalgs -e exec-mode -p protocol-name [-s]</pre>				
描述	Use the ipsecalgs command to query and modify the IPsec protocol and algorithms stored in /etc/inet/ipsecalgs. You can use the ipsecalgs command to do the following: ▪ list the currently defined IPsec protocols and algorithms ▪ modify IPsec protocols definitions ▪ modify IPsec algorithms definitions Never edit the /etc/inet/ipsecalgs file manually. The valid IPsec protocols and algorithms are described by the ISAKMP DOI. See RFC 2407. In the general sense, a Domain of Interpretation (DOI) defines data formats, network traffic exchange types, and conventions for naming security-relevant information such as security policies or cryptographic algorithms and modes. For ipsecalgs, the DOI defines naming and numbering conventions for algorithms and the protocols they belong to. These numbers are defined by the Internet Assigned Numbers Authority (IANA). Each algorithm belongs to a protocol. Algorithm information includes supported key lengths, block or MAC length, and the name of the cryptographic mechanism corresponding to that algorithm. This information is used by the IPsec modules, ipsecesp(7P) and ipsecah(7P), to determine the authentication and encryption algorithms that can be applied to IPsec traffic. The following protocols are predefined: <table><tr><td>IPSEC_PROTO_ESP</td><td>Defines the encryption algorithms (transforms) that can be used by IPsec to provide data confidentiality.</td></tr><tr><td>IPSEC_PROTO_AH</td><td>Defines the authentication algorithms (transforms) that can be used by IPsec to provide authentication.</td></tr></table>	IPSEC_PROTO_ESP	Defines the encryption algorithms (transforms) that can be used by IPsec to provide data confidentiality.	IPSEC_PROTO_AH	Defines the authentication algorithms (transforms) that can be used by IPsec to provide authentication.
IPSEC_PROTO_ESP	Defines the encryption algorithms (transforms) that can be used by IPsec to provide data confidentiality.				
IPSEC_PROTO_AH	Defines the authentication algorithms (transforms) that can be used by IPsec to provide authentication.				

The mechanism name specified by an algorithm entry must correspond to a valid Solaris Cryptographic Framework mechanism. You can obtain the list of available mechanisms by using the [cryptoadm\(1M\)](#) command.

Applications can retrieve the supported algorithms and their associated protocols by using the functions [getipsecalgbyname\(3NSL\)](#), [getipsecalgbynum\(3NSL\)](#), [getipsecprotobyname\(3NSL\)](#) and [getipsecprotobynum\(3NSL\)](#).

Modifications to the protocols and algorithm by default update only the contents of the `/etc/inet/ipsecalgs` configuration file. In order for the new definitions to be used for IPsec processing, the changes must be communicated to the kernel using the `-s` option. See NOTES for a description of how the `ipsecalgs` configuration is synchronized with the kernel at system restart.

When invoked without arguments, `ipsecalgs` displays the list of mappings that are currently defined in `/etc/inet/ipsecalgs`. You can obtain the corresponding kernel table of protocols and algorithms by using the `-l` option.

选项

`ipsecalgs` supports the following options:

- a
Adds an algorithm of the protocol specified by the `-P` option. The algorithm name(s) are specified with the `-n` option. The supported key lengths and block sizes are specified with the `-k`, `-i`, and `-b` options.
- b
Specifies the block or MAC lengths of an algorithm, in bytes. Set more than one block length by separating the values with commas.
- e
Designates the execution mode of cryptographic requests for the specified protocol in the absence of cryptographic hardware provider. See [cryptoadm\(1M\)](#). *exec-mode* can be one of the following values:

sync	Cryptographic requests are processed synchronously in the absence of a cryptographic hardware provider. This execution mode leads to better latency when no cryptographic hardware providers are available
async	Cryptographic requests are always processed asynchronously in the absence of cryptographic hardware provider. This execution can improve the resource utilization on a multi-CPU system, but can lead to higher latency when no cryptographic hardware providers are available.

This option can be specified when defining a new protocol or to modify the execution mode of an existing protocol. By default, the sync execution mode is used in the absence of a cryptographic hardware provider.

- f
Used with the -a option to force the addition of an algorithm or protocol if an entry with the same name or number already exists.
- i
Specifies the valid key length increments in bits. This option must be used when the valid key lengths for an algorithm are specified by a range with the -k option.
- K
Specifies the default key lengths for an algorithm, in bits. If the -K option is not specified, the minimum key length will be determined as follows:
 - If the supported key lengths are specified by range, the default key length will be the minimum key length.
 - If the supported key lengths are specified by enumeration, the default key length will be the first listed key length.
- k
Specifies the supported key lengths for an algorithm, in bits. You can designate the supported key lengths by enumeration or by range.

 Without the -i option, -k specifies the supported key lengths by enumeration. In this case, *keylen-list* consists of a list of one or more key lengths separated by commas, for example:
 128,192,256

 The listed key lengths need not be increasing, and the first listed key length will be used as the default key length for that algorithm unless the -K option is used.

 With the -i option, -k specifies the range of supported key lengths for the algorithm. The minimum and maximum key lengths must be separated by a dash ('-') character, for example:
 32-448
- l
Displays the kernel algorithm tables.
- m
Specifies the name of the cryptographic framework mechanism name corresponding to the algorithm. Cryptographic framework mechanisms are described in the [cryptoadm\(1M\)](#) man page.
- N
Specifies an algorithm number. The algorithm number for a protocol must be unique. IANA manages the algorithm numbers. See RFC 2407.
- n
Specifies one or more names for an algorithm. When adding an algorithm with the -a option, *alg-names* contains a string or a comma-separated list of strings, for example:

des-cbs,des

When used with the -r option to remove an algorithm, *alg-names* contains one of the valid algorithm names.

-P

Adds a protocol of the number specified by *protocol-number* with the name specified by the -p option. This option is also used to specify an IPsec protocol when used with the -a and the -R options. Protocol numbers are managed by the IANA. See RFC 2407.

-p

Specifies the name of the IPsec protocol.

-R

Removes an IPsec protocol from the algorithm table. The protocol can be specified by number by using the -P option or by name by using the -p option. The algorithms associated with the protocol are removed as well.

-r

Removes the mapping for an algorithm. The algorithm can be specified by algorithm number using the -N option or by algorithm name using the -A option.

-s

Synchronizes the kernel with the contents of */etc/inet/ipsecalg*s. The contents of */etc/inet/ipsecalg*s are always updated, but new information is not passed on to the kernel unless the -s is used. See NOTES for a description of how the *ipsecalg*s configuration is synchronized with the kernel at system restart.

The following options allow optional parameters to be configured. These are currently only used for combined mode algorithms, that is, algorithms that provide encryption and authentication in a single operation.

-I

The length of the Initialization Vector (IV) in bytes. The default IV length is the same as the block length.

-M

The length of the MAC or ICV in bytes for combined mode algorithms.

-S

The number of bytes of salt needed by the algorithm. The salt needs to be provided by the key management mechanism.

-F

Algorithm flags. These influence the way in which the kernel handles security tasks, especially authentication, in the kernel. They are also used by [ipseckey\(1M\)](#) and [ipsecconf\(1M\)](#). Flags can be specified as a comma-separated list of tokens; see the example below. The following tokens are supported:

COUNTERMODE

The algorithm uses counter mode.

COMBINED

The algorithm provides encryption and authentication in the same operation.

CCM

The cryptographic framework mechanism needs a CK_AES_CCM_PARAMS structure.

GMAC

The cryptographic framework mechanism needs a CK_AES_GMAC_PARAMS structure.

GCM

The cryptographic framework mechanism needs a CK_AES_GCM_PARAMS structure.

CBC

This flag indicates the algorithm uses Cipher-block chaining. The cryptographic framework mechanism does not need a params structure. This is also the default, this flag can be omitted.

The algorithm flags can be displayed with the `-l` option.

示例**示例 1** Adding a Protocol for IPsec Encryption

The following example shows how to add a protocol for IPsec encryption:

```
example# ipsecalg -P 3 -p "IPSEC_PROTO_ESP"
```

示例 2 Adding the Blowfish Algorithm

The following example shows how to add the Blowfish algorithm:

```
example# ipsecalg -a -P 3 -k 32-488 -K 128 -i 8 -n "blowfish" \
-b 8 -N 7 -m CKM_BF_CBC
```

示例 3 Updating the Kernel Algorithm Table

The following example updates the kernel algorithm table with the currently defined protocol and algorithm definitions:

```
example# svcadm refresh ipsecalg
```

示例 4 Adding the AES Galois/Counter Mode (GCM) Algorithm

The following command adds this algorithm.

```
example# ipsecalg -a -P3 -k 128-256 -K 128 -i 64 -N 20 -b 16 \
-n "aes-gcm16,aes-gcm" -m CKM_AES_GCM -M 16 -I 8 -S 4 \
-F GCM, COMBINED, COUNTER
```

文件 /etc/inet/ipsecalgs
File that contains the configured IPsec protocols and algorithm definitions. Never edit this file manually.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [cryptoadm\(1M\)](#), [ipseconf\(1M\)](#), [ipseckey\(1M\)](#), [svcadm\(1M\)](#), [getipsecalgbyname\(3NSL\)](#), [getipsecpROTObyname\(3NSL\)](#), [ike.config\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [ipsec\(7P\)](#), [ipsecesp\(7P\)](#)

Piper, Derrell, RFC 2407, The Internet IP Security Domain of Interpretation for ISAKMP. Network Working Group. November 1998.

附注 When protocols or algorithm definitions that are removed or altered, services that rely upon these definitions can become unavailable. For example, if the IPSEC_PROTO_ESP protocol is removed, then IPsec cannot encrypt and decrypt packets.

Synchronization of the ipsecalgs configuration with the kernel at system startup is provided by the following [smf\(5\)](#) service:

```
svc:/network/ipsec/ipsecalgs:default
```

The IPsec services are delivered as follows:

```
svc:/network/ipsec/policy:default (enabled)
svc:/network/ipsec/ipsecalgs:default (enabled)
svc:/network/ipsec/manual-key:default (disabled)
svc:/network/ipsec/ike:default (disabled)
```

Services that are delivered disabled are delivered that way because the system administrator must create configuration files for those services before enabling them. See [ipseckey\(1M\)](#) and [ike.config\(4\)](#). The default policy for the policy service is to allow all traffic to pass without IPsec protection. See [ipseconf\(1M\)](#).

The correct administrative procedure is to create the configuration file for each service, then enable each service using [svcadm\(1M\)](#), as shown in the following example:

```
example# svcadm enable ipsecalgs
```

The service's status can be queried using the [svcs\(1\)](#) command.

If the ipsecalgs configuration is modified, the new configuration should be resynchronized as follows:

```
example# svcadm refresh ipsecalgs
```

Administrative actions on this service, such as enabling, disabling, refreshing, and requesting restart can be performed using [svcadm\(1M\)](#). A user who has been assigned the authorization shown below can perform these actions:

```
solaris.smf.manage.ipsec
```

See [auths\(1\)](#), [user_attr\(4\)](#), [rbac\(5\)](#).

The ipsecalgs [smf\(5\)](#) service does not have any user-configurable properties.

The [smf\(5\)](#) framework records any errors in the service-specific log file. Use any of the following commands to examine the logfile property:

```
example# svcs -l ipsecalgs  
example# svcprop ipsecalgs  
example# svccfg -s ipsecalgs listprop
```

This command requires `sys_ip_config` privilege to operate and thus can run in the global zone and in exclusive-IP zones. All shared-IP zones share the same available set of algorithms; however, you can use [ipseconf\(1M\)](#) to set up system policy that uses differing algorithms for various shared-IP zones. All exclusive-IP zones have their own set of algorithms.

引用名 ipsecconf – configure system wide IPsec policy

用法概要

```

/usr/sbin/ipsecconf
/usr/sbin/ipsecconf -a file [-q]
/usr/sbin/ipsecconf -c file
/usr/sbin/ipsecconf -d [-i tunnel-name] {index, tunnel-name, index}
/usr/sbin/ipsecconf -f [-i tunnel-name]
/usr/sbin/ipsecconf -F
/usr/sbin/ipsecconf -Fa file [-q]
/usr/sbin/ipsecconf -l [-i tunnel-name] [-n]
/usr/sbin/ipsecconf -L [-n]
/usr/sbin/ipsecconf -r file [-q]
```

描述

The ipsecconf utility configures the IPsec policy for a host or for one of its tunnels. Once the policy is configured, all outbound and inbound datagrams are subject to policy checks as they exit and enter the host or tunnel. For the host policy, if no entry is found, no policy checks will be completed, and all the traffic will pass through. For a tunnel, if no entry is found and there is at least one entry for the tunnel, the traffic will automatically drop. The difference in behavior is because of the assumptions about IPsec tunnels made in many implementations. Datagrams that are being forwarded will not be subjected to policy checks that are added using this command. See [ifconfig\(1M\)](#) and [dladm\(1M\)](#) for information on how to protect forwarded packets. Depending upon the match of the policy entry, a specific action will be taken.

This command can be run only by superuser.

Each entry can protect traffic in either one direction (requiring a pair of entries) or by a single policy entry which installs the needed symmetric sadb rules.

When the command is issued without any arguments, the list of file policy entries loaded are shown. To display the (spd p.e.s) use the -l option. Both will display the index number for the entry. To specify a single tunnel's SPD, use the -i option in combination with -l. To specify all SPDs, both host and for all tunnels, use -L.

Note, since one file policy entry (FPE) can generate multiple SPD pol entries (SPEs), the list of FPEs may not show all the actual entries. However, it is still useful in determining what rules have been added to get the spd into its current state.

You can use the -d option with the index to delete a given policy in the system. If the -d option removes an FPE entry that produces multiple SPEs, only then SPD with the same policy index as the FPE will be removed. This can produce a situation where there may be SPEs when there are no FPEs.

As with -l, -d can use the -i flag to indicate a tunnel. An alternate syntax is to specify a tunnel name, followed by a comma (,), followed by an index. For example, ip.tun0,1.

With no options, the entries are displayed in the order that they were added, which is not necessarily the order in which the traffic match takes place.

To view the order in which the traffic match will take place, use the `-l` option. The rules are ordered such that all bypass rules are checked first, then ESP rules, then AH rules. After that, they are checked in the order entered.

Policy entries are not preserved across system restarts. Permanent policy entries should be added to `/etc/inet/ipsecinit.conf`. This file is read by the following [smf\(5\)](#) service:

```
svc:/network/ipsec/policy
```

See **NOTES** for more information on managing IPsec security policy and **SECURITY** for issues in securing `/etc/inet/ipsecinit.conf`.

选项

`ipsecconf` supports the following options:

`-a file`

Add the IPsec policy to the system as specified by each entry in the file. An IPsec configuration file contains one or more entries that specify the configuration. Once the policy is added, all outbound and inbound datagrams are subject to policy checks.

Entries in the files are described in the [操作数](#) section below. Examples can be found in the [示例](#) section below.

Policy is latched for TCP/UDP sockets on which a [connect\(3SOCKET\)](#) or [accept\(3SOCKET\)](#) is issued. So, the addition of new policy entries may not affect such endpoints or sockets. However, the policy will be latched for a socket with an existing non-null policy. Thus, make sure that there are no preexisting connections that will be subject to checks by the new policy entries.

The feature of policy latching explained above may change in the future. It is not advisable to depend upon this feature.

The default behavior is to append new rules to the existing policy. If a new rule conflicts with an existing rule, an error is reported, the new rule will not be added.

To add a new rule that conflicts with an existing rule, the existing rule must be removed (see below) or the existing policy flushed. If the policy is flushed, IPsec will not protect any network traffic until a new policy is added.

The `-F` (flush) flag can be combined with `-a` to perform an atomic policy replacement; the existing policy will be replaced by the new policy described in the config file. By combining these two flags, there is not even a small window when the system is without a policy, which is the case when the flush and add commands are run sequentially.

`-c file`

Check the syntax of the configuration file and report any errors without making any changes to the policy. This option is useful when debugging configurations and when [smf\(5\)](#) reports a configuration error. See **SECURITY**.

-d *index*

Delete the host policy denoted by the index. The index is obtained by invoking `ipsecconf` without any arguments, or with the `-l` option. See DESCRIPTION for more information. Once the entry is deleted, all outbound and inbound datagrams affected by this policy entry will not be subjected to policy checks. Be advised that with connections for which the policy has been latched, packets will continue to go out with the same policy, even if it has been deleted. It is advisable to use the `-l` option to find the correct policy index.

-d *name,index*

Delete the policy entry denoted by *index* on a tunnel denoted by *name*. Since tunnels affect traffic that might originate off-node, latching does not apply as it does in the host policy case. Equivalent to: `-d index -i name`.

-f

Flush all the policies in the system. Constraints are similar to the `-d` option with respect to latching and host versus per-tunnel behavior.

-F

Flush all policies on all tunnels and also flush all host policies. See discussion of combining the `-F` and `-a` options, under `-a`, above.

-i *name*

Specify a tunnel interface name for use with the `-d`, `-f`, or `-l` flags.

-l

Listing of a single policy table, defaulting to the host policy. When `ipsecconf` is invoked without any arguments, a complete list of policy entries with indexes added by the user since boot is displayed. The current table can differ from the previous one if, for example, a multi-homed entry was added or policy reordering occurred, or if a single rule entry generates two `spd` rules. In the case of a multi-homed entry, all the addresses are listed explicitly. If a mask was not specified earlier but was instead inferred from the address, it will be explicitly listed here. This option is used to view policy entries in the correct order. The outbound and inbound policy entries are listed separately.

-L

Lists all policy tables, including host policy and all tunnel instances (including configured but unplumbed).

If `-i` is specified, `-L` lists the policy table for a specific tunnel interface.

-n

Show network addresses, ports, protocols in numbers. The `-n` option may only be used with the `-l` option.

-q

Quiet mode. Suppresses the warning message generated when adding policies.

`-r file`

Remove IPsec policy rules from the system as specified by each entry in *file*. The format of the file contents is the same as is specified with the `-a` option. The file could be created with `ipsecconf -l` and then modified with an editor.

操作数

Each policy entry contains three parts specified as follows:

```
{pattern} action {properties}
```

or

```
{pattern} action {properties} ["or" action {properties}]*
```

Every policy entry begins on a new line and can span multiple lines. If an entry exceeds the length of a line, you should split it only within a “braced” section or immediately before the first (left-hand) brace of a braced section. Avoid using the backslash character (`\`). See EXAMPLES.

The *pattern* section, as shown in the syntax above, specifies the traffic pattern that should be matched against the outbound and inbound datagrams. If there is a match, a specific *action* determined by the second argument will be taken, depending upon the *properties* of the policy entry.

If there is an `or` in the rule (multiple action-properties for a given pattern), a transmitter will use the first action-property pair that works, while a receiver will use any that are acceptable.

pattern and *properties* are name-value pairs where name and value are separated by a `<space>`, `<tab>` or `<newline>`. Multiple name-value pairs should be separated by `<space>`, `<tab>` or `<newline>`. The beginning and end of the pattern and properties are marked by `{` and `}` respectively.

Files can contain multiple policy entries. An unspecified name-value pair in the *pattern* will be considered as a wildcard. Wildcard entries match any corresponding entry in the datagram.

One thing to remember is that UDP port 500 is always bypassed regardless of any policy entries. This is a requirement for [in.iked\(1M\)](#) to work.

File can be commented by using a `#` as the first character. Comments may be inserted either at the beginning or the end of a line.

The complete syntax of a policy entry is:

```
policy ::= { <pattern1> } <action1> { <properties1> } |
        { <pattern2> } <action2> { <properties2> }
        [ 'or' <action2> { <properties2> } ]*

pattern1 ::= <pattern_name_value_pair1>*

pattern2 ::= <pattern_name_value_pair2>*
```

```

action1 ::= apply | permit | bypass | pass
action2 ::= bypass | pass | drop | ipsec

properties1 ::= {<prop_name_value_pair1>}
properties2 ::= {<prop_name_value_pair2>}

pattern_name_value_pair1 ::=
    saddr <address>/<prefix> |
    src <address>/<prefix> |
    srcaddr <address>/<prefix> |
    smask <mask> |
    sport <port> |
    daddr <address>/<prefix> |
    dst <address>/<prefix> |
    dstaddr <address>/<prefix> |
    dmask <mask> |
    dport <port> |
    ulp <protocol> |
    proto <protocol> |
    type <icmp-type> |
    type <number>-<number> |
    code <icmp-code>
    code <number>-<number>
    tunnel <interface-name> |
    negotiate <tunnel,transport>

pattern_name_value_pair2 ::=
    raddr <address>/<prefix> |
    remote <address>/<prefix> |
    rport <port> |
    laddr <address>/<prefix> |
    local <address>/<prefix> |
    lport <port> |
    ulp <protocol> |
    type <icmp-type> |
    type <number>-<number> |
    code <icmp-code> |
    code <number>-<number>
    proto <protocol> |
    tunnel <interface-name> |
    negotiate <tunnel,transport> |
    dir <dir_val2>

address ::= <IPv4 dot notation> | <IPv6 colon notation> |
    <String recognized by gethostbyname>|
    <String recognized by getnetbyname>

```

```

prefix ::= <number>

mask ::= <0xhexdigit[hexdigit]> | <0Xhexdigit[hexdigit]> |
        <IPv4 dot notation>

port ::= <number>| <String recognized by getservbyname>

protocol ::= <number>| <String recognized by getprotobyname>

prop_name_value_pair1 ::=
    auth_algs <auth_alg> |
    encr_algs <encr_alg> |
    encr_auth_algs <auth_alg> |
    sa <sa_val> |
    dir <dir_val1>

prop_name_value_pair2 ::=
    auth_algs <auth_alg> |
    encr_algs <encr_alg> |
    encr_auth_algs <auth_alg> |
    sa <sa_val>

auth_alg ::= <auth_algname> ['(' <keylen> ')']
auth_algname ::= any | md5 | hmac-md5 | sha | sha1 | hmac-sha |
                hmac-sha1 | hmac-sha256 | hmac-sha384 |
                hmac-sha512 | aes-gmac128 | aes-gmac192 |
                aes-gmac256 | <number>

encr_alg ::= <encr_algname> ['(' <keylen> ')']
encr_algname ::= any | aes | aes-cbc | des | des-cbc | 3des |
                3des-cbc | blowfish | blowfish-cbc | aes-ccm |
                aes-gcm | aes-none-gmac | <number>

keylen ::= <number> | <number>'..' | '..'<number> | <number>'..' \
<number>

sa_val ::= shared | unique

dir_val1 ::= out | in
dir_val2 ::= out | in | both

number ::= < 0 | 1 | 2 ... 9> <number>
icmp-type ::= <number> | unreachable | echo | echoresp | squench |
              redir | timex | paramprob | timest | timestrep |
              inforeq | inforep | maskreq | maskrep | unreachable6 |
              pkttoobig6 | timex6 | paramprob6 | echo6 | echoresp6 |
              router-sol6 | router-ad6 | neigh-sol6 | neigh-ad6 |

```

redir6

```
icmp-code ::= <number> | net-unr | host-unr | proto-unr | port-unr |
needfrag | srcfail | net-unk | host-unk | isolate |
net-prohib | host-prohib | net-tos | host-tos |
filter-prohib | host-prec | cutoff-prec |
no-route6 | adm-prohib6 | addr-unr6 | port-unr6 |
hop-limex6 | frag-re-timex6 | err-head6 | unrec-head6 |
unreq-opt6
```

Policy entries may contain the following (name value) pairs in the *pattern* field. Each (name value) pair may appear only once in given policy entry.

laddr/plen

local/plen

The value that follows is the local address of the datagram with the prefix length. Only plen leading bits of the source address of the packet will be matched. plen is optional. Local means destination on incoming and source on outgoing packets. The source address value can be a hostname as described in [getaddrinfo\(3SOCKET\)](#) or a network name as described in [getnetbyname\(3XNET\)](#) or a host address or network address in the Internet standard dot notation. See [inet_addr\(3XNET\)](#). If a hostname is given and [getaddrinfo\(3SOCKET\)](#) returns multiple addresses for the host, then policy will be added for each of the addresses with other entries remaining the same.

raddr/plen

remote/plen

The value that follows is the remote address of the datagram with the prefix length. Only plen leading bits of the remote address of the packet will be matched. plen is optional. Remote means source on incoming packets and destination on outgoing packets. The remote address value can be a hostname as described in [getaddrinfo\(3SOCKET\)](#) or a network name as described in [getnetbyname\(3XNET\)](#) or a host address or network address in the Internet standard dot notation. See [inet_addr\(3XNET\)](#). If a hostname is given and [getaddrinfo\(3SOCKET\)](#) returns multiple addresses for the host, then policy will be added for each of the addresses with other entries remaining the same.

src/plen

srcaddr/plen

saddr/plen

The value that follows is the source address of the datagram with the prefix length. Only plen leading bits of the source address of the packet will be matched. plen is optional.

The source address value can be a hostname as described in [getaddrinfo\(3SOCKET\)](#) or a network name as described in [getnetbyname\(3XNET\)](#) or a host address or network address in the Internet standard dot notation. See [inet_addr\(3XNET\)](#).

If a hostname is given and [getaddrinfo\(3SOCKET\)](#) returns multiple addresses for the host, then policy will be added for each of the addresses with other entries remaining the same.

daddr/plen
dest/plen
dstaddr/plen

The value that follows is the destination address of the datagram with the prefix length. Only *plen* leading bits of the destination address of the packet will be matched. *plen* is optional.

See *saddr* for valid values that can be given. If multiple source and destination addresses are found, then a policy entry that covers each source address-destination address pair will be added to the system.

smask

For IPv4 only. The value that follows is the source mask. If prefix length is given with *saddr*, this should not be given. This can be represented either in hexadecimal number with a leading 0x or 0X, for example, 0xffff0000, 0Xffff0000 or in the Internet decimal dot notation, for example, 255.255.0.0 and 255.255.255.0. The mask should be contiguous and the behavior is not defined for non-contiguous masks.

smask is considered only when *saddr* is given.

For both IPv4 and IPv6 addresses, the same information can be specified as a *slen* value attached to the *saddr* parameter.

dmask

Analogous to *smask*.

lport

The value that follows is the local port of the datagram. This can be either a port number or a string searched with a NULL proto argument, as described in [getservbyname\(3XNET\)](#)

rport

The value that follows is the remote port of the datagram. This can be either a port number or a string searched with a NULL proto argument, as described in [getservbyname\(3XNET\)](#)

sport

The value that follows is the source port of the datagram. This can be either a port number or a string searched with a NULL proto argument, as described in [getservbyname\(3XNET\)](#)

dport

The value that follows is the destination port of the datagram. This can be either a port number or a string as described in [getservbyname\(3XNET\)](#) searched with NULL proto argument.

proto ulp

The value that follows is the Upper Layer Protocol that this entry should be matched against. It could be a number or a string as described in [getprotobyname\(3XNET\)](#). If no *smask* or *plen* is specified, a *plen* of 32 for IPv4 or 128 for IPv6 will be used, meaning a host. If the *ulp* is *icmp* or *ipv6-icmp*, any action applying IPsec must be the same for all *icmp* rules.

type *num* or *num-num*

The value that follows is the ICMP type that this entry should be matched against. *type* must be a number from 0 to 255, or one of the appropriate `icmp-type` keywords. Also, *ulp* must be present and must specify either `icmp` or `ipv6-icmp`. A range of types can be specified with a hyphen separating numbers.

code *num* or *num-num*

The value that follows is the ICMP code that this entry should be matched against. The value following the keyword `code` must be a number from 0 to 254 or one of the appropriate `icmp-code` keywords. Also, *type* must be present. A range of codes can be specified with a hyphen separating numbers.

tunnel *name*

Specifies a tunnel network interface, as configured with [ifconfig\(1M\)](#). If a tunnel of *name* does not yet exist, the policy entries are added anyway, and joined with the tunnel state when it is created. If a tunnel is unplumbed, its policy entries disappear.

negotiate *tunnel***negotiate** *transport*

For per-tunnel security, specify whether the IPsec SAs protecting the traffic should be tunnel-mode SAs or transport-mode SAs. If transport-mode SAs are specified, no addresses can appear in the policy entry. Transport-mode is backward compatible with Solaris 9, and tunnel IPsec policies configured with [ifconfig\(1M\)](#) will show up as transport mode entries here.

Policy entries may contain the following (name-value) pairs in the properties field. Each (name-value) pair may appear only once in a given policy entry.

auth_algs

An acceptable value following this implies that IPsec AH header will be present in the outbound datagram. Values following this describe the authentication algorithms that will be used while applying the IPsec AH on outbound datagrams and verified to be present on inbound datagrams. See RFC 2402.

This entry can contain either a string or a decimal number.

string

This should be either MD5 or HMAC-MD5 denoting the HMAC-MD5 algorithm as described in RFC 2403, and SHA1, or HMAC-SHA1 or SHA or HMAC-SHA denoting the HMAC-SHA algorithm described in RFC 2404. You can use the [ipsecalgs\(1M\)](#) command to obtain the complete list of authentication algorithms.

The string can also be ANY, which denotes no-preference for the algorithm. Default algorithms will be chosen based upon the SAs available at this time for manual SAs and the key negotiating daemon for automatic SAs. Strings are not case-sensitive.

number
A number in the range 1-255. This is useful when new algorithms can be dynamically loaded.

If *auth_algs* is not present, the AH header will not be present in the outbound datagram, and the same will be verified for the inbound datagram.

encr_algs
An acceptable value following this implies that IPsec ESP header will be present in the outbound datagram. The value following this describes the encryption algorithms that will be used to apply the IPsec ESP protocol to outbound datagrams and verify it to be present on inbound datagrams. See RFC 2406.

This entry can contain either a string or a decimal number. Strings are not case-sensitive.

string
Can be one of the following:

string value:	Algorithm Used:	See RFC:
DES or DES-CBC	DES-CBC	2405
3DES or 3DES-CBC	3DES-CBC	2451
BLOWFISH or BLOWFISH-CBC	BLOWFISH-CBC	2451
AES or AES-CBC	AES-CBC	2451
AES-CCM	AES-CCM	4309
AES-GCM	AES_GCM	4106

You can use the [ipsecalgs\(1M\)](#) command to obtain the complete list of authentication algorithms.

The value can be NULL, which implies a NULL encryption, pursuant to RFC 2410. This means that the payload will not be encrypted. The string can also be ANY, which indicates no-preference for the algorithm. Default algorithms will be chosen depending upon the SAs available at the time for manual SAs and upon the key negotiating daemon for automatic SAs. Strings are not case-sensitive.

number
A decimal number in the range 1-255. This is useful when new algorithms can be dynamically loaded.

encr_auth_algs
An acceptable value following *encr_auth_algs* implies that the IPsec ESP header will be present in the outbound datagram. The values following *encr_auth_algs* describe the

authentication algorithms that will be used while applying the IPsec ESP protocol on outbound datagrams and verified to be present on inbound datagrams. See RFC 2406. This entry can contain either a string or a number. Strings are case-insensitive.

string

Valid values are the same as the ones described for `auth_algs` above.

number

This should be a decimal number in the range 1-255. This is useful when new algorithms can be dynamically loaded.

If `encr_algs` is present and `encr_auth_algs` is not present in a policy entry, the system will use an ESP SA regardless of whether the SA has an authentication algorithm or not.

If `encr_algs` is not present and `encr_auth_algs` is present in a policy entry, null encryption will be provided, which is equivalent to `encr_algs` with NULL, for outbound and inbound datagrams.

If both `encr_algs` and `encr_auth_algs` are not present in a policy entry, ESP header will not be present for outbound datagrams and the same will be verified for inbound datagrams.

If both `encr_algs` and `encr_auth_algs` are present in a policy entry, ESP header with integrity checksum will be present on outbound datagrams and the same will be verified for inbound datagrams.

For `encr_algs`, `encr_auth_algs`, and `auth_algs` a key length specification may be present. This is either a single value specifying the only valid key length for the algorithm or a range specifying the valid minimum and/or maximum key lengths. Minimum or maximum lengths may be omitted.

dir

Values following this decides whether this entry is for outbound or inbound datagram. Valid values are strings that should be one of the following:

out

This means that this policy entry should be considered only for outbound datagrams.

in

This means that this policy entry should be considered only for inbound datagrams.

both

This means that this policy entry should be considered for both inbound and outbound datagrams

This entry is not needed when the action is “apply”, “permit” or “ipsec”. But if it is given while the action is “apply” or “permit”, it should be “out” or “in” respectively. This is mandatory when the action is “bypass”.

sa

Values following this decide the attribute of the security association. Value indicates whether a unique security association should be used or any existing SA can be used. If there is a policy requirement, SAs are created dynamically on the first outbound datagram using the key management daemon. Static SAs can be created using `ipseckey(1M)`. The values used here determine whether a new SA will be used/obtained. Valid values are strings that could be one of the following:

unique

Unique Association. A new/unused association will be obtained/used for packets matching this policy entry. If an SA that was previously used by the same 5 tuples, that is, {Source address, Destination address, Source port, Destination Port, Protocol (for example, TCP/UDP)} exists, it will be reused. Thus uniqueness is expressed by the 5 tuples given above. The security association used by the above 5 tuples will not be used by any other socket. For inbound datagrams, uniqueness will not be verified.

For tunnel-mode tunnels, unique is ignored. SAs are assigned per-rule in tunnel-mode tunnels. For transport-mode tunnels, unique is implicit, because the enforcement happens only on the outer-packet addresses and protocol value of either IPv4-in-IP or IPv6-in-IP.

shared

Shared association. If an SA exists already for this source-destination pair, it will be used. Otherwise a new SA will be obtained. This is the default.

This is mandatory only for outbound policy entries and should not be given for entries whose action is “bypass”. If this entry is not given for inbound entries, for example, when “dir” is in or “action” is permit, it will be assumed to be shared.

Action follows the pattern and should be given before properties. It should be one of the following and this field is mandatory.

ipsec

Use IPsec for the datagram as described by the properties, if the pattern matches the datagram. If ipsec is given without a dir spec, the pattern is matched to incoming and outgoing datagrams.

apply

Apply IPsec to the datagram as described by the properties, if the pattern matches the datagram. If apply is given, the pattern is matched only on the outbound datagram.

permit

Permit the datagram if the pattern matches the incoming datagram and satisfies the constraints described by the properties. If it does not satisfy the properties, discard the datagram. If permit is given, the pattern is matched only for inbound datagrams.

bypass

pass

Bypass any policy checks if the pattern matches the datagram. `dir` in the properties decides whether the check is done on outbound or inbound datagrams. All the `bypass` entries are checked before checking with any other policy entry in the system. This has the highest precedence over any other entries. `dir` is the only field that should be present when action is `bypass`.

`drop`

Drop any packets that match the pattern.

If the file contains multiple policy entries, for example, they are assumed to be listed in the order in which they are to be applied. In cases of multiple entries matching the outbound and inbound datagram, the first match will be taken. The system will reorder the policy entry, that is, add the new entry before the old entry, only when:

The level of protection is “stronger” than the old level of protection.

Currently, strength is defined as:

AH and ESP > ESP > AH

The standard uses of AH and ESP were what drove this ranking of “stronger”. There are flaws with this. ESP can be used either without authentication, which will allow cut-and-paste or replay attacks, or without encryption, which makes it equivalent or slightly weaker than AH. An administrator should take care to use ESP properly. See [ipsecesp\(7P\)](#) for more details.

If the new entry has `bypass` as action, `bypass` has the highest precedence. It can be added in any order, and the system will still match all the `bypass` entries before matching any other entries. This is useful for key management daemons which can use this feature to bypass IPsec as it protects its own traffic.

Entries with both AH (`auth_algs` present in the policy entry) and ESP (`encr_auth_algs` or `encr_auth_algs` present in the policy entry) protection are ordered after all the entries with AH and ESP and before any AH-only and ESP-only entries. In all other cases the order specified by the user is not modified, that is, newer entries are added at the end of all the old entries. See [示例](#).

A new entry is considered duplicate of the old entry if an old entry matches the same traffic pattern as the new entry. See [示例](#) for information on duplicates.

安全

If, for example, the policy file comes over the wire from an NFS mounted file system, an adversary can modify the data contained in the file, thus changing the policy configured on the machine to suit his needs. Administrators should be cautious about transmitting a copy of the policy file over a network.

To prevent non-privileged users from modifying the security policy, ensure that the configuration file is writable only by trusted users.

The configuration file is defined by a property of the policy [smf\(5\)](#) service. The default configuration file, is `/etc/inet/ipsecinit.conf`. This can be changed using the [svccprop\(1\)](#) command. See NOTES for more details.

The policy description language supports the use of tokens that can be resolved by means of a name service, using functions such as [gethostbyname\(3NSL\)](#). While convenient, these functions are only secure as the name service the system is configured to use. Great care should be taken to secure the name service if it is used to resolve elements of the security policy.

If your source address is a host that can be looked up over the network and your naming system itself is compromised, then any names used will no longer be trustworthy.

If the name switch is configured to use a name service that is not local to the system, bypass policy entries might be required to prevent the policy from preventing communication to the name service. See [nsswitch.conf\(4\)](#).

Policy is latched for TCP/UDP sockets on which a [connect\(3SOCKET\)](#) or [accept\(3SOCKET\)](#) has been issued. Adding new policy entries will not have any effect on them. This feature of latching may change in the future. It is not advisable to depend upon this feature.

The `ipsecconf` command can only be run by a user who has sufficient privilege to open the [pf_key\(7P\)](#) socket. The appropriate privilege can be assigned to a user with the Network IPsec Management profile. See [profiles\(1\)](#), [rbac\(5\)](#), [prof_attr\(4\)](#).

Make sure to set up the policies before starting any communications, as existing connections may be affected by the addition of new policy entries. Similarly, do not change policies in the middle of a communication.

Note that certain `ndd` tunables affect how policies configured with this tool are enforced; see [ipsecesp\(7P\)](#) for more details.

示例

示例 1 Protecting Outbound TCP Traffic With ESP and the AES Algorithm

The following example specified that any TCP packet from `spiderweb` to `arachnid` should be encrypted with AES, and the SA could be a shared one. It does not verify whether or not the inbound traffic is encrypted.

```
#
# Protect the outbound TCP traffic between hosts spiderweb
# and arachnid with ESP and use AES algorithm.
#
{
    laddr spiderweb
    raddr arachnid
    ulp tcp
    dir out
} ipsec {
```

示例 1 Protecting Outbound TCP Traffic With ESP and the AES Algorithm (续)

```

        encr_algs AES
    }

```

示例 2 Verifying Whether or Not Inbound Traffic is Encrypted

Example 1 does not verify whether or not the inbound traffic is encrypted. The entry in this example protects inbound traffic:

```

#
# Protect the TCP traffic on inbound with ESP/DES from arachnid
# to spiderweb
#
{
    laddr spiderweb
    raddr arachnid
    ulp tcp
    dir in
} ipsec {
    encr_algs AES
}

```

sa can be absent for inbound policy entries as it implies that it can be a shared one. Uniqueness is not verified on inbound. Note that in both the above entries, authentication was never specified. This can lead to cut and paste attacks. As mentioned previously, though the authentication is not specified, the system will still use an ESP SA with encr_auth_alg specified, if it was found in the SA tables.

示例 3 Protecting All Traffic Between Two Hosts

The following example protects both directions at once:

```

{
    laddr spiderweb
    raddr arachnid
    ulp tcp
} ipsec {
    encr_algs AES
}

```

示例 4 Authenticating All Inbound Traffic to the Telnet Port

This entry specifies that any inbound datagram to telnet port should come in authenticated with the SHA1 algorithm. Otherwise the datagram should not be permitted. Without this entry, traffic destined to port number 23 can come in clear. sa is not specified, which implies that it is shared. This can be done only for inbound entries. You need to have an equivalent entry to protect outbound traffic so that the outbound traffic is authenticated as well, remove the dir.

示例 4 Authenticating All Inbound Traffic to the Telnet Port (续)

```
#
# All the inbound traffic to the telnet port should be
# authenticated.
#
{
    lport telnet
    dir in
} ipsec {
    auth_algs sha1
}
```

示例 5 Verifying Inbound Traffic is Null-Encrypted

The first entry specifies that any packet with address host-B should not be checked against any policies. The second entry specifies that all inbound traffic from network-B should be encrypted with a NULL encryption algorithm and the MD5 authentication algorithm. NULL encryption implies that ESP header will be used without encrypting the datagram. As the first entry is bypass it need not be given first in order, as bypass entries have the highest precedence. Thus any inbound traffic will be matched against all bypass entries before any other policy entries.

```
#
# Make sure that all inbound traffic from network-B is NULL
# encrypted, but bypass for host-B alone from that network.
# Add the bypass first.
{
    raddr host-B
    dir in
} bypass {}

# Now add for network-B.
{
    raddr network-B/16
    dir in
} ipsec {
    encr_algs NULL
    encr_auth_algs md5
}
```

示例 6 Entries to Bypass Traffic from IPsec

The first two entries provide that any datagram leaving the machine with source port 53 or coming into port number 53 should not be subjected to IPsec policy checks, irrespective of any other policy entry in the system. Thus the latter two entries will be considered only for ports other than port number 53.

示例 6 Entries to Bypass Traffic from IPsec (续)

```
#
# Bypass traffic for port no 53
#
{lport 53} bypass {}
{rport 53} bypass {}
{raddr spiderweb } ipsec {encr_algs any sa unique}
```

示例 7 Protecting Outbound Traffic

```
#
# Protect the outbound traffic from all interfaces.
#
{raddr spiderweb dir out} ipsec {auth_algs any sa unique}
```

If the `gethostbyname(3XNET)` call for spiderweb yields multiple addresses, multiple policy entries will be added for all the source address with the same properties.

```
{
    laddr arachnid
    raddr spiderweb
    dir in
} ipsec {auth_algs any sa unique}
```

If the `gethostbyname(3XNET)` call for spiderweb and the `gethostbyname(3XNET)` call for arachnid yield multiple addresses, multiple policy entries will be added for each (saddr daddr) pair with the same properties. Use `ipsecconf -l` to view all the policy entries added.

示例 8 Bypassing Unauthenticated Traffic

```
#
# Protect all the outbound traffic with ESP except any traffic
# to network-b which should be authenticated and bypass anything
# to network-c
#
{raddr network-b/16 dir out} ipsec {auth_algs any}
{dir out} ipsec {encr_algs any}
{raddr network-c/16 dir out} bypass {} # NULL properties
```

Note that bypass can be given anywhere and it will take precedence over all other entries. NULL pattern matches all the traffic.

示例 9 Encrypting IPv6 Traffic with 3DES and MD5

The following entry on the host with the link local address `fe80::a00:20ff:fe21:4483` specifies that any outbound traffic between the hosts with IPv6 link-local addresses `fe80::a00:20ff:fe21:4483` and `fe80::a00:20ff:felf:e346` must be encrypted with 3DES and MD5.

示例 9 Encrypting IPv6 Traffic with 3DES and MD5 (续)

```
{
    laddr fe80::a00:20ff:fe21:4483
    raddr fe80::a00:20ff:felf:e346
    dir out
} ipsec {
    encr_algs 3DES
    encr_auth_algs MD5
}
```

示例 10 Verifying IPv6 Traffic is Authenticated with SHA1

The following two entries require that all IPv6 traffic to and from the IPv6 site-local network fec0:abcd::0/32 be authenticated with SHA1.

```
{raddr fec0:abcd::0/32} ipsec { auth_algs SHA1 }
```

示例 11 KeyLengths

```
# use aes at any key length
{raddr spiderweb} ipsec {encr_algs aes}

# use aes with a 192 bit key
{raddr spiderweb} ipsec {encr_algs aes(192)}

# use aes with any key length up to 192 bits
# i.e. 192 bits or less
{raddr spiderweb} ipsec {encr_algs aes(..192)}

# use aes with any key length of 192 or more
# i.e. 192 bits or more
{raddr spiderweb} ipsec {encr_algs aes(192..)}

#use aes with any key from 192 to 256 bits
{raddr spiderweb} ipsec {encr_algs aes(192..256)}

#use any algorithm with a key of 192 bits or longer
{raddr spiderweb} ipsec {encr_algs any(192..)}
```

示例 12 Correct and Incorrect Policy Entries

The following are examples of correctly formed policy entries:

```
{ raddr that_system rport telnet } ipsec { encr_algs 3des encr_auth_algs
sha1 sa shared}

{
    raddr that_system
    rport telnet
```

示例 12 Correct and Incorrect Policy Entries (续)

```

} ipsec {
    encr_algs 3des
    encr_auth_algs sha1
    sa shared
}

{ raddr that_system rport telnet } ipsec
  { encr_algs 3des encr_auth_algs sha1 sa shared}

{ raddr that_system rport telnet } ipsec
  { encr_algs 3des encr_auth_algs sha1 sa shared} or ipsec
  { encr_algs aes encr_auth_algs sha1 sa shared}

```

...and the following is an incorrectly formed entry:

```

{ raddr that_system rport telnet } ipsec
  { encr_algs 3des encr_auth_algs sha1 sa shared}
  or ipsec { encr_algs aes encr_auth_algs sha1 sa shared}

```

In the preceding, incorrect entry, note that the third line begins with “or ipsec”. Such an entry causes ipsecconf to return an error.

示例 13 Allowing Neighbor Discovery to Occur in the Clear

The following two entries require that all IPv6 traffic to and from the IPv6 site-local network fec0:abcd::0/32 be authenticated with SHA1. The second entry allows neighbor discovery to operate correctly.

```

{raddr fec0:abcd::0/32} ipsec { auth_algs SHA1 }
{raddr fec0:abcd::0/32 ulp ipv6-icmp type 133-137 dir both }
  pass { }

```

示例 14 Using “or”

The following entry allows traffic using the AES or Blowfish algorithms from the remote machine spiderweb:

```

{raddr spiderweb} ipsec {encr_algs aes} or ipsec {encr_algs blowfish}

```

示例 15 Configuring a Tunnel to be Backward-Compatible with Solaris 9

The following example is equivalent to “encr_algs aes encr_auth_algs md5” in [ifconfig\(1M\)](#):

```

{tunnel ip.tun0 negotiate transport} ipsec {encr_algs aes
                                           encr_auth_algs md5}

```


示例 16 Configuring a Tunnel to a VPN client with an Assigned Address

The following example assumes a distinct “inside” network with its own topology, such that a client’s default route goes “inside”.

```
# Unlike route(1m), the default route has to be spelled-out.
{tunnel ip.tun0 negotiate tunnel raddr client-inside/32
 laddr 0.0.0.0/0} ipsec {encr_algs aes encr_auth_algs sha1}
```

示例 17 Transit VPN router between Two Tunnelled Subnets and a Third

The following example specifies a configuration for a VPN router that routes between two tunnelled subnets and a third subnet that is on-link. Consider remote-site A, remote-site B, and local site C, each with a /24 address allocation.

```
# ip.tun0 between me (C) and remote-site A.
# Cover remote-site A to remote-site B.
{tunnel ip.tun0 negotiate tunnel raddr A-prefix/24 laddr
 B-prefix/24} ipsec {encr_algs 3des encr_auth_algs md5}

# Cover remote-site A traffic to my subnet.
{tunnel ip.tun0 negotiate tunnel raddr A-prefix/24 laddr
 C-prefix/24} ipsec {encr_algs 3des encr_auth_algs md5}

# ip.tun1 between me (C) and remote-site B.
# Cover remote-site B to remote-site A.
{tunnel ip.tun1 negotiate tunnel raddr B-prefix/24 laddr
 A-prefix/24} ipsec {encr_algs aes encr_auth_algs sha1}

# Cover remote-site B traffic to my subnet.
{tunnel ip.tun1 negotiate tunnel raddr B-prefix/24 laddr
 C-prefix/24} ipsec {encr_algs aes encr_auth_algs md5}
```

示例 18 Using Combined Mode Ciphers

Combined mode ciphers provide data privacy and message authentication in a single operation. They are treated as special versions of `encr_algs`. They provide message authentication without the need to specify `encr_auth_algs`. The two combined mode ciphers supported are:

`aes-ccm`

AES CCM Mode(Counter with CBC-MAC)

`aes-gcm`

AES GCM Mode (Galois/Counter)

The parameters used are the same as any other `encr_algs` value. In both examples, the number in the algorithm token indicates the length of the Integrity Check Vector (ICV). See

[ipsecalgs\(1M\)](#).

示例 18 Using Combined Mode Ciphers (续)

```
# simple example using transport mode
{laddr 192.168.99.2 raddr 192.168.99.3} ipsec
    {encr_algs aes-gcm sa shared}
# simple example using CCM mode and 128 bit keys
{laddr 192.168.99.100 raddr 192.168.99.200} ipsec
    {encr_algs aes-ccm(128) sa shared}
```

示例 19 Using AES GMAC

The AES GMAC algorithm is a hash algorithm that provides message authentication. An Integrity Check Vector (ICV) is calculated and transmitted as part of the authenticated packet.

The AES GMAC algorithm can be used in ESP mode, in which case the packet data and the ESP header are authenticated. When used with IPsec ESP, AES GMAC can only be specified as an encryption algorithm, even though it does not provide any encryption.

```
# simple example using AES GMAC and 128 bit keys for ESP
{laddr 192.168.99.100 raddr 192.168.99.200} ipsec
    {encr_algs aes-none-gmac(128) sa shared}
```

The above example is analogous to the following invalid example:

```
{laddr 192.168.99.100 raddr 192.168.99.200} ipsec
    {encr_algs null encr_auth_algs aes-gmac(128) sa shared}
```

When used with ESP, the `aes-none-gmac` algorithm takes the key length as an optional argument, the supported key lengths can be displayed using the [ipsecalgs\(1M\)](#) command.

The AES GMAC algorithm can be used in AH mode, in which case the whole packet, including the IP header, is authenticated. When used with IPsec AH, AES GMAC can only be specified as an authentication algorithm.

```
# simple example using AES GMAC and 128 bit keys for AH
{laddr 192.168.99.100 raddr 192.168.99.200} ipsec
    {auth_algs aes-gmac128 sa shared}
```

When used with AH, each key length has its own DOI number, the key length does not need to be specified as an argument.

文件

`/var/run/ipsecpolicy.conf`

Cache of IPsec policies currently configured for the system, maintained by `ipsecconf` command. Do not edit this file.

`/etc/inet/ipsecinit.conf`

File containing IPsec policies to be installed at system restart by the policy [smf\(5\)](#) service. See NOTES for more information.

/etc/inet/ipsecinit.sample
Sample input file for ipseconf.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[auths\(1\)](#), [profiles\(1\)](#), [svcprop\(1\)](#), [svcs\(1\)](#), [in.iked\(1M\)](#), [init\(1M\)](#), [ifconfig\(1M\)](#), [ipsecalgs\(1M\)](#), [ipseckey\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [gethostbyname\(3NSL\)](#), [accept\(3SOCKET\)](#), [connect\(3SOCKET\)](#), [gethostbyname\(3XNET\)](#), [getnetbyname\(3XNET\)](#), [getprotobyname\(3XNET\)](#), [getservbyname\(3XNET\)](#), [getaddrinfo\(3SOCKET\)](#), [socket\(3SOCKET\)](#), [ike.config\(4\)](#), [nsswitch.conf\(4\)](#), [prof_attr\(4\)](#), [user_attr\(4\)](#), [attributes\(5\)](#), [rbac\(5\)](#), [smf\(5\)](#), [ipsecah\(7P\)](#), [ipsecesp\(7P\)](#), [pf_key\(7P\)](#)

Glenn, R. and Kent, S. RFC 2410, The NULL Encryption Algorithm and Its Use With IPsec. The Internet Society. 1998.

Kent, S. and Atkinson, R. RFC 2402, IP Authentication Header. The Internet Society. 1998.

Kent, S. and Atkinson, R. RFC 2406, IP Encapsulating Security Payload (ESP). The Internet Society. 1998.

Madsen, C. and Glenn, R. RFC 2403, The Use of HMAC-MD5-96 within ESP and AH. The Internet Society. 1998.

Madsen, C. and Glenn, R. RFC 2404, The Use of HMAC-SHA-1-96 within ESP and AH. The Internet Society. 1998.

Madsen, C. and Doraswamy, N. RFC 2405, The ESP DES-CBC Cipher Algorithm With Explicit IV. The Internet Society. 1998.

Pereira, R. and Adams, R. RFC 2451, The ESP CBC-Mode Cipher Algorithms. The Internet Society. 1998.

Frankel, S. and Kelly, R. Glenn, The AES Cipher Algorithm and Its Use With IPsec. 2001.

诊断

Bad “string” on line N.
Duplicate “string” on line N.
 string refers to one of the names in pattern or properties. A Bad string indicates that an argument is malformed; a Duplicate string indicates that there are multiple arguments of a similar type, for example, multiple Source Address arguments.

Interface name already selected
 Dual use of -i *name* and *name,index* for an index.

Error before or at line *N*.

Indicates parsing error before or at line *N*.

Non-existent index

Reported when the *index* for delete is not a valid one.

spd_msg return: File exists

Reported when there is already a policy entry that matches the traffic of this new entry.

附注

IPsec manual keys are managed by the service management facility, [smf\(5\)](#). The services listed below manage the components of IPsec. These services are delivered as follows:

```
svc:/network/ipsec/policy:default (enabled)
svc:/network/ipsec/ipsecalgs:default (enabled)
svc:/network/ipsec/manual-key:default (disabled)
svc:/network/ipsec/ike:default (disabled)
```

The manual-key service is delivered disabled. The system administrator must create manual IPsec Security Associations (SAs), as described in [ipseckey\(1M\)](#), before enabling that service.

The policy service is delivered enabled, but without a configuration file, so that, as a starting condition, packets are not protected by IPsec. After you create the configuration file `/etc/inet/ipsecinit.conf`, as described in this man page, and refresh the service (`svcadm refresh`, see below), the policy contained in the configuration file is applied. If there is an error in this file, the service enters maintenance mode.

Services that are delivered disabled are delivered that way because the system administrator must create configuration files for those services before enabling them. See [ike.config\(4\)](#) for the `ike` service.

See [ipsecalgs\(1M\)](#) for the `ipsecalgs` service.

The correct administrative procedure is to create the configuration file for each service, then enable each service using [svcadm\(1M\)](#).

If the configuration needs to be changed, edit the configuration file then refresh the service, as follows:

```
example# svcadm refresh policy
```

The [smf\(5\)](#) framework will record any errors in the service-specific log file. Use any of the following commands to examine the `logfile` property:

```
example# svcs -l policy
example# svcprop policy
example# svccfg -s policy listprop
```

The following property is defined for the `policy` service:

```
config/config_file
```

This property can be modified using [svccfg\(1M\)](#) by users who have been assigned the following authorization:

```
solaris.smf.value.ipsec
```

See [auths\(1\)](#), [user_attr\(4\)](#), [rbac\(5\)](#).

The service needs to be refreshed using [svcadm\(1M\)](#) before the new property is effective. General non-modifiable properties can be viewed with the [svccprop\(1\)](#) command.

```
# svccfg -s ipsec/policy setprop config/config_file = /new/config_file
# svcadm refresh policy
```

Administrative actions on this service, such as enabling, disabling, refreshing, and requesting restart can be performed using [svcadm\(1M\)](#). A user who has been assigned the authorization shown below can perform these actions:

```
solaris.smf.manage.ipsec
```

The service's status can be queried using the [svcs\(1\)](#) command.

The ipsecconf command is designed to be managed by the policy [smf\(5\)](#) service. While the ipsecconf command can be run from the command line, this is discouraged. If the ipsecconf command is to be run from the command line, the policy [smf\(5\)](#) service should be disabled first. See [svcadm\(1M\)](#).

引用名

ipseckey – manually manipulate an IPsec Security Association Database (SADB)

用法概要

```
ipseckey [-nvp]
ipseckey [-nvp] -f filename
ipseckey -c filename
ipseckey [-nvp] [delete | delete-pair | get] SA_TYPE {EXTENSION value...}
ipseckey [-np] [monitor | passive_monitor | pmonitor]
ipseckey [-nvp] flush {SA_TYPE}
ipseckey [-nvp] dump {SA_TYPE}
ipseckey [-nvp] save SA_TYPE {filename}
ipseckey [-nvp] -s filename
```

描述

The ipseckey command is used to manually manipulate the security association databases of the network security services, [ipsecah\(7P\)](#) and [ipsecesp\(7P\)](#). You can use the ipseckey command to set up security associations between communicating parties when automated key management is not available.

While the ipseckey utility has only a limited number of general options, it supports a rich command language. The user may specify requests to be delivered by means of a programmatic interface specific for manual keying. See [pf_key\(7P\)](#). When ipseckey is invoked with no arguments, it will enter an interactive mode which prints a prompt to the standard output and accepts commands from the standard input until the end-of-file is reached. Some commands require an explicit security association (“SA”) type, while others permit the SA type to be unspecified and act on all SA types.

ipseckey uses a PF_KEY socket and the message types SADB_ADD, SADB_DELETE, SADB_GET, SADB_UPDATE, SADB_FLUSH, and SADB_X_PROMISC. Thus, you must be a superuser to use this command.

ipseckey handles sensitive cryptographic keying information. Please read the [Security](#) section for details on how to use this command securely.

选项

-c [*filename*]

Analogous to the -f option (see following), except that the input is not executed but only checked for syntactical correctness. Errors are reported to `stderr`. This option is provided to debug configurations without making changes. See [SECURITY](#) and “Service Management Facility” for more information.

-f [*filename*]

Read commands from an input file, *filename*. The lines of the input file are identical to the command line language. The load command provides similar functionality. The -s option or the save command can generate files readable by the -f argument.

- n
Prevent attempts to print host and network names symbolically when reporting actions. This is useful, for example, when all name servers are down or are otherwise unreachable.
- p
Paranoid. Do not print any keying material, even if saving SAs. Instead of an actual hexadecimal digit, print an X when this flag is turned on.
- s *[filename]*
The opposite of the -f option. If '-' is given for a *filename*, then the output goes to the standard output. A snapshot of all current SA tables will be output in a form readable by the -f option. The output will be a series of add commands, but with some names not used. This occurs because a single name may often indicate multiple addresses.
- v
Verbose. Print the messages being sent into the PF_KEY socket, and print raw seconds values for lifetimes.

Commands

- add
Add an SA. Because it involves the transfer of keying material, it cannot be invoked from the shell, lest the keys be visible in `ps(1)` output. It can be used either from the interactive `ipseckey>` prompt or in a command file specified by the -f command. The add command accepts all extension-value pairs described below.
- update
Update SA lifetime, and in the cases of larval SAs (leftover from aborted automated key management), keying material and other extensions. Like add, this command cannot be invoked from the shell because keying material would be seen by the `ps(1)` command. It can be used either from the interactive `ipseckey>` prompt or in a command file specified by the -f command. The update command accepts all extension-value pairs, but normally is only used for SA lifetime updates.
- update-pair
As update, but apply the update to the SA and its paired SA, if there is one.
- delete
Delete a specific SA from a specific SADB. This command requires the spi extension, and the dest extension for IPsec SAs. Other extension-value pairs are superfluous for a delete message. If the SA to be deleted is paired with another SA, the SA is deleted and the paired SA is updated to indicate that it is now unpaired.
- delete-pair
Delete a specific SA from a specific SADB. If the SA is paired with another SA, delete that SA too. This command requires the spi extension and the dest extension for the IPsec SA, or its pair.
- get
Lookup and display a security association from a specific SADB. Like delete, this command only requires spi and dest for IPsec.

flush

Remove all SA for a given SA_TYPE, or all SA for all types.

monitor

Continuously report on any PF_KEY messages. This uses the SADB_X_PROMISC message to enable messages that a normal PF_KEY socket would not receive to be received. See [pf_key\(7P\)](#).

passive_monitor

Like monitor, except that it does not use the SADB_X_PROMISC message.

pmonitor

Synonym for passive_monitor.

dump

Will display all SAs for a given SA type, or will display all SAs. Because of the large amount of data generated by this command, there is no guarantee that all SA information will be successfully delivered, or that this command will even complete.

save

Is the command analog of the -s option. It is included as a command to provide a way to snapshot a particular SA type, for example, esp or ah.

help

Prints a brief summary of commands.

SA_TYPE**all**

Specifies all known SA types. This type is only used for the flush and dump commands. This is equivalent to having no SA type for these commands.

ah

Specifies the IPsec Authentication Header (“AH”) SA.

esp

Specifies the IPsec Encapsulating Security Payload (“ESP”) SA.

Extension Value Types

Commands like add, delete, get, and update require that certain extensions and associated values be specified. The extensions will be listed here, followed by the commands that use them, and the commands that require them. Requirements are currently documented based upon the IPsec definitions of an SA. Required extensions may change in the future. <number> can be in either hex (0xnnn), decimal (nnn) or octal (0nnn). <string> is a text string. <hexstr> is a long hexadecimal number with a bit-length. Extensions are usually paired with values; however, some extensions require two values after them.

spi <number>

Specifies the security parameters index of the SA. This extension is required for the add, delete, get and update commands.

`pair-spi <number>`

When `pair-spi` is used with the `add` or `update` commands, the SA being added or updated will be paired with the SA defined by `pair-spi`. A pair of SAs can be updated or deleted with a single command.

The two SAs that make up the pair need to be in opposite directions from the same pair of IP addresses. The command will fail if either of the SAs specified are already paired with another SA.

If the `pair-spi` token is used in a command and the SA defined by `pair-spi` does not exist, the command will fail. If the command was `add` and the pairing failed, the SA to be added will instead be removed.

`inbound | outbound`

These optional flags specify the direction of the SA. When the `inbound` or `outbound` flag is specified with the `add` command, the kernel will insert the new SA into the specified hash table for faster lookups. If the flag is omitted, the kernel will decide into which hash table to insert the new SA based on its knowledge the IP addresses specified with the `src` and `dst` extensions.

When these flags are used with the `update`, `delete`, `update-pair` or `get` commands, the flags provide a hint as to the hash table in which the kernel should find the SA.

`replay <number>`

Specifies the replay window size. If not specified, the replay window size is assumed to be zero. It is not recommended that manually added SAs have a replay window. This extension is used by the `add` and `update` commands.

`replay_value <number>`

Specifies the replay value of the SA. This extension is used by the `add` and `update` commands.

`state <string>|<number>`

Specifies the SA state, either by numeric value or by the strings “larval”, “mature”, “dying” or “dead”. If not specified, the value defaults to `mature`. This extension is used by the `add` and `update` commands.

`auth_alg <string>|<number>`

`authalg <string>|<number>`

Specifies the authentication algorithm for an SA, either by numeric value, or by strings indicating an algorithm name. Current authentication algorithms include:

 HMAC-MD5

 md5, hmac-md5

 HMAC-SH-1

 sha, sha-1, hmac-sha1, hmac-sha

HMAC-SHA-256

sha256, sha-256, hmac-sha256, hmac-sha-256

HMAC-SHA-384

sha384, sha-384, hmac-sha384, hmac-sha-384

HMAC-SHA-512

sha512, sha-512, hmac-sha512, hmac-sha-512

Often, algorithm names will have several synonyms. This extension is required by the `add` command for certain SA types. It is also used by the `update` command.

Use the [ipsecalgs\(1M\)](#) command to obtain the complete list of authentication algorithms.

`encr_alg <string>|<number>`

`encr_alg <string>|<number>`

Specifies the encryption algorithm for an SA, either by numeric value, or by strings indicating an algorithm name. Current encryption algorithms include DES (“des”), Triple-DES (“3des”), Blowfish (“blowfish”), and AES (“aes”). This extension is required by the `add` command for certain SA types. It is also used by the `update` command.

Use the [ipsecalgs\(1M\)](#) command to obtain the complete list of encryption algorithms.

The next six extensions are lifetime extensions. There are two varieties, “hard” and “soft”. If a hard lifetime expires, the SA will be deleted automatically by the system. If a soft lifetime expires, an SADB_EXPIRE message will be transmitted by the system, and its state will be downgraded to dying from mature. See [pf_key\(7P\)](#). The `monitor` command to key allows you to view SADB_EXPIRE messages.

`idle_addtime <number>`

`idle_usetime <number>`

Specifies the number of seconds that this SA can exist if the SA is not used before the SA is revalidated. If this extension is not present, the default value is half of the `hard_addtime` (see below). This extension is used by the `add` and `update` commands.

`soft_bytes <number>`

`hard_bytes <number>`

Specifies the number of bytes that this SA can protect. If this extension is not present, the default value is zero, which means that the SA will not expire based on the number of bytes protected. This extension is used by the `add` and `update` commands.

`soft_addtime <number>`

`hard_addtime <number>`

Specifies the number of seconds that this SA can exist after being added or updated from a larval SA. An update of a mature SA does not reset the initial time that it was added. If this extension is not present, the default value is zero, which means the SA will not expire based on how long it has been since it was added. This extension is used by the `add` and `update` commands.

`soft_usetime <number>`

`hard_usetime <number>`

Specifies the number of seconds this SA can exist after first being used. If this extension is not present, the default value is zero, which means the SA will not expire based on how long it has been since it was added. This extension is used by the add and update commands.

`saddr address | name`

`srcaddr address | name`

`saddr6 IPv6 address`

`srcaddr6 IPv6 address`

`src address | name`

`src6 IPv6 address`

`srcaddr address` and `src address` are synonyms that indicate the source address of the SA. If unspecified, the source address will either remain unset, or it will be set to a wildcard address if a destination address was supplied. To not specify the source address is valid for IPsec SAs. Future SA types may alter this assumption. This extension is used by the add, update, get and delete commands.

`daddr <address>|<name>`

`dstaddr <address>|<name>`

`daddr6 <IPv6 address>|<name>`

`dstaddr6 <IPv6 address>|<name>`

`dst <addr>|<name>`

`dst6 <IPv6 address>|<name>`

`dstaddr <addr>` and `dst <addr>` are synonyms that indicate the destination address of the SA. If unspecified, the destination address will remain unset. Because IPsec SAs require a specified destination address and `spi` for identification, this extension, with a specific value, is required for the add, update, get and delete commands.

If a name is given, `ipseckey` will attempt to invoke the command on multiple SAs with all of the destination addresses that the name can identify. This is similar to how `ipsecconf` handles addresses.

If `dst6` or `dstaddr6` is specified, only the IPv6 addresses identified by a name are used.

`sport <portnum>`

`sport` specifies the source port number for an SA. It should be used in combination with an upper-layer protocol (see below), but it does not have to be.

`dport <portnum>`

`dport` specifies the destination port number for an SA. It should be used in combination with an upper-layer protocol (see below), but it does not have to be.

`encap <protocol>`

Identifies the protocol used to encapsulate NAT-traversal IPsec packets. Other NAT-traversal parameters (`nat_*`) are below. The only acceptable value for `<protocol>` currently is `udp`.

`proto <protocol number>`

`ulp <protocol number>`

`proto`, and its synonym `ulp`, specify the IP protocol number of the SA.

`nat_loc <address>|<name>`

If the local address in the SA (source or destination) is behind a NAT, this extension indicates the NAT node's globally-routable address. This address can match the SA's local address if there is a `nat_lport` (see below) specified.

`nat_rem <address>|<name>`

If the remote address in the SA (source or destination) is behind a NAT, this extension indicates that node's internal (that is, behind-the-NAT) address. This address can match the SA's local address if there is a `nat_rport` (see below) specified.

`nat_lport <portnum>`

Identifies the local UDP port on which encapsulation of ESP occurs.

`nat_rport <portnum>`

Identifies the remote UDP port on which encapsulation of ESP occurs.

`isrc <address>|<name>[/<prefix>]`

`innersrc <address>|<name>[/<prefix>]`

`isrc6 <address>|<name>[/<prefix>]`

`innersrc6 <address>|<name>[/<prefix>]`

`proxyaddr <address>|<name>[/<prefix>]`

`proxy <address>|<name>[/<prefix>]`

`isrc <address>[/<prefix>]` and `innersrc <address>[/<prefix>]` are synonyms. They indicate the inner source address for a tunnel-mode SA.

An inner-source can be a prefix instead of an address. As with other address extensions, there are IPv6-specific forms. In such cases, use only IPv6-specific addresses or prefixes.

Previous versions referred to this value as the proxy address. The usage, while deprecated, remains.

`idst <address>|<name>[/<prefix>]`

`innerdst <address>|<name>[/<prefix>]`

`idst6 <address>|<name>[/<prefix>]`

`innerdst6 <address>|<name>[/<prefix>]`

`idst <address>[/<prefix>]` and `innerdst <address>[/<prefix>]` are synonyms. They indicate the inner destination address for a tunnel-mode SA.

An inner-destination can be a prefix instead of an address. As with other address extensions, there are IPv6-specific forms. In such cases, use only IPv6-specific addresses or prefixes.

`innersport <portnum>`

`isport <portnum>`

`innersport` specifies the source port number of the inner header for a tunnel-mode SA. It should be used in combination with an upper-layer protocol (see below), but it does not have to be.

`innerdport` <portnum>

`idport` <portnum>

`innerdport` specifies the destination port number of the inner header for a tunnel-mode SA. It should be used in combination with an upper-layer protocol (see below), but it does not have to be.

`ipproto` <protocol number> `iuip` <protocol number>

`ipproto`, and its synonym `iuip`, specify the IP protocol number of the inner header of a tunnel-mode SA.

`authkey` <hexstring>

Specifies the authentication key for this SA. The key is expressed as a string of hexadecimal digits, with an optional / at the end, for example, 123/12. Bits are counted from the most-significant bits down. For example, to express three '1' bits, the proper syntax is the string "e/3". For multi-key algorithms, the string is the concatenation of the multiple keys. This extension is used by the add and update commands.

`encrkey` <hexstring>

Specifies the encryption key for this SA. The syntax of the key is the same as `authkey`. A concrete example of a multi-key encryption algorithm is 3des, which would express itself as a 192-bit key, which is three 64-bit parity-included DES keys. This extension is used by the add and update commands.

`reserved_bits` <number>

The last <number> bits of the `encrkey` string are marked as reserved in the PF_KEY message. This option is only for testing certain encryption algorithms.

Certificate identities are very useful in the context of automated key management, as they tie the SA to the public key certificates used in most automated key management protocols. They are less useful for manually added SAs. Unlike other extensions, `srcidtype` takes two values, a *type*, and an actual *value*. The type can be one of the following:

`prefix`

An address prefix.

`fqdn`

A fully-qualified domain name.

`domain`

Domain name, synonym for `fqdn`.

`user_fqdn`

User identity of the form *user@fqdn*.

`mailbox`

Synonym for `user_fqdn`.

The *value* is an arbitrary text string that should identify the certificate.

srcidtype <*type*, *value*>

Specifies a source certificate identity for this SA. This extension is used by the `add` and `update` commands.

dstidtype <*type*, *value*>

Specifies a destination certificate identity for this SA. This extension is used by the `add` and `update` commands

Label extensions are used on Trusted Extensions to associate sensitivity labels with the traffic carried inside a security association. These extensions are not allowed unless Trusted Extensions is enabled.

label label

Defines the sensitivity label of traffic carried by this SA. Disallowed on systems not using Trusted Extensions.

outer-label label

Defines the sensitivity of the ciphertext traffic belonging to this SA; this label will appear in the outer packet header. Disallowed on systems not using Trusted Extensions. Incorrect use of this extension might allow label policy to be circumvented.

implicit-label label

Defines the sensitivity of the ciphertext traffic belonging to this SA and request that this SA not contain an explicit on-the-wire label. Disallowed on systems not using Trusted Extensions. Incorrect use of this extension might allow label policy to be circumvented.

Tunnel Mode versus Transport Mode SAs

An IPsec SA is a Tunnel Mode SA if the “proto” value is either 4 (ipip) or 41 (ip6) *and* there is an inner-address or inner-port value specified. Otherwise, the SA is a Transport Mode SA.

安全

Keying material is very sensitive and should be generated as randomly as possible. Some algorithms have known weak keys. IPsec algorithms have built-in weak key checks, so that if a weak key is in a newly added SA, the `add` command will fail with an invalid value.

The `ipseckey` command allows a privileged user to enter cryptographic keying information. If an adversary gains access to such information, the security of IPsec traffic is compromised. The following issues should be taken into account when using the `ipseckey` command.

1. Is the TTY going over a network (interactive mode)?
 - If it is, then the security of the keying material is the security of the network path for this TTY's traffic. Using `ipseckey` over a clear-text telnet or `rlogin` session is risky.
 - Even local windows might be vulnerable to attacks where a concealed program that reads window events is present.
2. Is the file accessed over the network or readable to the world (-f option)?
 - A network-mounted file can be sniffed by an adversary as it is being read.
 - A world-readable file with keying material in it is also risky.

3. The ipseckey command is designed to be managed by the manual-key [smf\(5\)](#) service. Because the [smf\(5\)](#) log files are world-readable, the ipseckey does not record any syntax errors in the log files, as these errors might include secret information.

If a syntax error is found when the manual-key [smf\(5\)](#) service is enabled, the service enters maintenance mode. The log file will indicate that there was a syntax error, but will not specify what the error was.

The administrator should use `ipseckey -c filename` from the command line to discover the cause of the errors. See [OPTIONS](#).

If your source address is a host that can be looked up over the network and your naming system itself is compromised, then any names used will not be trustworthy.

Security weaknesses often lie in misapplication of tools, not in the tools themselves. Administrators are urged to be cautious when using ipseckey. The safest mode of operation is probably on a console or other hard-connected TTY.

For further thoughts on this subject, see the afterward by Matt Blaze in Bruce Schneier's *Applied Cryptography: Protocols, Algorithms, and Source Code in C*.

Service Management Facility

IPsec manual keys are managed by the service management facility, [smf\(5\)](#). The services listed below manage the components of IPsec. These services are delivered as follows:

```
svc:/network/ipsec/policy:default (enabled)
svc:/network/ipsec/ipsecalg:default (enabled)
svc:/network/ipsec/manual-key:default (disabled)
svc:/network/ipsec/ike:default (disabled)
```

The manual-key service is delivered disabled. The system administrator must create manual IPsec Security Associations (SAs), as described in this man page, before enabling that service.

The policy service is delivered enabled, but without a configuration file, so that, as a starting condition, packets are not protected by IPsec. After you create the configuration file `/etc/inet/ipsecinit.conf` and refresh the service (`svcadm refresh`, see below), the policy contained in the configuration file is applied. If there is an error in this file, the service enters maintenance mode. See [ipseconf\(1M\)](#).

Services that are delivered disabled are delivered that way because the system administrator must create configuration files for those services before enabling them. See [ike.config\(4\)](#) for the `ike` service.

See [ipsecalgs\(1M\)](#) for the `ipsecalgs` service.

The correct administrative procedure is to create the configuration file for each service, then enable each service using [svcadm\(1M\)](#).

If the configuration needs to be changed, edit the configuration file then refresh the service, as follows:

```
example# svcadm refresh manual-key
```

Warning: To prevent ipseckey complaining about duplicate Associations, the ipseckey command flushes the Security Association Data Base (SADB) when the ipseckey command is run from [smf\(5\)](#), before adding any new Security Associations defined in the configuration file. This differs from the command line behavior where the SADB is not flushed before adding new Security Associations.

The [smf\(5\)](#) framework will record any errors in the service-specific log file. Use any of the following commands to examine the logfile property:

```
example# svcs -l manual-key
example# svcprop manual-key
example# svccfg -s manual-key listprop
```

The following property is defined for the manual-key service:

```
config/config_file
```

This property can be modified using [svccfg\(1M\)](#) by users who have been assigned the following authorization:

```
solaris.smf.value.ipsec
```

See [auths\(1\)](#), [user_attr\(4\)](#), [rbac\(5\)](#).

The service needs to be refreshed using [svcadm\(1M\)](#) before the new property is effective. General non-modifiable properties can be viewed with the [svcprop\(1\)](#) command.

```
# svccfg -s ipsec/manual-key setprop config/config_file = \
/new/config_file
# svcadm refresh manual-key
```

Administrative actions on this service, such as enabling, disabling, refreshing, and requesting restart can be performed using [svcadm\(1M\)](#). A user who has been assigned the authorization shown below can perform these actions:

```
solaris.smf.manage.ipsec
```

The service's status can be queried using the [svcs\(1\)](#) command.

The ipseckey command is designed to be run under [smf\(5\)](#) management. While the ipseconf command can be run from the command line, this is discouraged. If the ipseckey command is to be run from the command line, the manual-key [smf\(5\)](#) service should be disabled first. See [svcadm\(1M\)](#).

示例

示例1 Emptying Out All SAs

To empty out all SA:

```
example# ipseckey flush
```


示例 2 Flushing Out IPsec AH SAs Only

To flush out only IPsec AH SAs:

```
example# ipseckey flush ah
```

示例 3 Saving All SAs To Standard Output

To save all SAs to the standard output:

```
example# ipseckey save all
```

示例 4 Saving ESP SAs To The File /tmp/snapshot

To save ESP SAs to the file /tmp/snapshot:

```
example# ipseckey save esp /tmp/snapshot
```

示例 5 Deleting an IPsec SA

To delete an IPsec SA, only the SPI and the destination address are needed:

```
example# ipseckey delete esp spi 0x2112 dst 224.0.0.1
```

An alternative would be to delete the SA and the SAs pair if it has one:

```
example# ipseckey delete-pair esp spi 0x2112 dst 224.0.0.1
```

示例 6 Getting Information on an IPsec SA

Likewise, getting information on a SA only requires the destination address and SPI:

```
example# ipseckey get ah spi 0x5150 dst mypeer
```

示例 7 Adding or Updating IPsec SAs

Adding or updating SAs requires entering interactive mode:

```
example# ipseckey
ipseckey> add ah spi 0x90125 src me.domain.com dst you.domain.com \
    authalg md5 authkey 1234567890abcdef1234567890abcdef
ipseckey> update ah spi 0x90125 dst you.domain.com hard_bytes \
    16000000
ipseckey> exit
```

Adding two SAs that are linked together as a pair:

```
example# ipseckey
ipseckey> add esp spi 0x2345 src me.domain.com dst you.domain.com \
    authalg md5 authkey bde359723576fdea08e56cbe876e24ad \
    encrkey des encrkey be02938e7def2839
ipseckey> add esp spi 0x5432 src me.domain.com dst you.domain.com \
    authalg md5 authkey bde359723576fdea08e56cbe876e24ad \
    encrkey des encrkey be02938e7def2839 pair-spi 0x2345
```

示例 7 Adding or Updating IPsec SAs (续)

```
ipseckey> exit
```

示例 8 Adding an SA in the Opposite Direction

In the case of IPsec, SAs are unidirectional. To communicate securely, a second SA needs to be added in the opposite direction. The peer machine also needs to add both SAs.

```
example# ipseckey
ipseckey> add ah spi 0x2112 src you.domain.com dst me.domain.com \
        authalg md5 authkey bde359723576fdea08e56cbe876e24ad \
        hard_bytes 16000000
ipseckey> exit
```

示例 9 Monitoring PF_KEY Messages

Monitoring for PF_KEY messages is straightforward:

```
example# ipseckey monitor
```

示例 10 Using Commands in a File

Commands can be placed in a file that can be parsed with the -f option. This file may contain comment lines that begin with the “#” symbol. For example:

```
# This is a sample file for flushing out the ESP table and
# adding a pair of SAs.
```

```
flush esp
```

```
### Watch out! I have keying material in this file. See the
### SECURITY section in this manual page for why this can be
### dangerous .
```

```
add esp spi 0x2112 src me.domain.com dst you.domain.com \
    authalg md5 authkey bde359723576fdea08e56cbe876e24ad \
    encralg des encrkey be02938e7def2839 hard_usetime 28800
add esp spi 0x5150 src you.domain.com dst me.domain.com \
    authalg md5 authkey 930987dbe09743ade09d92b4097d9e93 \
    encralg des encrkey 8bd4a52e10127deb hard_usetime 28800
```

```
## End of file - This is a gratuitous comment
```

示例 11 Adding SAs for IPv6 Addresses

The following commands from the interactive-mode create an SA to protect IPv6 traffic between the site-local addresses

```
example # ipseckey
ipseckey> add esp spi 0x6789 src6 fec0:bbbb::4483 dst6 fec0:bbbb::7843\
        authalg md5 authkey bde359723576fdea08e56cbe876e24ad \
```

示例 11 Adding SAs for IPv6 Addresses (续)

```
encralg des encrkey be02938e7def2839 hard_usetime 28800
ipseckey>exit
```

示例 12 Linking Two SAs as a Pair

The following command links two SAs together, as a pair:

```
example# ipseckey update esp spi 0x123456 dst 192.168.99.2 \
pair-spi 0x654321
```

文件

/etc/inet/secret/ipseckeys
Default configuration file used at boot time. See “Service Management Facility” and SECURITY for more information.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[ps\(1\)](#), [svcprop\(1\)](#), [svcs\(1\)](#), [ipseconf\(1M\)](#), [ipsecalgs\(1M\)](#), [route\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [ike.config\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [ipsec\(7P\)](#), [ipsecah\(7P\)](#), [ipsecesp\(7P\)](#), [pf_key\(7P\)](#)

Schneier, B., Applied Cryptography: Protocols, Algorithms, and Source Code in C. Second ed. New York, New York: John Wiley & Sons, 1996.

诊断

The ipseckey command parses the configuration file and reports any errors. In the case of multiple errors, ipseckey reports as many of these as possible.

The ipseckey command does not attempt to use a COMMAND that has a syntax error. A COMMAND might be syntactically correct but can nevertheless generate an error because the kernel rejected the request made to [pf_key\(7P\)](#). This might occur because a key had an invalid length or because an unsupported algorithm was specified.

If there are any errors in the configuration file, ipseckey reports the number of valid COMMANDS and the total number of COMMANDS parsed.

Parse error on line N.

If an interactive use of ipseckey would print usage information, this would print instead. Usually preceded by another diagnostic. Because COMMANDS can cover more than a single line in the configuration file by using the backslash character to delimit lines, its not always possible to pinpoint in the configuration file the exact line that caused the error.

Unexpected end of command line.

An additional argument was expected on the command line.

Unknown

A value for a specific extension was unknown.

Address type *N* not supported.

A name-to-address lookup returned an unsupported address family.

***N* is not a bit specifier**

bit length *N* is too big for

string is not a hex string

Keying material was not entered appropriately.

Can only specify single

A duplicate extension was entered.

Don't use extension for *<string>* for *<command>*.

An extension not used by a command was used.

One of the entered values is incorrect: Diagnostic code *NN*: *<msg>*

This is a general invalid parameter error. The diagnostic code and message provides more detail about what precise value was incorrect and why.

附注

In spite of its IPsec-specific name, ipseckey is analogous to [route\(1M\)](#), in that it is a command-line interface to a socket-based administration engine, in this case, PF_KEY. PF_KEY was originally developed at the United States Naval Research Laboratory.

To have machines communicate securely with manual keying, SAs need to be added by all communicating parties. If two nodes wish to communicate securely, both nodes need the appropriate SAs added.

In the future ipseckey may be invoked under additional names as other security protocols become available to PF_KEY.

This command requires sys_ip_config privilege to operate and thus can run in the global zone and in exclusive-IP zones. The global zone can set up security associations with ipseckey to protect traffic for shared-IP zones on the system.

引用名	iscsiadm – 允许管理 iSCSI 启动器
用法概要	<code>iscsiadm subcommand direct-object [options] [operand]</code>
描述	<code>iscsiadm</code> 命令允许在主机上管理 iSCSI (Internet SCSI) 启动器。<code>iscsiadm</code> 作为一组子命令实现，许多子命令都具有其自己的选项，并在其各自的部分中进行说明。“选项”下面介绍了不与特定子命令关联的选项。 仅当下列服务处于联机状态时，<code>iscsiadm</code> 才会运行： <pre>svc:/network/iscsi/initiator:default</pre> <code>iscsiadm</code> 命令支持下列子命令，下面的小节中详细介绍了这些子命令： <code>add</code> 向对象添加元素。 <code>list</code> 列出对象的元素。 <code>modify</code> 修改对象的属性。 <code>remove</code> 删除对象的元素。 <code>iscsiadm</code> 子命令基于 <i>direct-object</i> 运行。各个子命令的相应部分中分别对这些子命令进行了说明。 <code>iscsiadm</code> 命令支持 Internet 存储名称服务 (Internet Storage Name Service, iSNS) 以便搜索 iSCSI 目标。此命令支持质询握手身份验证协议 (Challenge Handshake Authentication Protocol, CHAP) 以便进行身份验证。
add 子命令	<code>add</code> 子命令的语法如下： <pre># iscsiadm add direct_object [operands...]</pre> <code>add</code> 子命令添加下列 <i>direct_object</i>： <code>discovery-address discovery-address [...]</code> 向搜索地址列表添加目标。搜索地址（如在下面所示的语法中）是在 <code>SendTargets</code> 搜索会话中使用的 IP <i>address:port</i> 组合。通过使用此搜索方法，目标设备可告知启动器由该设备公开的各个目标的目标地址和目标名称。除非已在主机上启用搜索的 <code>SendTargets</code> 方法，否则不会尝试连接到目标。应使用 <code>modify</code> 子命令启用此方法。 <i>discovery-address</i> 参数的格式如下： <pre><IP address>[:port]</pre> 如果未指定 <i>port</i>，则将使用缺省值 3260。 <code>isns-server isns-server [...]</code> 向 iSNS 服务器地址列表添加 iSNS 服务器。iSNS 服务器地址（在下面所示的语法中指定）是在 iSNS 搜索会话中使用的 IP 地址-端口组合。通过使用 iSNS 搜索，iSNS 服务器可以向启动器提供有关门户以及与该启动器隶属于同一搜索域中的各个目标

的名称的信息。除非已在主机上启用搜索的 iSNS 方法，否则不会尝试连接到 iSNS 服务器。应使用 `modify` 子命令启用此方法，如下所述。

`isns-server` 参数的格式如下：

`IP_address[:port]`

如果未指定端口，则将使用缺省值 3205。

`static-config static_target [...]`

向静态配置目标的列表添加目标。除非已启用搜索的静态配置方法，否则不会尝试连接到目标。

`static_target` 参数的格式如下：

`<target-name>,<target address>[:port-number][,tpgt]`

`<target-name>` 最多可包含 223 个字符。

list 子命令

`list` 子命令的语法如下：

```
# iscsiadm list direct-object [options]
```

`list` 子命令显示下列 *direct-object* 的数据：

discovery

列出搜索方法及其当前激活状态（即启用或禁用）。搜索方法如下所示：

- iSNS（Internet Storage Name Service，Internet 存储名称服务）
- 静态
- SendTargets

initiator-node

列出主机上的启动器节点的信息。iSCSI 启动器节点表示逻辑 HBA，它是 iSCSI 目标的逻辑主机连接点。响应中列出的参数值是启动器的缺省参数设置。启动器的每个连接目标都具有不同于启动器节点上的参数值的参数值。

`static-config [static_target[,...]]`

列出指定目标的目标名称和地址，或者如果未指定静态目标，则列出所有静态搜索目标。

`target [-S] [-v] [target[,...]]`

列出目标的当前参数、连接状态以及用于目标搜索的方法。报告指定目标的信息，或者如果未指定目标，则列出搜索到的所有目标或已使用 `modify target` 子命令修改其参数的所有目标。

当与指定目标的 `-S` 选项配合使用时，此子命令将返回：

- 目标名称
- 逻辑单元号
- 供应商 ID
- 产品 ID

- OS 设备名称（例如，`/dev/rdisk/c0t2d0s0`）

-v 选项提供了更多详细信息，例如，当前登录参数、详细连接信息以及用于搜索目标的搜索方法。

返回 NA 作为搜索方法参数表示目标是使用 `iscsiadm modify target-param` 命令创建的，并且不会作为搜索目标存在。要删除此类目标，请使用 `iscsiadm remove target-param`。

`target-param [-v] target [...]`

列出目标的缺省参数和用户定义参数。

`discovery-address [-v] [discovery-address[, ...]]`

列出使用 `iscsiadm add discovery-address` 子命令添加的 `discovery-address` 对象。

当与 -v 选项配合使用时，列出位于指定 `discovery-address` 的所有已知目标。-v 选项返回一个或多个目标名称、零个或多个目标地址以及关联的目标门户组标记 (target portal group tag, TPGT)（如果适用）。

`isns-server [-v] [isns-server[, ...]]`

列出使用 `iscsiadm add isns-server` 子命令添加的 `isns-server` 对象。

当与 -v 选项配合使用时，此子命令列出位于指定 `isns-server` 地址的所有已知目标。-v 选项返回一个或多个目标名称、零个或多个目标地址以及关联的目标门户组标记 (target portal group tag, TPGT)（如果适用）。

modify 子命令

modify 子命令的语法如下：

```
# iscsiadm modify direct_object [options]
```

modify 子命令支持下列 *direct_object*：

`discovery [options]`

启用搜索方法可使用此方法启动搜索。禁用当前已启用的搜索方法不会影响使用该方法已搜索到的任何目标的连接。

modify discovery 的选项如下所示：

-i, -iSNS *enable|disable*
启用或禁用 iSNS 搜索。

-s, --static *enable|disable*
启用或禁用静态搜索。

-t, --sendtargets *enable|disable*
启用或禁用 SendTargets 搜索。

`initiator-node [options]`

修改启动器的属性。如果当前已连接目标，此操作即可成功。但是，在现有连接会话不再存在并且已建立新连接之前，修改后的参数集对此目标将不会生效。选项 -C 和 --CHAP-secret 在响应提示时需要输入 CHAP 密钥。

要在禁用 Solaris I/O 多路径功能（以前称为 Sun StorEdge Traffic Manager [STMS] 或 MPxIO）时进行 iSCSI 引导，您可以仅修改下列启动器节点选项：

- `-r, --radius-server`
- `-R, --radius-access`
- `-P, --radius-shared-secret`

要在启用 Solaris I/O 多路径功能时进行 iSCSI 引导，您可以仅修改下列启动器节点选项：

- `-h, --headerdigest`
- `-d, --datadigest`
- `-c, --configured-sessions`

`modify initiator-node` 的选项如下所示：

`-A, --node-alias <initiator node alias>`

修改启动器节点别名。最大长度为 223 个字符。

`-a, --authentication chap | none`

设置身份验证模式。

`-C, --CHAP-secret`

设置 CHAP 密钥值。没有缺省值。最大长度为 16 个字符；最小必需长度为 12 个字符。

`-c, --configured-sessions <num_sessions> | <IP address>[, <IP address>...]`

设置将为每个 iSCSI 目标创建的已配置 iSCSI 会话的数目。此功能应当与 [scsi_vhci\(7D\)](#) 中介绍的 Solaris I/O 多路径功能结合使用。

`-d, --datadigest none | CRC32`

设置是否启用 CRC32 以便检查 SCSI 数据传送。

`-H, --CHAP-name CHAP name`

指定 CHAP 用户名。如果未使用此选项，则会在初始化时将 CHAP 名称设置为启动器节点名称。将身份验证方法设置为 CHAP（请参见上面的

`-a/--authentication` 选项）时，使用命令 `iscsiadm list initiator-node` 可显示 CHAP 用户名。

`-h, --headerdigest none | CRC32`

设置是否启用 CRC32 以便检查 SCSI 包标头。

`-m, --max-connections number_connections`

修改 iSCSI 会话的最大连接数。缺省值为 1。各个会话的最大连接数为 65535。

`-N, --node-name <initiator node name>`

修改启动器节点名称。最多包含 223 个字符。

注—安装 Solaris 时，启动器节点名称设置为全局唯一值。更改此值可能会对 iSCSI 网络中的操作产生负面影响。

-P, --radius-shared-secret (exclusive)

设置 RADIUS 共享密钥。

-R, --radius-access enable | disable

设置是否将使用 RADIUS 服务器。

-r, --radius-server <IP address>[:<port>]

设置要使用的 radius 服务器的 IP 地址和端口。

-T, --tunable-param <<tunable-prop>=<value>, ...>

指定启动器节点连接的所有目标的一个或多个可调参数。

注 – 只有明确此参数在 iSCSI 网络中的影响的管理人员才能修改这些值。

支持的 tunable-prop 选项包括：

recv-login-rsp-timeout

会话登录响应时间

recv-login-rsp-timeout 选项指定 iSCSI 启动器将等待来自 iSCSI 目标的 iSCSI 会话登录请求响应的的时间。有效值介于 0 至 60*60 之间，缺省值为 60 秒。

conn-login-max

最长连接重试时间

如果在给定时间段内发生 IO 超时或连接故障，conn-login-max 选项允许 iSCSI 启动器重新建立到目标的连接。有效值介于 0 至 60*60 之间，缺省值为 180 秒。

polling-login-delay

登录重试时间间隔

polling-login-delay 选项指定在 iSCSI 启动器定位 IO 超时或连接故障时每次重试登录的时间间隔。有效值介于 0 至 60*60 之间，缺省值为 60 秒。

target-param [options] target

修改目标的参数。如果当前已连接目标，修改操作将成功，即使修改后的设置可能不会在几秒后生效也是如此。要确认这些设置是否处于活动状态，请使用 iscsiadm list target -v。如果指定目标不与任何搜索方法相关联，则会使用指定参数创建目标对象。使用此命令修改目标参数之后，将保存新参数，直到在此目标上使用 iscsiadm remove target-param 命令修改或删除这些参数为止。选项 -C 和 --CHAP-secret 在响应提示时需要输入 CHAP 密钥。

modify target-param 的选项如下所示：

-B, --bi-directional-authentication enable | disable

设置双向选项。如果设置为 enable，启动器将对指定目标执行双向身份验证。

-C, --CHAP-secret

设置目标的 CHAP 密钥值。没有缺省值。最大可接受长度为 16 个字符。

-c, --configured-sessions <num_sessions> | <IP address>[,<IP address>...]
 设置将为每个 iSCSI 目标创建的已配置 iSCSI 会话的数目。此功能应当与 [scsi_vhci\(7D\)](#) 中介绍的 Solaris I/O 多路径功能结合使用。

-d, --datadigest none | CRC32
 设置是为数据启用还是禁用 CRC32。

-H, --CHAP-name *CHAP name*
 设置 CHAP 用户名。如果未使用此选项，则会在初始化时将 CHAP 名称设置为目标名称。将身份验证方法设置为 CHAP（请参见上面的 `initiator-node` 直接对象下的 `-a/--authentication` 选项）时，使用命令 `iscsiadm list initiator-node` 可显示 CHAP 用户名。

-h, --headerdigest none | CRC32
 设置是为头启用还是禁用 CRC32。

-p, --login-param
 指定一个或多个登录参数设置。

注 – 只有明确此参数在 iSCSI 网络中的影响的管理人员才能修改这些值。

登录参数派生自 iSCSI 建议的标准 RFC 3720。有效值包括：

dataseqinorder	yes 或 no
defaulttime2retain	0–3600
defaulttime2wait	0–3600
firstburstlength	512 至 2^{24-1}
immediatedata	yes 或 no
initialr2t	yes 或 no
maxburstlength	512 至 2^{24-1}
datapduinorder	yes 或 no
maxoutstandingr2t	1 - 65535
maxrecvdataseglen	512 至 2^{24-1}

-T, --tunable-param <<tunable-prop>=<value>,...>
 指定启动器节点连接的所有目标的一个或多个可调参数。

注 – 只有明确此参数在 iSCSI 网络中的影响的管理人员才能修改可调参数值。

支持的 *tunable-prop* 选项包括：

recv-login-rsp-timeout
 会话登录响应时间

`recv-login-rsp-timeout` 选项指定 iSCSI 启动器将等待来自 iSCSI 目标的 iSCSI 会话登录请求响应的的时间。有效值介于 0 至 60*60 之间，缺省值为 60 秒。

`conn-login-max`
最长连接重试时间

如果在给定时间段内发生 IO 超时或连接故障，`conn-login-max` 选项允许 iSCSI 启动器重新建立到目标的连接。有效值介于 0 至 60*60 之间，缺省值为 180 秒。

`polling-login-delay`
登录重试时间间隔

`polling-login-delay` 选项指定在 iSCSI 启动器定位 IO 超时或连接故障时每次重试登录的时间间隔。有效值介于 0 至 60*60 之间，缺省值为 60 秒。

remove 子命令

remove 子命令的语法如下：

```
# iscsiadm remove direct_object
```

remove 子命令支持下列 *direct_object*：

`discovery-address` *discovery-address*, ...

从搜索地址列表中删除目标设备。搜索地址（如在下面所示的语法中）是在 SendTargets 搜索会话中使用的 IP 地址-端口组合。通过使用此搜索方法，目标设备可告知启动器由该设备公开的各个目标的目标地址和目标名称。如果当前已挂载搜索地址公开的任何目标或者设备上存在活动 I/O，则会返回“逻辑单元正在使用”错误，并且此操作将失败。如果关联设备未在使用，则会删除这些设备。

必须将 *discovery-address* 的格式设置为：

```
<IP address>[:<port>]
```

没有与此直接对象相关联的选项。

`isns-server` *isns-server*, ...

从 iSNS 服务器地址列表中删除 iSNS 服务器。iSNS 服务器地址（在下面所示的语法中指定）是在 iSNS 搜索会话中使用的 IP 地址-端口组合。通过使用 iSNS 搜索，iSNS 服务器可以向启动器提供有关门户以及与该启动器隶属于同一搜索域中的各个目标的名称的信息。如果当前已挂载通过 iSNS 搜索的任何目标或者设备上存在活动 I/O，则会返回“逻辑单元正在使用”错误，并且此操作将失败。如果关联设备未在使用，则会删除这些设备。

必须将 *isns-server* 的格式设置为：

```
IP_address[:port]
```

没有与此直接对象相关联的选项。

static-config static_target, ...

从静态搜索目标的列表中删除目标。如果正在删除当前已挂载的目标或者设备上存在活动 I/O，则会返回“逻辑单元正在使用”错误，并且此操作将失败。如果设备未在使用，则将删除此设备。

必须将 *static_target* 的格式设置为：

<target-name>, <target-address>[:port-number][, tpgt]

没有与此直接对象相关联的选项。

target-param target-name

删除 *target-name* 指定的目标。目标名称的格式如下：

<target-name>

没有与此直接对象相关联的选项。要在启用 Solaris I/O 多路径功能（以前称为 Sun StorEdge Traffic Manager [STMS] 或 MPxIO）时进行 iSCSI 引导，则不能删除目标。

正确使用搜索方法

请勿将目标配置为可供静态和动态搜索方法搜索。使用冗余搜索方法可能会导致在与 iSCSI 目标设备通信时降低性能。

选项

支持下列常规选项：

-V, --version 显示版本信息。停止后续参数的解释。

-, --help 显示帮助信息。可以在没有参数的 *iscsiadm* 命令、子命令或子命令直接对象组合后面使用。使用适用于您的输入的帮助信息进行响应。例如，如果输入以下内容：

```
# iscsiadm modify initiator-node --help
```

...*iscsiadm* 通过显示可用于子命令和直接对象的此组合的选项来进行响应。

示例

示例1 添加搜索地址

下面的命令使用 *add* 子命令添加搜索地址。

```
# iscsiadm add discovery-address 10.0.0.1:3260 10.0.0.2:3260
```

示例2 添加静态目标

下面的命令使用 *add* 子命令添加静态目标。

```
# iscsiadm add static-config \
iqn.1999-08.com.array:sn.01234567,10.0.0.1:3260
```

示例3 列出当前搜索设置

下面的命令使用 *list* 子命令列出当前搜索设置。

示例3 列出当前搜索设置 (续)

```
# iscsiadm list discovery
Discovery:
    Static: enabled
    Send Targets: disabled
    iSNS: enabled
```

示例4 获取详细搜索输出

下面的命令使用 `-v` 选项（其中一个命令使用 `v` 选项，另一个命令不使用 `v` 选项）和 `list` 子命令获取详细输出。

```
# iscsiadm list discovery-address
Discovery Address: 10.0.0.1:3260
Discovery Address: 10.0.0.2:3260

# iscsiadm list discovery-address -v 10.0.0.1:3260
Discovery Address: 10.0.0.1:3260
Target name: eui.210000203787d1f7
Target address: 10.0.0.1:3260
Target name: eui.210000203787a693
Target address: 10.0.0.1:3260
```

示例5 显示有关启动器的信息

下面的命令使用 `list` 子命令显示有关启动器的信息。

```
# iscsiadm list initiator-node
Initiator node name: iqn.1986-03.com.company.central.interopv20-1
Initiator node alias: interopv20-1
Login Parameters (Default/Configured):
    Header Digest: NONE/NONE
    Data Digest: NONE/NONE
Authentication Type: CHAP
    CHAP Name: iqn.1986-03.com.company.central.interopv20-1
RADIUS Server: NONE
RADIUS access: disabled
Tunable Parameters (Default/Configured):
    Session Login Response Time: 60/-
    Maximum Connection Retry Time: 180/-
    Login Retry Time Interval: 60/-
Configured Sessions: 1
```

示例6 显示静态配置信息

下面的命令使用 `list` 子命令显示有关静态配置的信息。

```
# iscsiadm list static-config
Static target: eui.210000203787a693,10.0.0.1:3260
```

示例7 显示目标信息

下面的命令显示如何使用 `list` 子命令和不同选项以便显示有关目标的信息。

iscsiadm list target

```
Target: iqn.2004-05.com.abcStorage:Tgt-1
Alias: -
TPGT: 12288
ISID: 4000002a0000
Connections: 1# iscsiadm list target -v iqn.2004-05.com.abcStorage:Tgt-1
```

```
Target: iqn.2004-05.com.abcStorage:Tgt-1
Alias: -
TPGT: 12288
ISID: 4000002a0000
Connections: 1
  CID: 0
    IP address (Local): 10.4.52.158:32803
    IP address (Peer): 10.4.49.70:3260
    Discovery Method: SendTargets
    Login Parameters (Negotiated):
      Data Sequence In Order: yes
      Data PDU In Order: yes
      Default Time To Retain: 20
      Default Time To Wait: 2
      Error Recovery Level: 0
      First Burst Length: 65536
      Immediate Data: yes
      Initial Ready To Transfer (R2T): yes
      Max Burst Length: 262144
      Max Outstanding R2T: 1
      Max Receive Data Segment Length: 65536
      Max Connections: 1
      Header Digest: NONE
      Data Digest: NONE
```

iscsiadm list target -S iqn.2004-05.com.abcStorage:Tgt-1

```
Target: iqn.2004-05.com.abcStorage:Tgt-1
Alias: -
TPGT: 12288
ISID: 4000002a0000
Connections: 1
  LUN: 6
    Vendor: ABCStorage
    Product: iSCSI Target
    OS Device Name: /dev/rdisk/c3t1d0s2
  LUN: 5
    Vendor: ABCStorage
    Product: iSCSI Target
    OS Device Name: /dev/rdisk/c3t0d0s2
```

示例8 显示目标参数信息

下面的命令使用 `list` 子命令显示特定目标的目标信息。

```
# iscsiadm list target-param -v iqn.2004-05.com.abcStorage:Tgt-1
Target: iqn.2004-05.com.abcStorage:Tgt-1
Alias: -
Bi-directional Authentication: disabled
Authentication Type: NONE
Login Parameters (Default/Configured):
    Data Sequence In Order: yes/-
    Data PDU In Order: yes/-
    Default Time To Retain: 20/-
    Default Time To Wait: 2/-
    Error Recovery Level: 0/-
    First Burst Length: 65536/-
    Immediate Data: yes/-
    Initial Ready To Transfer (R2T): yes/-
    Max Burst Length: 262144/-
    Max Outstanding R2T: 1/-
    Max Receive Data Segment Length: 65536/-
    Max Connections: 1/-
    Header Digest: NONE/-
    Data Digest: NONE/-
Tunable Parameters (Default/Configured):
    Session Login Response Time: 60/-
    Maximum Connection Retry Time: 180/-
    Login Retry Time Interval: 60/-
Configured Sessions: 1
```

示例9 启用静态搜索方法

下面的命令使用 `modify` 子命令启用静态搜索方法。

```
# iscsiadm modify discovery --static enable
```

示例10 设置Radius服务器的IP地址

下面的命令使用 `modify` 子命令设置 radius 服务器的 IP 地址，该服务器将用于 CHAP 身份验证。

```
# iscsiadm modify initiator --radius-server 10.0.0.1
```

示例11 设置启动器的节点名称

下面的命令使用 `modify` 子命令设置启动器节点的节点名称。

```
# iscsiadm modify initiator-node -N iqn.2004-10.com.SUN.host-1
```

示例 12 设置启动器节点的最大连接数

下面的命令使用 `modify` 子命令设置启动器节点的最大连接数。这允许在一个会话中存在多个连接。

```
# iscsiadm modify initiator-node -m 3
```

示例 13 更改目标参数

下面的命令使用 `modify` 子命令更改指定目标的目标参数的最大连接数。

```
# iscsiadm modify target-param -m 3 eui.210000203787a693
```

示例 14 删除搜索地址

下面的命令使用 `remove` 子命令删除搜索地址。

```
# iscsiadm remove discovery-address 10.0.0.1:3260
```

示例 15 删除目标参数

下面的命令使用 `remove` 子命令删除一组目标参数。

```
# iscsiadm remove target-param eui.210000203787a693
```

示例 16 修改最大连接数

下面的命令修改启动器的属性中的每个会话的最大连接数。修改后的值将用于所有会话中的所有目标。

```
# iscsiadm modify initiator-node --max-connections 4
```

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/storage/iscsi/iscsi-initiator
接口稳定性	Committed（已确定）

另请参见

[attributes\(5\)](#)、[iscsi\(7D\)](#)、[scsi_vhci\(7D\)](#)

《Oracle Solaris 管理：设备和文件系统》

引用名 isns – Internet 存储名称服务

用法概要 /usr/sbin/isns

描述 Internet 存储名称服务 (Internet Storage Name Service, iSNS) 在 IP 网络中提供针对 Internet SCSI (iSCSI) 和 Internet 光纤通道协议 (Internet Fibre Channel Protocol, iFCP) 设备的统一搜索服务。iSNS 使用客户机/服务器机制：服务器存储客户机的配置信息，并根据客户机的请求提供该信息。IETF RFC 4171 描述了服务器和客户机之间的协议。

本手册页简要介绍了 Solaris iSNS 服务器实现。当前实现不支持 iFCP 设备。

Solaris iSNS 服务器是作为守护进程 isns 实现的，该守护进程绑定到服务客户机请求熟知的端口 3205。该守护进程由服务管理工具 ([smf\(5\)](#)) 启动，并使用以下故障管理资源标识符 (Fault Management Resource Identifier, FMRI)：

```
svc:/network/isns_server
```

使用 [svcadm\(1M\)](#) 启用 isns。启用该服务意味着每次引导系统时它都自动启动和运行。可使用 [svcs\(1\)](#) 命令显示服务的状态。

可使用 [svccfg\(1M\)](#) 管理下面列出的服务属性。缺省值是根据 RFC 4171 和实现选择分配的。

data_store_location

配置数据存储位置。缺省位置是 /etc/isns/isnsdata.xml。

ESI_retry_threshold_count

实体状态查询重试阈值计数器。缺省计数为 3。

Management_SCNs_Enabled

用以确定是否启用管理状态更改通知的布尔值。缺省值为 yes。

Authorized_Control_Nodes

控制节点名称。

在更改属性值之后，必须使用 [svcadm\(1M\)](#) refresh，isns 才能识别新值。如果更改了 data_store_location 属性，必须输入 svcadm restart 命令才能使更改生效。

为了管理未分配给任何用户定义的搜索域的客户机，RFC 4171 定义了缺省搜索域、缺省域集和“缺省 DD/DDS”设置。服务器采取与缺省搜索域和缺省搜索域集相关的以下操作：

- 将未分配的客户机添加到缺省搜索域中。新注册的客户机或删除了其最后一个搜索域成员身份的客户机都被视为未分配的客户机。
- 将客户机分配给用户定义的搜索域时，服务器会从缺省搜索域中删除该客户机。
- 为了使管理员能够对缺省搜索域中的客户机进行搜索控制，允许通过管理操作激活或取消激活缺省搜索域集。
- 不允许通过管理操作将客户机添加到缺省搜索域中，也不允许通过管理操作将用户定义的搜索域添加到缺省搜索域集中。

- 缺省搜索域集的缺省状态为非活动状态。

isns 服务器支持允许您管理 isns 活动的某些 rbac(5) 授权。这些授权包括以下 auth_attr(4) 特权：

solaris.isnsmgr.write
此特权是执行下列操作时所必需的：创建搜索域或搜索域集、启用/禁用搜索域集，以及更改搜索域中 iSNS 客户机的分组或更改搜索域集中搜索域的分组。

solaris.smf.manage.isns
此特权是通过 smf(5) 管理 isns 服务器时所必需的。

solaris.smf.value.isns
此特权是更改与 isns 关联的 SMF 服务属性时所必需的。

"iSNS Server Management"（iSNS 服务器管理）配置文件（请参见 prof_attr(4)）中包括了上述所有授权。有关角色和授权的概述，请参见 rbac(5)。

选项

isns 守护进程不支持任何选项。

示例

示例1 启动 isns 服务器
以下命令可启动 isns 服务器。

svcadm enable svc:/network/isns_server

示例2 停止 isns 服务器
以下命令可停止 isns 服务器。

svcadm disable svc:/network/isns_server

示例3 更改 isns 属性
以下命令序列可更改 ESI_retry_threshold_count 属性的值。

**# svccfg -s svc:/network/isns_server setprop \
config/ESI_retry_threshold_count = 6
svcadm refresh svc:/network/isns_server**

文件

/usr/sbin/isns
iSNS 守护进程二进制文件。

属性

有关以下属性的说明，请参见 attributes(5)：

属性类型	属性值
可用性	service/storage/isns
接口稳定性	Committed（已确定）
标准	请参见 standards(5)。

另请参见

[svcs\(1\)](#)、[isnsadm\(1M\)](#)、[svcadm\(1M\)](#)、[svccfg\(1M\)](#)、[auth_attr\(4\)](#)、[prof_attr\(4\)](#)、[attributes](#)

附注

强烈建议您在更改服务属性之后重新启动服务器 (`svcadm restart`)。这使得服务器可以为现有客户机和新客户机应用统一的设置。

如 RFC 4171 中所述，控制节点不是管理服务器时所必需的。可以通过本地主机上的 [isnsadm\(1M\)](#) 命令接口来实现控制节点操作。例如，`isnsadm` 允许您创建搜索域和搜索域集并为其添加成员，以创建搜索域和搜索域集关联。

引用名	isnsadm – 管理 Internet 存储名称服务 (Internet Storage Name Service, iSNS) 服务器
用法概要	<code>isnsadm options</code> <code>isnsadm subcommand [subcommand_options] [operand]</code>
描述	<code>isnsadm</code> 命令是 Internet 存储名称服务 (Internet Storage Name Service, iSNS) 服务器的命令行接口。<code>isnsadm</code> 包括一组子命令（将分别在它们的各自部分中进行描述），每个子命令用于完成 iSNS 服务器管理功能中的一种。 对于将更改 iSNS 配置的任何操作都需要 <code>solaris.isnsmgr.write</code> 授权。请参见 isns(1M)。对于读取操作，该命令不需要特殊的授权。 <code>isnsadm</code> 有一组常规选项和一组特定于子命令的选项。第一个类别在“选项”部分中进行了描述，第二个类别在每个子命令描述上下文中进行了描述。
选项	所支持的选项如下： <code>-?, --help</code> 显示上下文帮助。停止解析任何后续参数。 <code>-V, --version</code> 显示版本信息。停止解析任何后续参数。
操作数	一个或多个 <code>isnsadm</code> 子命令使用下列操作数。 <i>iscsi-node-name</i> iSCSI 目标或 iSCSI 启动器符号名称。最大长度为 223 个字符的字符串。 <i>discovery-domain-name</i> 搜索域符号名称。最大长度为 256 个字符的字符串。 <i>discovery-domain-set-name</i> 搜索域集符号名称。最大长度为 256 个字符的字符串。
子命令	<code>isnsadm</code> 命令支持下面描述的子命令。
add-dd	<code>add-dd</code> 子命令用于将搜索域添加到搜索域集中。 <code>add-dd</code> 子命令的语法如下所示： <pre># isnsadm add-dd option discovery-domain-name, ...</pre> <code>add-dd</code> 具有以下选项： <code>-s discovery-domain-set-name</code> 指定一个搜索域集。
add-node	<code>add-node</code> 子命令用于将节点添加到指定的搜索域。 <code>add-node</code> 子命令的语法如下所示： <pre># isnsadm add-node option iscsi-node-name, ...</pre>

`add-node` 具有以下选项：

`-d, --dd discovery-domain-name`
指定一个搜索域。

`create-dd` `create-dd` 子命令用于使用您指定的名称创建搜索域。

`create-dd` 子命令的语法如下所示：

```
# isnsadm create-dd discovery-domain-name, ...
```

`create-dd` 没有选项。

`create-dd-set` `create-dd-set` 子命令用于使用您指定的名称创建搜索域集。

`create-dd-set` 子命令的语法如下所示：

```
# isnsadm create-dd-set discovery-domain-set-name, ...
```

`create-dd-set` 没有选项。

`delete-dd` `delete-dd` 子命令用于删除具有您指定的名称的搜索域。

`delete-dd` 子命令的语法如下所示：

```
# isnsadm delete-dd discovery-domain-name, ...
```

`delete-dd` 没有选项。

`delete-dd-set` `delete-dd-set` 子命令用于删除具有您指定的名称的搜索域集。

`delete-dd-set` 子命令的语法如下所示：

```
# isnsadm delete-dd-set discovery-domain-set-name, ...
```

`delete-dd-set` 没有选项。

`disable-dd-set` `disable-dd-set` 子命令用于禁用搜索域集。

`disable-dd-set` 子命令的语法如下所示：

```
# isnsadm disable-dd-set discovery-domain-set-name, ...
```

`disable-dd-set` 没有选项。

`enable-dd-set` `enable-dd-set` 子命令用于启用搜索域集。

`enable-dd-set` 子命令的语法如下所示：

```
# isnsadm enable-dd-set discovery-domain-set-name, ...
```

`enable-dd-set` 没有选项。

`list-dd` `list-dd` 子命令用于显示关于搜索域的信息。如果未指定操作数，则它会列出 iSNS 服务器上当前存在的所有搜索域。

`list-dd` 子命令的语法如下所示：

```
# isnsadm list-dd [option] [discovery-domain-name, ...]
```

`list-dd` 支持以下选项：

`-v, --verbose`

显示搜索域的成员内容。

`list-dd-set`

`list-dd-set` 子命令用于列出 iSNS 服务器上存在的搜索域集（包括启用的和禁用的）。请注意，没有 `dd-set` 注册。如果未指定操作数，它将列出所有搜索域集。

`list-dd-set` 子命令的语法如下所示：

```
# isnsadm list-dd-set [option] [discovery-domain-set-name, ...]
```

`list-dd-set` 支持以下选项：

`-v, --verbose`

显示搜索域集内的所有搜索域。

`list-node`

`list-node` 子命令用于显示当前已向 iSNS 服务器注册的节点的相关信息，以及未向其注册且属于非缺省搜索域的节点的相关信息。对于后一种情况，节点将其类型字段显示为 `unknown`（未知）。如果未指定操作数，`list-node` 将列出 iSNS 服务器已知的所有节点。

`list-node` 子命令的语法如下所示：

```
# isnsadm list-node [options] [iscsi-node-name, ...]
```

`list-node` 支持以下选项：

`-t, --target`

对列表进行过滤以便只显示 iSCSI 目标节点。

`-i, --initiator`

对列表进行过滤以便只显示 iSCSI 启动器节点。

`-v, --verbose`

显示节点的详细信息。不使用此选项时，只会显示名称、别名和类型信息。

`modify-dd`

`modify-dd` 子命令用于修改指定的搜索域的属性。

`modify-dd` 子命令的语法如下所示：

```
# isnsadm modify-dd option discovery-domain-name
```

`modify-dd` 具有以下选项：

`-n discovery-domain-name`

指定要向现有的搜索域应用的搜索域新名称。

modify-dd-set	modify-dd-set 子命令用于修改搜索域集。 modify-dd-set 子命令的语法如下所示： <pre># isnsadm modify-dd-set option discovery-domain-set-name</pre> modify-dd-set 具有以下选项： -n <i>discovery-domain-set-name</i> 指定要向现有的搜索域集应用的搜索域集新名称。
remove-dd	remove-dd 子命令用于删除与指定的搜索域集的关联。 remove-dd 子命令的语法如下所示： <pre># isnsadm remove-dd option discovery-domain-name, ...</pre> remove-dd 具有以下选项： -s <i>discovery-domain-set-name</i> 指定将从中删除搜索域的搜索域集。
remove-node	remove-node 子命令用于删除节点。 remove-node 子命令的语法如下所示： <pre># isnsadm remove-node option iscsi-node-name, ...</pre> remove-node 具有以下选项： -d <i>discovery-domain-name</i> 指定将从中删除节点的搜索域。
show-config	show-config 子命令用于显示 iSNS 服务器的管理设置。注意，可以通过服务管理工具（请参见 smf(5)）修改这些设置。请参见 isns(1M)。 show-config 子命令的语法如下所示： <pre># isnsadm show-config</pre> show-config 没有选项。

示例

示例1 显示客户机

下面的示例使用 list-node 子命令显示客户机。

```
# isnsadm list-node -v
iSCSI Name: iqn.1986-03.com.sun:01:000e0c9f10da.45173FEA.engr
    Alias: STK5320_NAS
    Type: Target
    Network Entity: SE5310
    Portal: 172.20.57.95:3260
        Portal Group: 1
```

示例1 显示客户机 (续)

```

        Portal: 172.20.56.95:3260
        Portal Group: 1
        DD Name: Default
iSCSI Name: iqn.1986-03.com.sun:01:000e0c9f10da.454F00A2.acct
        Alias:
        Type: Target
        Network Entity: SE5310
        Portal: 172.20.57.95:3260
        Portal Group: 1
        Portal: 172.20.56.95:3260
        Portal Group: 1
        DD Name: Default
iSCSI Name: iqn.1986-03.com.sun:01:e00000000000.46fd8e2b
        Alias: host-x2100
        Type: Initiator
        Network Entity: iqn.1986-03.com.sun:01:e00000000000.46fd8e2b
        Portal: 172.20.236.123:58530
        Portal Group: 1
        DD Name: Default

```

示例2 显示搜索域

下面的示例使用 `list-dd` 子命令显示搜索域。

```

# isnsadm list-dd -v
DD name: Default
        DD set(s): Default
        iSCSI Name: iqn.1986-03.com.sun:01:000e0c9f10da.45173FEA.engr
        iSCSI Name: iqn.1986-03.com.sun:01:000e0c9f10da.454F00A2.acct
        iSCSI name: iqn.1986-03.com.sun:01:e00000000000.46fd8e2b
DD name: acct-dd
DD name: engineering-dd

```

示例3 添加节点

下面的示例使用 `add-node` 子命令将节点添加到搜索域中，从而创建搜索域成员身份。

```

# isnsadm add-node -d engineering-dd \
iqn.1986-03.com.sun:01:000e0c9f10da.454F00A2.engr

```

示例4 删除节点

下面的示例使用 `remove-node` 子命令将节点从搜索域中删除，从而删除搜索域成员身份。

```

# isnsadm remove-node -d acct-dd \
iqn.1986-03.com.sun:01:000e0c9f10da.454F00A2.acct

```


示例5 创建搜索域集

下面的示例使用 `create-dd-set` 子命令创建搜索域集。

```
# isnsadm create-dd-set operation-dd-set
```

示例6 显示搜索域集

下面的示例使用 `list-dd-set` 子命令显示搜索域集。

```
# isnsadm list-dd-set -v
DD Set name: Default
    State: Disabled
    DD Name: Default
DD Set name: operation-dd-set
    State: Disabled
```

示例7 添加搜索域

下面的示例使用 `add-dd` 子命令将搜索域添加到搜索域集中。

```
# isnsadm add-dd -s operation-dd-set engineering-dd
```

示例8 显示搜索域集

下面的示例使用 `list-dd-set` 子命令显示搜索域集的属性。

```
# isnsadm list-dd-set
DD Set name: Default
    State: Disabled
    DD Name: Default
DD Set name: operation-dd-set
    State: Disabled
    DD Name: engineering-dd
```

示例9 启用搜索域集

下面的示例使用 `enable-dd-set` 子命令启用搜索域集。

```
# isnsadm enable-dd-set Default
```

示例10 禁用搜索域集

下面的示例使用 `disable-dd-set` 子命令禁用搜索域集。

```
# isnsadm disable-dd-set Default
```

示例11 显示管理设置

下面的示例使用 `show-config` 子命令显示当前的管理设置。

```
# isnsadm show-config
Data Store Location: /etc/isns/isnsdata.xml
Entity Status Inquiry Non-Response Threshold: 3
```

示例 11 显示管理设置 (续)

```
Management SCN Enabled: yes
Authorized Control Node Names: -
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/storage/isns
接口稳定性	Volatile (可变)

另请参见 [iscsiadm\(1M\)](#)、[isns\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

附注 如果调用某个子命令时使用了多个操作数，并且在一个或多个（并非全部）操作数上失败，则 `isnsadm` 会显示一条指示部分失败的一般消息，并显示失败操作数的列表。要查看有关某个特定操作数的错误，可以对该失败操作数发出同一子命令。

引用名

itadm – 管理 iSCSI 目标

用法概要

```

itadm create-target [-a,--auth-method radius | chap | none | default]
                  [-s,--chap-secret] [-S,--chap-secret-file path]
                  [-u,--chap-user chap-user-name] [-n,--node-name target_node_name]
                  [-l,--alias alias] [-t,--tpg tpg-name[,tpg-name]]

itadm modify-target [-a,--auth-method radius | chap | none | default]
                  [-s,--chap-secret] [-S,--chap-secret-file path]
                  [-u,--chap-user chap-user-name] [-n,--node-name new_target_node_name]
                  [-l,--alias alias] [-t,--tpg tpg-name[,tpg-name]] target_node_name

itadm delete-target [-f,--force] target_node_name

itadm list-target [-v,--verbose] [target_node_name]

itadm create-tpg tpg_name IP-address[:port] [IP-address[:port]]...

itadm list-tpg [-v,--verbose] [tpg_name]

itadm delete-tpg [-f,--force] tpg_name

itadm create-initiator [-s,--chap-secret] [-S,--chap-secret-file path]
                    [-u,--chap-user chap-user-name] initiator_node_name

itadm modify-initiator [-s,--chap-secret] [-S,--chap-secret-file path]
                    [-u,--chap-user chap-user-name] initiator_node_name

itadm list-initiator [-v,--verbose] initiator_node_name

itadm delete-initiator initiator_node_name

itadm modify-defaults [-a,--auth-method radius | chap | none]
                    [-r,--radius-server IP-address[:port]] [-d,--radius-secret]
                    [-D,--radius-secret-file path][-i,--isns enable | disable]
                    [-I,--isns-server IP-address[:port][,IP-address[:port]]]

itadm list-defaults

```

描述

itadm 命令用于管理 [stmfadm\(1M\)](#) 和 [libstmf\(3LIB\)](#) 中介绍的 SCSI 目标模式框架中的 Internet SCSI (iSCSI) 目标节点。这允许 iSCSI 启动器使用 iSCSI 协议访问 STMF 逻辑单元。除了 iSCSI 目标节点外，itadm 还管理其他两类托管对象：iSCSI 目标门户组与 iSCSI 启动器节点上下文。

itadm 是作为一组子命令以及用于各个子命令的选项和操作数实现的。下文中的相应部分分别描述了这些子命令。除了其子命令，itadm 还有一个帮助命令，用以显示该实用程序的用法信息。该帮助命令是使用 `-?` 选项调用的。

iSCSI 目标门户组

iSCSI 目标网络门户是指可由启动器节点用来连接到 iSCSI 目标的 IP 地址和 TCP 端口。这些门户的集合称为目标门户组 (Target Portal Group, TPG)。您可以使用 TPG 来限制对 iSCSI 目标的访问。可以使用 `itadm modify -t` 命令将特定 iSCSI 目标绑定到 TPG。将在属于 TPG 的每个 IP 地址上创建一个 iSCSI 侦听器，它侦听与 iSCSI 目标的连接。

TPG 是由在创建 TPG 时提供的唯一名称所标识的。在创建 TPG 时会自动生成一个范围介于 2-65535 的数字“目标门户组标记”。“目标门户组标记 1”是为没有在目标上设置显式“目标门户组”时使用的“缺省”目标门户组保留的。缺省 TPG 的门户与来自端口 3260 上所有网络接口的请求相匹配。

iSCSI 启动器节点上下文

某些操作，例如通过质询握手身份验证协议 (Challenge Handshake Authentication Protocol, CHAP) 进行的验证，需要与远程 iSCSI 启动器节点关联的参数。这些参数与 iSCSI 启动器节点上下文关联。iSCSI 启动器节点上下文由其启动器节点名称标识，该名称的格式为 IQN 或 EUI 格式（请参见 RFC 3720）。例如：

```
iqn.1986-03.com.sun:01:e00000000000.47d55444
eui.02004567A425678D
```

指定 IP 地址

许多 `itadm` 子命令要求您指定一个或多个具有可选端口号的 IP 地址。对于 IPv4，请使用标准的点分十进制记法。对于 IPv6，请用方括号将地址括起来。以下是示例指定方法。

```
IPv4:  10.2.4.1
        10.2.4.1:3260
IPv6:  [1080:0:0:0:8:800:200C:417A]
        [1080:0:0:0:8:800:200C:417A]:3260
```

子命令

下面是 `itadm` 子命令及其选项。

itadm create-target

```
itadm create-target [-a,--auth-method radius | chap | none | default]
                    [-s,--chap-secret]
                    [-S,--chap-secret-file path] [-u,--chap-user chap-user-name]
                    [-n,--node-name target_node_name] [-l,--alias alias]
                    [-t,--tpg tpg-name[,tpg-name,...]]
```

使用指定选项创建 iSCSI 目标。选项如下所述。

-a,--auth-method radius | chap | none | default

指定用于目标的验证方式。有效值是 `radius`、`chap` 和 `none`。`chap` 表示必须使用质询握手身份验证协议 (Challenge Handshake Authentication Protocol, CHAP) 对连接到此目标的启动器进行验证。`radius` 表示还应当使用 CHAP 对启动器进行验证，但应当从中央 RADIUS 服务器获取所需的验证参数（请参见 `radius-server` 和 `radius-secret` 选项）。`none` 表示不需要对与目标的连接进行验证。`default` 表示目标将使用此属性的全局设置。（请参见 `modify-defaults` 子命令）。

-s,--chap-secret

交互 CHAP 验证期间要发送的 CHAP 密钥。此属性没有缺省值。最大长度为 255 个字符；最小必需长度为 12 个字符。

-S,--chap-secret-file path

如 `-s` 选项中所述，包含 CHAP 密钥的临时文件的路径。

-u,--chap-user *chap-user-name*

为目标指定在交互 CHAP 验证中使用的 CHAP 用户名。此值只允许用于目标，不能全局设置，并且只会在启动器节点被配置为使用交互 CHAP 验证时使用。如果没有指定值，则将目标节点名称用作用户名。请参见 [iscsiadm\(1M\)](#)。

-n,--node-name *target_node_name*

iSCSI 目标节点由其目标节点名称标识，该名称的格式为 IQN 或 EUI 格式（请参见 RFC 3720）。此选项用于创建该名称。

-l,--alias *alias*

与目标节点关联的替代标识符。此标识符不需要是唯一的。

-t,--tpg *tpg-name*[,*tpg-name*,...]

一个目标门户组 (Target Portal Group, TPG) 列表（用以指定启动器可以用来访问特定目标的 TGP）或者关键字 `default`。如果指定了 `default`，目标将使用缺省门户 INADDR_ANY:3260。

```
itadm modify-target  itadm modify-target [-a,--auth-method radius | chap | none | default]
                    [-s,--chap-secret] [-S,--chap-secret-file path]
                    [-u,--chap-user chap-user-name] [-n,--node-name new_tgt_node_name]
                    [-l,--alias alias] [-t,--tpg tpg-name[,tpg-name]] target_node_name
```

根据指定的选项修改 iSCSI 目标。选项如下所述。

-a,--auth-method *radius* | *chap* | *none* | *default*

如上文中的 `create-target` 子命令部分所述。

-s,--chap-secret

如上文中的 `create-target` 子命令部分所述。

-S,--chap-secret-file *path*

如上文中的 `create-target` 子命令部分所述。

-u,--chap-user *chap-user-name*

如上文中的 `create-target` 子命令部分所述。要删除显式设置的 CHAP 用户名，请使用 `-u none`。

-n,--node-name *target_node_name*

重命名目标。另请参见上文中的 `create-target` 子命令部分中对 `-n` 的说明。

-l,--alias *alias*

如上文中的 `create-target` 子命令部分所述。要删除显式设置的别名，请使用 `-l none`。

-t,--tpg *tpg-name*[,*tpg-name*,...]

如上文中的 `create-target` 子命令部分所述。

```
itadm list-target  itadm list-target [-v,--verbose] [target_node_name]
```

列出所配置的目标的相关信息。如果指定了 *target_node_name*，则只会列出关于该目标的信息。选项如下所述。

-v,--verbose
详细模式。

itadm delete-target **itadm delete-target** [-f,--force] *target_node_name*

删除由 *target_node_name* 指定的目标。目标必须处于 **offline**（脱机）状态才能删除。选项如下所述。

-f,--force
如果目标仍然保持联机状态，此选项将尝试在删除目标前使其脱机。

itadm create-tpg **itadm create-tpg** *tpg_name IP-address[:port]...*

创建由指定的门户组成的 iSCSI 目标门户组，并为其分配标识符 *tpg_name*。每个门户都是一个 IP 地址与端口对。IPv4 门户是以点分地址表示法指定的，例如 172.31.255.255。IPv6 门户地址必须用方括号括起来。

此子命令没有选项。

itadm list-tpg **itadm list-tpg** [-v,--verbose] [*tpg_name*]

列出所配置的目标门户组的相关信息。如果指定了 *tpg_name*，则只会列出与 *tpg_name* 关联的目标门户组的相关信息。选项如下所述。

-v,--verbose
详细模式。

itadm delete-tpg **itadm delete-tpg** [-f,--force] *tpg_name*

删除与 *tpg_name* 关联的目标门户组。选项如下所述。

-f,--force
如果 TPG 与任何目标关联，则除非指定了此选项，否则删除的请求将被拒绝。

itadm create-initiator **itadm create-initiator** [-s,--chap-secret] [-S,--chap-secret-file *path*]
 [-u,--chap-user *chap-user-name*] *initiator_node_name*

配置与名为 *initiator_node_name* 的远程启动器关联的参数。选项如下所述。

-s,--chap-secret
如上文中的 **create-target** 子命令部分所述。

-S,--chap-secret-file *path*
如上文中的 **create-target** 子命令部分所述。

-u,--chap-user *chap-user-name*
为启动器指定用于 CHAP 验证的 CHAP 用户名。如果没有指定值，则将启动器节点名称用作用户名。

itadm modify-initiator **itadm modify-initiator** [-s,--chap-secret] [-S,--chap-secret-file *path*]
 [-u,--chap-user *chap-user-name*] *initiator_node_name*

修改与名为 *initiator_node_name* 的远程启动器关联的参数。选项如下所述。

-s,--chap-secret

如上文中的 `create-target` 子命令部分所述。

-S,--chap-secret-file *path*

如上文中的 `create-target` 子命令部分所述。

-u,--chap-user *chap-user-name*

为启动器指定用于 CHAP 验证的 CHAP 用户名。如果没有指定值，则将启动器节点名称用作用户名。

itadm delete-initiator `itadm delete-initiator initiator_node_name`

删除与名为 *initiator_node_name* 的远程启动器关联的参数。此子命令没有选项。

itadm list-initiator `itadm list-initiator [-v,--verbose] initiator_node_name`

列出与名为 *initiator_node_name* 的启动器关联的参数。选项如下所述。

-v,--verbose

详细模式。

itadm modify-defaults `itadm modify-defaults [-a,--auth-method radius | chap | none]
 [-r,--radius-server IP-address[:port]] [-d,--radius-secret]
 [-D,--radius-secret-file path] [-i,--isns enable | disable]
 [-I,--isns-server IP-address[:port] [,IP-address[:port]]]`

修改缺省参数。选项如下所述。

-a,--auth-method *radius | chap | none*

指定用于所有目标的缺省验证方式。有效值是 *radius*、*chap* 和 *none*。*chap* 表示必须使用质询握手身份验证协议 (Challenge Handshake Authentication Protocol, CHAP) 对连接到此目标的启动器进行验证。*radius* 表示还应当使用 CHAP 对启动器进行验证，但应当从中央 RADIUS 服务器获取所需的验证参数。（请参见 `--radius-server` 和 `--radius-secret` 选项）。*none* 表示不需要对与目标的连接进行验证。各个目标都可使用 `create-target` 和 `modify-target` 子命令的 `-a` 选项覆盖此全局设置。

-d,--radius-secret

用于集中 CHAP 验证的 RADIUS 共享密钥。

-D,--radius-secret-file *path*

如 `-d` 选项中所述，包含 CHAP 密钥的临时文件的路径。

-i,--isns enable | disable

指定是否向所定义的 iSCSI 名称服务 (iSCSI Name Service, iSNS) 服务器集注册目标。

-I,--isns-server *IP-address[:port]* [,*IP-address[:port]*,...]

定义当设置了与相应目标关联的 *isns* 选项时，要向其注册 iSCSI 目标节点的 iSNS 服务器的列表。最多可以指定八个 iSNS 服务器。要删除所有 iSNS 服务器，请使用 `-I none`。

```
-r,--radius-server IP-address[:port]
```

指定用于集中 CHAP 验证的 RADIUS 服务器的 IP 地址。

itadm list-defaults

```
itadm list-defaults
```

列出关于缺省属性的信息。此子命令没有选项。

示例

示例1 创建目标

以下命令用于创建目标。

```
# itadm create-target
Target iqn.1986-03.com.sun:02:72e1b181-7bce-c0e6-851e-ec0d8cf14b7a
successfully created
```

示例2 使用特定名称创建目标

以下命令用于使用特定 IQN 创建目标。

```
# itadm create-target -n eui.20387ab8943ef7548
or:
# itadm create-target \
-n iqn.1986-03.com.sun:02:a9a366f8-cc2b-f291-840948c7f29e
```

示例3 更改名称

以下命令用于更改现有目标的 IQN。

```
# itadm modify-target -n eui.20387ab8943ef7548 \
iqn.1986-03.com.sun:02:a9a366f8-909b-cc2b-f291-840948c7f29e
```

示例4 设置 CHAP 验证

以下命令用于使用缺省 CHAP 用户名为目标设置 CHAP 验证。

```
# itadm modify-initiator -s iqn.1986-03.com.sun:01:e00000000000.47d55444
Enter CHAP secret: *****
Re-enter secret: *****
```

```
# itadm modify-target -a chap eui.20387ab8943ef7548
```

示例5 创建目标门户组

以下命令用于创建两个目标门户组 A 和 B，并且为 TPG 2 中的地址使用端口 8000。

```
# itadm create-tpg A 192.168.0.1 192.168.0.2
# itadm create-tpg B 192.168.0.2:8000 192.168.0.2:8000
```

示例6 将目标配置为使用 TPG

以下命令用于将目标配置为使用 TPG A 和 B。

```
# itadm modify-target -t A,B eui.20387ab8943ef7548
```


示例7 为特定目标设置 RADIUS 验证

以下命令用于为特定目标设置 RADIUS 验证。

```
# itadm modify-defaults -r 192.168.10.1 -d
Enter RADIUS secret: *****
Re-enter secret: *****

# itadm modify-target -a radius eui.20387ab8943ef7548
```

示例8 为所有目标设置 RADIUS 验证

以下命令用于为所有目标设置 RADIUS 验证。

```
# itadm modify-defaults -d -r 192.168.10.1 -a radius
Enter RADIUS secret: *****
Re-enter secret: *****
```

上面的命令假定所有目标都是使用 `-a default` 创建的。

示例9 列出缺省属性

以下命令用于列出缺省属性。

```
# itadm list-defaults
iSCSI Target Default Properties:

alias:          none
auth:           none
radiusserver:   none
radiussecret:   unset
isns:           disabled
isnsserver:     2.3.4.5,4.5.6.7
```

示例10 列出目标

以下命令用于列出目标。

```
# itadm list-target
TARGET NAME                                     STATE    SESSIONS
iqn.1986-03.com.sun:02:72e1b181-7bce-c0e6-851e-ec0d8cf14b7a online    0
iqn.1986-03.com.sun:02:2cb0c526-c05a-e279-e396-a367006f4227 online    0
iqn.1986-03.com.sun:02:d14125bb-1c9d-c28d-97b0-f89259b642f3 online    0
iqn.1986-03.com.sun:02:03ff9fc5-794a-e9b4-a081-bb82917c292a online    0
```

示例11 列出目标（详细）

以下命令用于使用详细选项列出目标。

```
# itadm list-target -v
TARGET NAME                                     STATE    SESSIONS
iqn.1986-03.com.sun:02:d23e68d8-2d79-c988-98e7-a6361689d33c online    0
alias:                                          -
```

示例 11 列出目标（详细）（续）

```
auth: none (defaults)
targetchapuser: -
targetchapsecret: unset
tpg-tags: default
iqn.1986-03.com.sun:02:94ec46d4-c8e1-6993-ef03-ffc1dcd66606 online 1
alias: -
auth: chap
targetchapuser: -
targetchapsecret: unset
tpg-tags: nge1_ipv4 = 3
```

示例 12 列出特定目标

以下命令用于使用详细选项列出目标。

```
# itadm list-target -v \
iqn.1986-03.com.sun:02:2cb0c526-c05a-e279-e396-a367006f4227
TARGET NAME STATE SESSIONS
iqn.1986-03.com.sun:02:2cb0c526-c05a-e279-e396-a367006f4227 online 1
alias: -
auth: chap
targetchapuser: -
targetchapsecret: unset
tpg-tags: nge1_ipv4 = 3
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/storage/iscsi/iscsi-target
接口稳定性	Committed（已确定）

另请参见 [iscsiadm\(1M\)](#)、[stmfadm\(1M\)](#)、[libstmf\(3LIB\)](#)、[attributes\(5\)](#)

引用名	itu – convert packages to Driver Update format and patch Solaris install media for Install Time Update
用法概要	<pre>itu makedu -r solaris_release [-v] [-f] [-d output_dir] [-o iso_file] [-l iso_label] package [package...] itu updatemedia -R media_root [-v] [-f] [-o iso_file] [-l iso_label] pkg_or_patch [pkg_or_patch...] itu makeiso -o iso_file [-v] [-l iso_label] media_root</pre>
描述	The <code>itu</code> utility converts driver packages to Driver Update (DU) format and patches a Solaris install media with driver packages and patches for Install Time Update (ITU). <code>itu</code> has three subcommands: <code>makedu</code> , <code>updatemedia</code> and <code>makeiso</code> .
选项	The following options are supported: -d <i>output_dir</i> Directory where the Driver Update directory is to be created. -f If <i>output_dir</i>/DU or <i>iso_file</i> already exists, remove it without asking first. -l <i>iso_label</i> Label/volume name of the ISO image (if -o option is specified). -o <i>iso_file</i> Path of the ISO image file to be created. For subcommands <code>updatemedia</code> and <code>makeiso</code>, it will be a bootable ISO image. This option must be specified for subcommand <code>makeiso</code>. -R <i>media_root</i> Top-level directory of on-disk image of Solaris installation media. This option must be specified for subcommand <code>updatemedia</code>. -r <i>solaris_release</i> Solaris release number for which the Driver Update is intended. It takes the form of the output of <code>uname -r</code>, for example, 5.10. This option must be specified for subcommand <code>makedu</code>. -v Verbose. Multiple -v options increase verbosity.
子命令	The <code>itu</code> subcommands are described as follows.
makedu	The <code>makedu</code> subcommand takes one or more driver packages as input and converts them to DU format. At the beginning of an interactive Solaris installation session, these driver updates can be applied to the running kernel, which will then also automatically apply them to the newly installed Solaris at the end of the installation process. The -r option is required to specify the Solaris release number for which the driver updates apply. The <i>solaris_release</i> option argument takes the form <code>uname -r</code> output, for example, 5.10 or 5.11.

If the `-d` option is specified, the resulting DU directory tree is placed in the directory *output_dir*.

If the `-o` option is specified, a (non-bootable) ISO image of the DU directory tree is written in the file *iso_file*. This ISO image can be burned onto a CD/DVD using `cdrw(1)` or `cdrecord(1)` (not a SunOS man page).

At least one of `-d` and `-o` option must be specified. If both are specified, then both an ISO image and a directory tree are generated.

updatemedia

The `updatemedia` subcommand takes a list of driver packages and patches as input and applies them to the miniroot of a Solaris install media. It also places them in a subdirectory called ITUs under the Solaris install media's top-level directory:

media_root/ITUs

When booting a system from the updated media, the patches and packages will be part of the booted Solaris image. They will also be applied to the target system being installed at the end of the installation process.

The `-R` option must be entered on the command line to specify the Solaris install media. Note that the install media must be on a location that is writable by `itu`.

If the `-o` option is specified, a bootable ISO image of the patched install media is also created in the file *iso_file*. The ISO image can then be burned onto a CD or DVD.

makeiso

The `makeiso` subcommand runs `mkisofs(8)` to create a bootable Solaris ISO image of the Solaris install media *media_root* and writes it to the file *iso_file*. The ISO image file can then be burned onto a CD or DVD with utilities such as `cdrw(1)` or `cdrecord(1)`. (Note that `mkisofs(8)` and `cdrecord(1)` are not SunOS man pages.)

注意 – The Solaris install media *media-root* must contain the file `boot/grub/stage2_eltorito`, which will be written to the media boot sectors. This file will be modified with some boot information, thus it has to be writable. If necessary, first save a copy, prior to running this subcommand.

操作数

The following operands are supported:

package [*package*...]

One or more driver packages.

pkg_or_patch [*pkg_or_patch*...]

One or more patches or packages.

media_root

The top-level directory of a Solaris install media.

示例

示例 1 Creating a DU CD/DVD

The following commands create a Driver Update CD/DVD containing the packages SAMPLEpkg1 and SAMPLEpkg2.

```
# itu makedu -r 5.10 -o my.iso SAMPLEpkg1 SAMPLEpkg2
# cdrw -i my.iso
```

示例 2 Patching the Solaris Install Media

The following command patches the Solaris install media in /export/s10u1 with patch /opt/patches/123456-07 and driver package /opt/pkgs/MYdriver. The command also creates a bootable ISO image with ISO label “MyS10U1” in the file /tmp/dvd.iso.

```
# /usr/bin/itu updatemedia -R /export/s10u1 -o /tmp/dvd.iso -l MyS10U1 \
/opt/patches/123456-07 /opt/pkgs/MYdriver
```

示例 3 Creating a Bootable ISO Image

The following commands create the bootable ISO image mydvd.iso of the Solaris install image /export/solaris-10u1 with ISO label “Special-S10”.

```
# /usr/bin/itu makeiso -o mydvd.iso -l "Special-S10" \
/export/solaris-10u1
# cdrw -i mydvd.iso
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[cdrw\(1\)](#), [pkgadd\(1M\)](#), [attributes\(5\)](#)

[mkisofs\(8\)](#), ([/usr/share/man/man8/mkisofs.8](#)), in the SUNWfsman package (not a SunOS man page)

引用名

k5srvutil – host key table (keytab) manipulation utility

用法概要

`/usr/sbin/k5srvutil operation [-ik] [-f filename]`

描述

The `k5srvutil` command allows a system manager to list or change keys currently in his keytab or to add new keys to the keytab.

The operand *operation* must be one of the following:

- `list` Lists the keys in a keytab, showing version number and principal name.
- `change` Changes all the keys in the keytab to new randomly-generated keys, updating the keys in the Kerberos server's database to match those by using the `kadmin` protocol. If a key's version number does not match the version number stored in the Kerberos server's database, the operation fails. The old keys are retained so that existing tickets continue to work. If the `-i` flag is specified, `k5srvutil` prompts for yes or no before changing each key. If the `-k` option is used, the old and new keys are displayed.
- `delold` Deletes keys that are not the most recent version from the keytab. This operation should be used at some point after a change operation to remove old keys. If the `-i` flag is specified, `k5srvutil` asks the user whether the old keys associated with each principal should be removed.
- `delete` Deletes particular keys in the keytab, interactively prompting for each key.

In all cases, the default keytab file is `/etc/krb5.keytab` file unless this is overridden by the `-f` option.

`k5srvutil` uses the [kadmin\(1M\)](#) program to edit the keytab in place. However, old keys are retained, so they are available in case of failure.

选项

The following options are supported:

- `-f filename` Specify a keytab file other than the default file, `/etc/krb5.keytab`.
- `-i` Prompts user before changing keys when using the `change` or `delold` operands.
- `-k` Displays old and new keys when using the `change` operand.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/security/kerberos-5
Interface Stability	Committed

另请参见

[ktutil\(1\)](#), [kadmin\(1M\)](#), [attributes\(5\)](#)

引用名	kadb – a kernel debugger	
用法概要		
SPARC	ok boot <i>device_specifier</i> kadb [-d] [<i>boot-flags</i>]	
x86	select (b)oot or (i)nterpreter: b kadb [-d] [<i>boot-flags</i>]	
描述	kadb, an interactive kernel debugger, has been replaced by kmdb(1). For backwards compatibility, the methods used to load kadb will be interpreted as requests to load kmdb(1). Unlike with the compatibility link from adb(1) to mdb(1), kmdb(1) will always load in its native user interface mode, regardless of the name used to load it. kmdb(1) is based on mdb(1), and thus shares mdb's user interface style and feature set. The mdb(1) man page describes the features and operation of mdb. The kmdb(1) man page describes the differences between mdb and kmdb. This man page describes the major changes and incompatibilities between kadb and kmdb. Consult the Solaris Modular Debugger Guide for a detailed description of both mdb and kmdb.	
Major changes	This section briefly lists the major differences between kadb and kmdb. It is not intended to be exhaustive.	
	Debugger Loading and Unloading	kmdb(1) may be loaded at boot, as with kadb. It may also be loaded after boot, thus allowing for kernel debugging and execution control without requiring a system reboot. If kmdb(1) is loaded after boot, it may be unloaded.
	mdb Feature Set	The features introduced by mdb(1) , including access to kernel type data, debugger commands (dcmds), debugger modules (dmods), and enhanced execution control facilities, are available under kmdb(1) . Support for changing the representative CPU (:x) is available for both SPARC and x86. Furthermore, full execution-control facilities are available after the representative CPU has been changed.
Significant Incompatibilities	This section lists the significant features that have changed incompatibly between kadb and kmdb(1) . It is not intended to be exhaustive. All kmdb(1) commands referenced here are fully described in the kmdb(1) man page. A description as well as examples can be found in the Solaris Modular Debugger Guide.	
	Deferred Breakpoints	The kadb-style “module#symbol:b” syntax is not supported under kmdb(1) . Instead, use “::bp module‘symbol”.
	Watchpoints	The ::wp dcmd is the preferred way to set watchpoint with kmdb. Various options are available to control the type of watchpoint set, including -p for physical watchpoints

(SPARC only), and -i for I/O port watchpoints (x86 only). \$l is not supported, therefore, the watchpoint size must be specified for each watchpoint created.

Access to I/O Ports (x86 only) The commands used to access I/O ports under kadb have been replaced with the ::in and ::out dcmts. These two dcmts allow both read and write of all I/O port sizes supported by kadb.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/library/processor

另请参见 [adb\(1\)](#), [mdb\(1\)](#), [kmdb\(1\)](#), [attributes\(5\)](#)

Solaris Modular Debugger Guide

引用名

kadmin, kadmin.local – Kerberos database administration program

用法概要

```
/usr/sbin/kadmin [-r realm] [-p principal] [-q query]
  [-s admin_server [:port]] [ [-c credential_cache]
  | [-k [-t keytab] | [-n]] | [-w password]] [-x db_args]...

/usr/sbin/kadmin.local [-r realm] [-p principal]
  [-q query] [-d dbname] [-e "enc:salt..."] [-m] [-D]
```

描述

kadmin and kadmin.local are interactive command-line interfaces to the Kerberos V5 administration system. They provide for the maintenance of Kerberos principals, policies, and service key tables (keytabs). kadmin and kadmin.local provide identical functionality; the difference is that kadmin.local can run only on the master KDC and does not use Kerberos authentication.

Except as explicitly noted otherwise, this man page uses kadmin to refer to both versions.

By default, both versions of kadmin attempt to determine your user name and perform operations on behalf of your “username/admin” instance. Operations performed are subject to privileges granted or denied to this user instance by the Kerberos ACL file (see [kadm5.acl\(4\)](#)). You may perform administration as another user instance by using the -p option.

The remote version, kadmin, uses Kerberos authentication and an encrypted RPC to operate securely from anywhere on the network. It normally prompts for a password and authenticates the user to the Kerberos administration server, kadmind, whose service principal is kadmin/fqdn. Some options specific to the remote version permit the password prompt to be bypassed. The -c option searches the named credentials cache for a valid ticket for the kadmin/fqdn service and uses it to authenticate the user to the Kerberos admin server without a password. The -k option searches a keytab for a credential to authenticate to the kadmin/fqdn service, and again no password is collected. If kadmin has collected a password, it requests a kadmin/fqdn Kerberos service ticket from the KDC, and uses that service ticket to interact with kadmind.

The local version, kadmin.local, must be run with an effective UID of root, and normally uses a key from the /var/krb5/.k5.realm stash file (see [kdb5_util\(1M\)](#)) to decrypt information from the database rather than prompting for a password. The -m option will bypass the .k5.realm stash file and prompt for the master password.

选项

The following options are supported:

-c *credentials_cache*

Search *credentials_cache* for a service ticket for the kadmin/fqdn service; it can be acquired with the [kinit\(1\)](#) program. If this option is not specified, kadmin requests a new service ticket from the KDC, and stores it in its own temporary credentials cache.

-d *dbname*

Specify a non-standard database name. [Local only]

-D

Turn on debug mode. [Local only]

- e *"enc:salt ..."*
Specify a different encryption type and/or key salt. [Local only]
- k [-t *keytab*]
Use the default keytab (-k) or a specific keytab (-t *keytab*) to decrypt the KDC response instead of prompting for a password. In this case, the default principal will be host/hostname. This is primarily used for keytab maintenance.
- m
Accept the database master password from the keyboard rather than using the /var/krb5/.k5.realm stash file. [Local only]
- n
Requests anonymous processing. Two types of anonymous principals are supported. For fully anonymous Kerberos, configure pkinit on the KDC and configure pkinit_anchors in the client's krb5.conf. Then use the -n option with a principal of the form @REALM (an empty principal name followed by the at-sign and a realm name). If permitted by the KDC, an anonymous ticket will be returned. A second form of anonymous tickets is supported; these realm-exposed tickets hide the identity of the client but not the client's realm. For this mode, use kinit -n with a normal principal name. If supported by the KDC, the principal (but not realm) will be replaced by the anonymous principal. As of release 1.8, the MIT Kerberos KDC supports only fully anonymous operation.
- p *principal*
Authenticate *principal* to the kadmin/fqdn service. Otherwise, kadmin will append /admin to the primary principal name of the default credentials cache, the value of the USER environment variable, or the username as obtained with getpwuid, in that order of preference.
- q *query*
Pass *query* directly to kadmin, which will perform *query* and then exit. This can be useful for writing scripts.
- r *realm*
Use *realm* as the default database realm.
- s *admin_server[:port]*
Administer the specified *admin* server at the specified port number (*port*). This can be useful in administering a realm not known to your client.
- w *password*
Use *password* instead of prompting for one. Note that placing the password for a Kerberos principal with administration access into a shell script can be dangerous if unauthorized users gain read access to the script or can read arguments of this command through ps(1).
- x *db_args*
Pass database-specific arguments to kadmin. Supported arguments are for LDAP and the Berkeley-db2 plug-in. These arguments are:

binddn=binddn

LDAP simple bind DN for authorization on the directory server. Overrides the `ldap_kadmin_dn` parameter setting in `krb5.conf(4)`.

bindpwd=bindpwd

Bind password.

dbname=name

For the Berkeley-db2 plug-in, specifies a name for the Kerberos database.

nconns=num

Maximum number of server connections.

port=num

Directory server connection port.

Commands

`list_requests`

Lists all the commands available for `kadmin`. Aliased by `lr` and `?`.

`get_privs`

Lists the current Kerberos administration privileges (ACLs) for the principal that is currently running `kadmin`. The privileges are based on the `/etc/krb5/kadm5.acl` file on the master KDC. Aliased by `getprivs`.

`add_principal [options] newprinc`

Creates a new principal, *newprinc*, prompting twice for a password. If the `-policy` option is not specified and a policy named `default` exists, then the `default` policy is assigned to the principal; note that the assignment of the `default` policy occurs automatically only when a principal is first created, so the `default` policy must already exist for the assignment to occur. The automatic assignment of the `default` policy can be suppressed with the `-clearpolicy` option. This command requires the `add` privilege. Aliased by `addprinc` and `ank`. The options are:

`-expire expdate`

Expiration date of the principal. See the `Time Formats` section for the valid absolute time formats that you can specify for *expdate*.

`-pwexpire pwexpdate`

Password expiration date. See the `Time Formats` section for the valid absolute time formats that you can specify for *pwexpdate*.

`-maxlife maxlife`

Maximum ticket life for the principal. See the `Time Formats` section for the valid time duration formats that you can specify for *maxlife*.

`-maxrenewlife maxrenewlife`

Maximum renewable life of tickets for the principal. See the `Time Formats` section for the valid time duration formats that you can specify for *maxrenewlife*.

`-kvno kvno`

Explicitly set the key version number.

-policy *policy*

Policy used by the principal. If both the **-policy** and **-clearpolicy** options are not specified, the default policy is used if it exists; otherwise, the principal will have no policy. Also note that the password and principal name must be different when you add a new principal with a specific policy or the default policy.

-clearpolicy

-clearpolicy prevents the default policy from being assigned when **-policy** is not specified. This option has no effect if the default policy does not exist.

{-|+}allow_postdated

-allow_postdated prohibits the principal from obtaining postdated tickets. (Sets the KRB5_KDB_DISALLOW_POSTDATED flag.) **+allow_postdated** clears this flag.

{-|+}allow_forwardable

-allow_forwardable prohibits the principal from obtaining forwardable tickets. (Sets the KRB5_KDB_DISALLOW_FORWARDABLE flag.) **+allow_forwardable** clears this flag.

{-|+}allow_renewable

-allow_renewable prohibits the principal from obtaining renewable tickets. (Sets the KRB5_KDB_DISALLOW_RENEWABLE flag.) **+allow_renewable** clears this flag.

{-|+}allow_proxiabile

-allow_proxiabile prohibits the principal from obtaining proxiabile tickets. (Sets the KRB5_KDB_DISALLOW_PROXIABLE flag.) **+allow_proxiabile** clears this flag.

{-|+}allow_dup_skey

-allow_dup_skey disables user-to-user authentication for the principal by prohibiting this principal from obtaining a session key for another user. (Sets the KRB5_KDB_DISALLOW_DUP_SKEY flag.) **+allow_dup_skey** clears this flag.

{-|+}requires_preauth

+requires_preauth requires the principal to preauthenticate before being allowed to kinit. (Sets the KRB5_KDB_REQUIRES_PRE_AUTH flag.) **-requires_preauth** clears this flag.

{-|+}requires_hwauth

+requires_hwauth requires the principal to preauthenticate using a hardware device before being allowed to kinit. (Sets the KRB5_KDB_REQUIRES_HW_AUTH flag.) **-requires_hwauth** clears this flag.

{-|+}allow_svr

-allow_svr prohibits the issuance of service tickets for the principal. (Sets the KRB5_KDB_DISALLOW_SVR flag.) **+allow_svr** clears this flag.

{-|+}allow_tgs_req

-allow_tgs_req specifies that a Ticket-Granting Service (TGS) request for a service ticket for the principal is not permitted. This option is useless for most things.

`+allow_tgs_req` clears this flag. The default is `+allow_tgs_req`. In effect, `-allow_tgs_req` sets the `KRB5_KDB_DISALLOW_TGT_BASED` flag on the principal in the database.

`{-|+}allow_tix`
`-allow_tix` forbids the issuance of any tickets for the principal. `+allow_tix` clears this flag. The default is `+allow_tix`. In effect, `-allow_tix` sets the `KRB5_KDB_DISALLOW_ALL_TIX` flag on the principal in the database.

`{-|+}needchange`
`+needchange` sets a flag in attributes field to force a password change; `-needchange` clears it. The default is `-needchange`. In effect, `+needchange` sets the `KRB5_KDB_REQUIRES_PWCHANGE` flag on the principal in the database.

`{-|+}password_changing_service`
`+password_changing_service` sets a flag in the attributes field marking this as a password change service principal (useless for most things).
`-password_changing_service` clears the flag. This flag intentionally has a long name. The default is `-password_changing_service`. In effect, `+password_changing_service` sets the `KRB5_KDB_PWCHANGE_SERVICE` flag on the principal in the database.

`{-|+}ok_as_delegate`
`+ok_as_delegate` sets the OK-AS-DELEGATE flag on tickets issued for use with this principal as the service, which clients may use as a hint that credentials can and should be delegated when authenticating to the service. (Sets the `KRB5_KDB_OK_AS_DELEGATE` flag.) `-ok_as_delegate` clears this flag.

`{-|+}ok_to_auth_as_delegate`
`+ok_to_auth_as_delegate` sets the service to allow the use of `S4U2Self`.
`-ok_to_auth_as_delegate` clears this flag.

`-randkey`
 Sets the key of the principal to a random value.

`-pw password`
 Sets the key of the principal to the specified string and does not prompt for a password. Note that using this option in a shell script can be dangerous if unauthorized users gain read access to the script.

`-e "enc:salt..."`
 Override the list of enctype:salttype pairs given in `kdc.conf(4)` for setting the key of the principal. The quotes are necessary if there are multiple enctype:salttype pairs. One key for each similar enctype and same salttype will be created and the first one listed will be used. For example, in a list of two similar encryptions with the same salt, "des-cbc-crc:normal des-cbc-md5:normal", one key will be created and it will be of type `des-cbc-crc:normal`.

-unlock

Unlocks the principal so that it can successfully authenticate. If the principal had previously been locked due to reaching `maxfailure` in `failurecountinterval` time then the principal will be locked for `lockoutduration` time.

Example:

```
kadmin: addprinc tlyu/admin
WARNING: no policy specified for "tlyu/admin@ACME.COM";
defaulting to no policy.
Enter password for principal tlyu/admin@ACME.COM:
Re-enter password for principal tlyu/admin@ACME.COM:
Principal "tlyu/admin@ACME.COM" created.
kadmin:
```

Errors:

KADM5_AUTH_ADD (requires add privilege)

KADM5_BAD_MASK (should not happen)

KADM5_DUP (principal exists already)

KADM5_UNK_POLICY (policy does not exist)

KADM5_PASS_Q_* (password quality violations)

delete_principal [-force] *principal*

Deletes the specified principal from the database. This command prompts for deletion, unless the `-force` option is given. This command requires the `delete` privilege. Aliased by `delprinc`.

Example:

```
kadmin: delprinc mwm_user
Are you sure you want to delete the principal
"mwm_user@ACME.COM"? (yes/no): yes
Principal "mwm_user@ACME.COM" deleted.
Make sure that you have removed this principal from
all kadmind ACLs before reusing.
kadmin:
```

Errors:

KADM5_AUTH_DELETE (requires delete privilege)

KADM5_UNK_PRINC (principal does not exist)

modify_principal [*options*] *principal*

Modifies the specified principal, changing the fields as specified. The options are as above for `add_principal`, except that password changing is forbidden by this command. In addition, the option `-clearpolicy` will clear the current policy of a principal. This command requires the `modify` privilege. Aliased by `modprinc`.

Errors:

KADM5_AUTH_MODIFY (requires modify privilege)

KADM5_UNK_PRINC (principal does not exist)

KADM5_UNK_POLICY (policy does not exist)

KADM5_BAD_MASK (should not happen)

change_password [*options*] *principal*

Changes the password of *principal*. Prompts for a new password if neither -randkey or -pw is specified. Requires the changepw privilege, or that the principal that is running the program to be the same as the one changed. Aliased by cpw. The following options are available:

-randkey

Sets the key of the principal to a random value.

-pw *password*

Sets the password to the specified string. Not recommended.

-e "enc:salt ..."

Override the list of enctype:salttype pairs given in [kdc.conf\(4\)](#) for setting the key of the principal. The quotes are necessary if there are multiple enctype:salttype pairs. For each key, the first matching similar enctype and same salttype in the list will be used to set the new key(s).

-keepold

Keeps the previous kvno's keys around. There is no easy way to delete the old keys, and this flag is usually not necessary except perhaps for TGS keys as it will allow existing valid TGTs to continue to work.

Example:

```
kadmin: cpw systest
Enter password for principal systest@ACME.COM:
Re-enter password for principal systest@ACME.COM:
Password for systest@ACME.COM changed.
kadmin:
```

Errors:

KADM5_AUTH_MODIFY (requires the modify privilege)

KADM5_UNK_PRINC (principal does not exist)

KADM5_PASS_Q_* (password policy violation errors)

KADM5_PASS_REUSE (password is in principal's password history)

KADM5_PASS_T00S00N (current password minimum life not expired)

`get_principal [-terse] principal`
Gets the attributes of *principal*. Requires the `inquire` privilege, or that the principal that is running the program to be the same as the one being listed. With the `-terse` option, outputs fields as quoted tab-separated strings. Aliased by `getprinc`.

Examples:

```
kadmin: getprinc tlyu/admin
Principal: tlyu/admin@ACME.COM
Expiration date: [never]
Last password change: Thu Jan 03 12:17:46 CET 2008
Password expiration date: [none]
Maximum ticket life: 24855 days 03:14:07
Maximum renewable life: 24855 days 03:14:07
Last modified: Thu Jan 03 12:17:46 CET 2008 (root/admin@ACME.COM)
Last successful authentication: [never]
Last failed authentication: [never]
Failed password attempts: 0
Number of keys: 5
Key: vno 2, AES-256 CTS mode with 96-bit SHA-1 HMAC, no salt
Key: vno 2, AES-128 CTS mode with 96-bit SHA-1 HMAC, no salt
Key: vno 2, Triple DES cbc mode with HMAC/sha1, no salt
Key: vno 2, ArcFour with HMAC/md5, no salt
Key: vno 2, DES cbc mode with RSA-MD5, no salt
Attributes: REQUIRES_PRE_AUTH
Policy: [none]
kadmin: getprinc -terse tlyu/admin
"tlyu/admin@ACME.COM" 0 1199359066 0 2147483647
"root/admin@ACME.COM" 1199359066 128 2 0 "[none]" 21474836
47 0 0 0 5 1 2 18 0 1 2
17 0 1 2 16 0 1 2 23 0 12
3 0
kadmin:
```

Errors:

- KADM5_AUTH_GET (requires the `get [inquire]` privilege)
- KADM5_UNK_PRINC (principal does not exist)

`list_principals [expression]`
Retrieves all or some principal names. *expression* is a shell-style glob expression that can contain the wild-card characters `?`, `*`, and `[]`'s. All principal names matching the expression are printed. If no expression is provided, all principal names are printed. If the expression does not contain an `@` character, an `@` character followed by the local realm is appended to the expression. Requires the `list` privilege. Aliased by `listprincs`, `get_principals`, and `getprincs`.

Examples:

```
kadmin: listprincs test*
test3@ACME.COM
test2@ACME.COM
test1@ACME.COM
testuser@ACME.COM
kadmin:
```

`add_policy [options] policy`

Adds the named policy to the policy database. Requires the add privilege. Aliased by `addpol`. The following options are available:

- maxfailure *maxnumber*
sets the maximum number of failures before the principal is locked after authentication failures in *failurecountinterval* time.
- failurecountinterval *failuretime*
sets the time after which the authentication failure count is reset 0. See the “Time Formats” section, below, for the valid time duration formats that you can specify for *failuretime*.
- lockoutduration *lockouttime*
sets the time in which the principal is locked from authenticating if *maxfailure* authentication failures occur within *failurecountinterval* time. See the “Time Formats” section, below, for the valid time duration formats that you can specify for *lockouttime*.
- maxlife *maxlife*
sets the maximum lifetime of a password. See the Time Formats section for the valid time duration formats that you can specify for *maxlife*.
- minlife *minlife*
sets the minimum lifetime of a password. See the Time Formats section for the valid time duration formats that you can specify for *minlife*.
- minlength *length*
sets the minimum length of a password.
- minclasses *number*
sets the minimum number of character classes allowed in a password. The valid values are:
 - 1
only letters (himom)
 - 2
both letters and numbers (hi2mom)
 - 3
letters, numbers, and punctuation (hi2mom!)

-history *number*
 sets the number of past keys kept for a principal.

Errors:

KADM5_AUTH_ADD (requires the add privilege)

KADM5_DUP (policy already exists)

delete_policy [-force] *policy*

Deletes the named policy. Unless the -force option is specified, prompts for confirmation before deletion. The command will fail if the policy is in use by any principals. Requires the delete privilege. Aliased by `delpol`.

Example:

```
kadmin: del_policy guests
Are you sure you want to delete the
policy "guests"? (yes/no): yes
Policy "guests" deleted.
kadmin:
```

Errors:

KADM5_AUTH_DELETE (requires the delete privilege)

KADM5_UNK_POLICY (policy does not exist)

KADM5_POLICY_REF (reference count on policy is not zero)

modify_policy [*options*] *policy*

Modifies the named policy. Options are as above for `add_policy`. Requires the modify privilege. Aliased by `modpol`.

Errors:

KADM5_AUTH_MODIFY (requires the modify privilege)

KADM5_UNK_POLICY (policy does not exist)

get_policy [-terse] *policy*

Displays the values of the named policy. Requires the inquire privilege. With the -terse flag, outputs the fields as quoted strings separated by tabs. Aliased by `getpol`.

Examples:

```
kadmin: get_policy admin
Policy: admin
Maximum password life: 180 days 00:00:00
Minimum password life: 00:00:00
Minimum password length: 6
Minimum number of password character classes: 2
Number of old keys kept: 5
Reference count: 17
```

```

Maximum password failures before logout: 3
Password failure count reset interval: 180
Password logout duration: 60
kadmin: get_policy -terse admin
admin admin    15552000    0    6    2    5    17 3 180 60
kadmin:

```

Errors:

KADM5_AUTH_GET (requires the get privilege)

KADM5_UNK_POLICY (policy does not exist)

list_policies [*expression*]

Retrieves all or some policy names. *expression* is a shell-style glob expression that can contain the wild-card characters `?`, `*`, and `[]`'s. All policy names matching the expression are printed. If no expression is provided, all existing policy names are printed. Requires the `list` privilege. Aliased by `listpols`, `get_policies`, and `getpols`.

Examples:

```

kadmin: listpols
test-pol dict-only once-a-min test-pol-nopw
kadmin: listpols t*
test-pol test-pol-nopw kadmin:

```

ktadd [-k *keytab*] [-q] [-e *keysaltlist*] [-norandkey] [[*principal* | -glob *princ-exp*]] [...]

Adds a principal or all principals matching *princ-exp* to a keytab. It randomizes each principal's key in the process, to prevent a compromised admin account from reading out all of the keys from the database. However, `kadmin.local` has the `-norandkey` option, which leaves the keys and their version numbers unchanged, similar to the Kerberos V4 `ext_srvtab` command. That allows users to continue to use the passwords they know to login normally, while simultaneously allowing scripts to login to the same account using a keytab. There is no significant security risk added since `kadmin.local` must be run by root on the KDC anyway.

`ktadd` requires the `inquire` and `changepw` privileges. An entry for each of the principal's unique encryption types is added, ignoring multiple keys with the same encryption type but different salt types. If the `-k` argument is not specified, the default keytab file, `/etc/krb5/krb5.keytab`, is used.

The “`-e enctype:salt`” option overrides the list of *enctypes* given in `krb5.conf(4)`, in the `permitted_enctypes` parameter. If “`-e enctype:salt`” is not used and `permitted_enctypes` is not defined in `krb5.conf(4)`, a key for each *enctype* supported by the system on which `kadmin` is run will be created and added to the keytab. Restricting the *enctypes* of keys in the keytab is useful when the system for which keys are being created does not support the same set of *enctypes* as the KDC. Note that `ktadd` modifies the *enctype* of the keys in the principal database as well.

If the `-q` option is specified, less status information is displayed. Aliased by `xst`. The `-glob` option requires the `list` privilege. Also, note that if you use `-glob` to create a keytab, you need to remove `/etc/krb5/kadm5.keytab` and create it again if you want to use `-p */admin` with `kadmin`.

`princ-exp`

princ-exp follows the same rules described for the `list_principals` command.

Example:

```
kadmin: ktadd -k /tmp/new-keytab nfs/chicago
Entry for principal nfs/chicago with kvno 2,
encryption type DES-CBC-CRC added to keytab
WRFILE:/tmp/new-keytab.
kadmin:
```

`ktremove [-k keytab] [-q] principal [kvno | all | old]`

Removes entries for the specified principal from a keytab. Requires no privileges, since this does not require database access. If `all` is specified, all entries for that principal are removed; if `old` is specified, all entries for that principal except those with the highest kvno are removed. Otherwise, the value specified is parsed as an integer, and all entries whose kvno match that integer are removed. If the `-k` argument is not specified, the default keytab file, `/etc/krb5/krb5.keytab`, is used. If the `-q` option is specified, less status information is displayed. Aliased by `ktrem`.

Example:

```
kadmin: ktremove -k /tmp/new-keytab nfs/chicago
Entry for principal nfs/chicago with kvno 2
removed from keytab
WRFILE:/tmp/new-keytab.
kadmin:
```

`quit`

Quits `kadmin`. Aliased by `exit` and `q`.

Time Formats

Various commands in `kadmin` can take a variety of time formats, specifying time durations or absolute times. The `kadmin` option variables *maxrenewlife*, *maxlife*, and *minlife* are time durations, whereas *expdate* and *pwexpdate* are absolute times.

Examples:

```
kadmin: modprinc -expire "12/31 7pm" jdb
kadmin: modprinc -maxrenewlife "2 fortnight" jdb
kadmin: modprinc -pexpire "this sunday" jdb
kadmin: modprinc -expire never jdb
kadmin: modprinc -maxlife "7:00:00pm tomorrow" jdb
```

Note that times which do not have the “ago” specifier default to being absolute times, unless they appear in a field where a duration is expected. In that case, the time specifier will be interpreted as relative. Specifying “ago” in a duration can result in unexpected behavior.

The following time formats and units can be combined to specify a time. The time and date format examples are based on the date and time of July 2, 1999, 1:35:30 p.m.

Time Format	Examples
<i>hh[:mm][:ss][am/pm/a.m./p.m.]</i>	1p.m., 1:35, 1:35:30pm

Variable	Description
<i>hh</i>	hour (12-hour clock, leading zero permitted but not required)
<i>mm</i>	minutes
<i>ss</i>	seconds

Date Format	Examples
<i>mm/dd[/yy]</i>	07/02, 07/02/99
<i>yyyy-mm-dd</i>	1999-07-02
<i>dd-month-yyyy</i>	02-July-1999
<i>month [,yyyy]</i>	Jul 02, July 02,1999
<i>dd month[yyyy]</i>	02 JULY, 02 july 1999

Variable	Description
<i>dd</i>	day
<i>mm</i>	month
<i>yy</i>	year within century (00-38 is 2000 to 2038; 70-99 is 1970 to 1999)
<i>yyyy</i>	year including century
<i>month</i>	locale's full or abbreviated month name

Time Units	Examples
[+ - #] year	"-2 year"
[+ - #] month	"2 months"
[+ - #] fortnight	

[+ - #] week	
[+ - #] day	
[+ - #] hour	
[+ - #] minute	
[+ - #] min	
[+ - #] second	
[+ - #] sec	
tomorrow	
yesterday	
today	
now	
this	“this year”
last	“last saturday”
next	“next month”
sunday	
monday	
tuesday	
wednesday	
thursday	
friday	
saturday	
never	

You can also use the following time modifiers: first, second, third, fourth, fifth, sixth, seventh, eighth, ninth, tenth, eleventh, twelfth, and ago.

环境变量

See [environ\(5\)](#) for descriptions of the following environment variables that affect the execution of kadmin:

PAGER

The command to use as a filter for paging output. This can also be used to specify options. The default is [more\(1\)](#).

- 文件

/var/krb5/principal

Kerberos principal database.

/var/krb5/principal.u^log

The update log file for incremental propagation.

/var/krb5/principal.kadm5

Kerberos administrative database. Contains policy information.

/var/krb5/principal.kadm5.lock

Lock file for the Kerberos administrative database. This file works backwards from most other lock files (that is, kadmin will exit with an error if this file does *not* exist).

/var/krb5/kadm5.dict

Dictionary of strings explicitly disallowed as passwords.

/etc/krb5/kadm5.acl

List of principals and their kadmin administrative privileges.

/etc/krb5/kadm5.keytab

Keytab for kadminind principals: kadmin/*fqdn*, changepw/*fqdn*, and kadmin/changepw.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/security/kerberos-5
Interface Stability	Committed

另请参见

[kpasswd\(1\)](#), [more\(1\)](#), [gkadmin\(1M\)](#), [kadminind\(1M\)](#), [kdb5_util\(1M\)](#), [kdb5_ldap_util\(1M\)](#), [kproplog\(1M\)](#), [kadm5.acl\(4\)](#), [kdc.conf\(4\)](#), [krb5.conf\(4\)](#), [attributes\(5\)](#), [environ\(5\)](#), [kerberos\(5\)](#), [krb5envvar\(5\)](#)

History

The kadmin program was originally written by Tom Yu at MIT, as an interface to the OpenVision Kerberos administration program.

诊断

The kadmin command is currently incompatible with the MIT kadminind daemon interface, so you cannot use this command to administer an MIT-based Kerberos database. However, clients running the Solaris implementation of Kerberos can still use an MIT-based KDC.

引用名	kadmind – Kerberos administration daemon
用法概要	<pre>/usr/lib/krb5/kadmind [-nofork] [-m] [-port <i>port-number</i>] [-r <i>realm</i>] [-x <i>db_args</i>] [-P <i>pid_file</i>]</pre>
描述	kadmind runs on the master key distribution center (KDC), which stores the principal and policy databases. kadmind accepts remote requests to administer the information in these databases. Remote requests are sent, for example, by kpasswd(1), gkadmin(1M), and kadmin(1M) commands, all of which are clients of kadmind. When you install a KDC, kadmind is set up in the <code>init</code> scripts to start automatically when the KDC is rebooted. kadmind requires a number of configuration files to be set up for it to work: <code>/etc/krb5/kdc.conf</code> The KDC configuration file contains configuration information for the KDC and the Kerberos administration system. kadmind understands a number of configuration variables (called relations) in this file, some of which are mandatory and some of which are optional. In particular, kadmind uses the <code>acl_file</code>, <code>dict_file</code>, <code>admin_keytab</code>, and <code>kadmind_port</code> relations in the <code>[realms]</code> section. Refer to the kdc.conf(4) man page for information regarding the format of the KDC configuration file. <code>/etc/krb5/kadm5.keytab</code> kadmind requires a keytab (key table) containing correct entries for the <code>kadmin/fqdn</code>, <code>kadmin/changepw</code> and <code>kadmin/changepw</code> principals for every realm that kadmind answers requests. The keytab can be created with the kadmin(1M) or kdb5_util(1M) command. The location of the keytab is determined by the <code>admin_keytab</code> relation in the <code>kdc.conf(4)</code> file. <code>/etc/krb5/kadm5.acl</code> kadmind uses an ACL (access control list) to determine which principals are allowed to perform Kerberos administration actions. The path of the ACL file is determined by the <code>acl_file</code> relation in the <code>kdc.conf</code> file. See kdc.conf(4). For information regarding the format of the ACL file, refer to kadm5.acl(4). The kadmind daemon will need to be restarted to reread the <code>kadm5.acl</code> file after it has been modified. You can do this, as root, with the following command: <pre># svcadm restart svc:/network/security/kadmin:default</pre> After kadmind begins running, it puts itself in the background and disassociates itself from its controlling terminal. kadmind can be configured for incremental database propagation. Incremental propagation allows slave KDC servers to receive principal and policy updates incrementally instead of receiving full dumps of the database. These settings can be changed in the kdc.conf(4) file: <pre>sunw_dbprop_enable = [true false]</pre> Enable or disable incremental database propagation. Default is false.

`sunw_dbprop_master_ulogsize = N`

Specifies the maximum amount of log entries available for incremental propagation to the slave KDC servers. The maximum value that this can be is 2500 entries. Default value is 1000 entries.

The `kiprop/<hostname>@<REALM>` principal must exist in the master's `kadm5.keytab` file to enable the slave to authenticate incremental propagation from the master. In the principal syntax above, `<hostname>` is the master KDC's host name and `<REALM>` is the realm in which the master KDC resides.

Kerberos client machines can automatically migrate Unix users to the default Kerberos realm specified in the local `krb5.conf(4)`, if the user does not have a valid kerberos account already. You achieve this by using the `pam_krb5_migrate(5)` service module for the service in question. The Kerberos service principal used by the client machine attempting the migration needs to be validated using the `u` privilege in `kadm5.acl(4)`. When using the `u` privilege, `kadmind` validates user passwords using PAM, specifically using a `PAM_SERVICE` name of `k5migrate` by calling `pam_authenticate(3PAM)` and `pam_acct_mgmt(3PAM)`.

A suitable PAM stack configuration example for `k5migrate` would look like:

```
k5migrate      auth      required      pam_unix_auth.so.1
k5migrate      account  required      pam_unix_account.so.1
```

选项

The following options are supported:

`-nofork`

Specifies that `kadmind` does not put itself in the background and does not disassociate itself from the terminal. In normal operation, you should use the default behavior, which is to allow the daemon to put itself in the background.

`-m`

Specifies that the master database password should be retrieved from the keyboard rather than from the stash file. When using `-m`, the `kadmind` daemon receives the password prior to putting itself in the background. If used in combination with the `-d` option, you must explicitly place the daemon in the background.

`-port port-number`

Specifies the port on which the `kadmind` daemon listens for connections. The default is controlled by the `kadmind_port` relation in the `kdc.conf(4)` file.

`-P pid_file`

Specifies the file to which the PID of `kadmind` process should be written to after it starts up. This can be used to identify whether `kadmind` is still running and to allow `init` scripts to stop the correct process.

`-r realm`

Specifies the default realm that `kadmind` serves. If `realm` is not specified, the default `realm` of the host is used. `kadmind` answers requests for any realm that exists in the local KDC database and for which the appropriate principals are in its `keytab`.

-x *db_args*
Pass database-specific arguments to kadmind. Supported arguments are for LDAP and the Berkeley-db2 plug-in. These arguments are:

binddn=binddn
LDAP simple bind DN for authorization on the directory server. Overrides the `ldap_kadmind_dn` parameter setting in `krb5.conf(4)`.

bindpwd=bindpwd
Bind password.

dbname=name
For the Berkeley-db2 plug-in, specifies a name for the Kerberos database.

nconns=num
Maximum number of server connections.

port=num
Directory server connection port.

/var/krb5/principal
Kerberos principal database.

/var/krb5/principal.ulong
The update log file for incremental propagation.

/var/krb5/principal.kadm5
Kerberos administrative database containing policy information.

/var/krb5/principal.kadm5.lock
Kerberos administrative database lock file. This file works backwards from most other lock files (that is, kadmind exits with an error if this file does not exist).

/var/krb5/kadm5.dict
Dictionary of strings explicitly disallowed as passwords.

/etc/krb5/kadm5.acl
List of principals and their kadmind administrative privileges.

/etc/krb5/kadm5.keytab
Keytab for kadmind principals: `kadmind/fqdn`, `changepw/fqdn`, and `kadmind/changepw`.

/etc/krb5/kdc.conf
KDC configuration information.

文件

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/security/kerberos-5

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

另请参见

```
kpasswd(1),svcs(1),gkadmin(1M),kadmin(1M),kadmin(1M),kdb5_util(1M),
kdb5_ldap_util(1M),kproplog(1M),svcadm(1M),pam_acct_mgmt(3PAM),
pam_authenticate(3PAM),kadm5.acl(4),kdc.conf(4),krb5.conf(4),attributes(5),
kerberos(5),krb5envvar(5),pam_krb5_migrate(5),smf(5)
```

附注

The Kerberos administration daemon (`kadmind`) is now compliant with the change-password standard mentioned in RFC 3244, which means it can now handle change-password requests from non-Solaris Kerberos clients.

The `kadmind` service is managed by the service management facility, `smf(5)`, under the service identifier:

```
svc:/network/security/kadmin
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using `svcadm(1M)`. The service's status can be queried using the `svcs(1)` command.

The `-d` and `-p` arguments are made obsolete with the `-no-fork` and `-port` arguments, respectively. The `-d` and `-p` arguments might be removed in a future release of the Solaris operating system.

引用名 kcfd – kernel-level cryptographic framework daemon

用法概要 kcfd

描述 The kcfd daemon helps in managing CPU usage by cryptographic operations performed in software by kernel threads. The system utilization associated with these threads is charged to the kcfd process. It also does module verification for kernel cryptographic modules.

Only a privileged user can run this daemon.

The kcfd daemon is automatically invoked in run level 1, after /usr is mounted. A previously invoked kcfd daemon that is still running must be stopped before invoking another kcfd command.

Manually starting and restarting kcfd is not recommended. If it is necessary to do so, use the [cryptoadm\(1M\)](#) start and stop subcommands.

退出状态 The following exit values are returned:

0 Daemon started successfully.

> 1 Daemon failed to start.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [cryptoadm\(1M\)](#), [attributes\(5\)](#)

引用名

kclient – set up a machine as a Kerberos client

用法概要

```
/usr/sbin/kclient [-n] [-R realm] [-k kdc] [-a adminuser]
                  [-c filepath] [-d dnsarg] [-f fqdn_list] [-h logical_host_name]
                  [-k kdc_list] [-m master_kdc] [-p profile] [-s pam_service]
                  [-T kdc_vendor]
```

描述

By specifying the various command options, you can use the `kclient` utility to:

- Configure a machine as a Kerberos client for a specified realm and for KDC by setting up `krb5.conf(4)`.
- Add the Kerberos host principal to the local host's keytab file (`/etc/krb5/krb5.keytab`).
- Set up the machine to do kerberized NFS.
- Bring over a master `krb5.conf` copy from a specified pathname.
- Setup a machine to do server and/or host/domain name-to-realm mapping lookups by means of DNS.
- Configure a Kerberos client to use an MS Active Directory server. This generates a keytab file with the Kerberos client's service keys populated.
- Setup a Kerberos client that has no service keys. This is useful when the client does not require service keys, because the client does not wish to host a service that uses Kerberos for security.
- Configure a Kerberos client that is part of a cluster. This option requires the logical host name of the cluster so that the proper service keys are created and populated in the client's keytab file.
- Setup a Kerberos client to join an environment that consists of Kerberos servers that are non-Solaris and non-MS Active Directory servers.
- Configure `pam.conf(4)` to use Kerberos authentication for specified services.
- Configure the client as a simple NTP broadcast/multicast client.
- Specify custom domain/host name-to-realm name mappings.
- Setup the Kerberos client to use multiple KDC servers.

The `kclient` utility needs to be run on the client machine with root permission and can be run either interactively or non-interactively. In the non-interactive mode, the user feeds in the required inputs by means of a profile, command-line options, or a combination of profile and command-line options. The user is prompted for “required” parameter values (`realm` and `adminuser`), if found missing in the non-interactive run. The interactive mode is invoked when the utility is run without any command-line arguments.

Both the interactive and non-interactive forms of `kclient` can add the `host/fqdn` entry to the local host's keytab file. They also can require the user to enter the password for the administrative user requested, to obtain the Kerberos Ticket Granting Ticket (TGT) for `adminuser`. The `host/fqdn`, `nfs/fqdn`, and `root/fqdn` principals can be added to the KDC database (if not already present) before their possible addition to the local host's keytab.

The `kclient` utility assumes that the local host has been setup for DNS and requires the presence of a valid `resolv.conf(4)`. Also, `kclient` can fail if the localhost time is not synchronized with that of the KDC. For Kerberos to function the localhost time must be within five minutes of that of the KDC. It is advised that both systems run some form of time synchronization protocol, such as the Network Time Protocol (NTP). See the `ntpd` man page, delivered in the `SUNWntpu` package (not a SunOS man page).

选项

The non-interactive mode supports the following options:

`-n`

Set up the machine for kerberized NFS. This involves making changes to `krb5*` security flavors in `nfssec.conf(4)`. This option will also add `nfs/fqdn` and `root/fqdn` entries to the local host's keytab file if the `-K` option has not been specified.

`-R [ realm ]`

Specifies the Kerberos realm.

`-k kdc_list`

The `-k` option specifies the KDC host names for the Kerberos client. `kdc_list` is a comma-separated list of KDCs. If the `-m` option is not used, it is assumed that the first (or only) host in `kdc_list` is the master KDC host name. Note that the list specified is used verbatim. This is helpful when specifying non-fully qualified KDC host names that can be canonicalized by DNS.

`-a [ adminuser ]`

Specifies the Kerberos administrative user.

`-T kdc_vendor`

Configure the Kerberos client to associate with a third party server. Valid `kdc_vendor` currently supported are:

`ms_ad`

Microsoft Active Directory

`mit`

MIT KDC server

`heimdal`

Heimdal KDC server

`shishi`

Shishi KDC server

Knowing the administrative password will be required to associate the client with the server if the `ms_ad` option is specified.

`-c [ filepath ]`

Specifies the pathname to the `krb5.conf(4)` master file, to be copied over to the local host. The path specified normally points to a master copy on a remote host and brought over to the local host by means of NFS.

-d [*dnsarg*]

Specifies the DNS lookup option to be used and specified in the `krb5.conf(4)` file. Valid *dnsarg* entries are: `none`, `dns_lookup_kdc`, `dns_lookup_realm` and `dns_fallback`. Any other entry is considered invalid. The latter three *dnsarg* values assume the same meaning as those described in `krb5.conf`. `dns_lookup_kdc` implies DNS lookups for the KDC and the other servers. `dns_lookup_realm` is for host/domain name-to-realm mapping by means of DNS. `dns_fallback` is a superset and does DNS lookups for both the servers and the host/domain name-to-realm mapping. A lookup option of `none` specifies that DNS is not be used for any kind of mapping lookup.

-D *domain_list*

Specifies the host and/or domain names to be mapped to the Kerberos client's default realm name. *domain_list* is a comma-separated list, for example "`example.com,host1.example.com`". If the `-D` option is not used, then only the client's domain is used for this mapping. For example, if the client is `host1.eng.example.com`, then the domain that is mapped to the `EXAMPLE.COM` realm is `example.com`.

-K

Configure the Kerberos client without service keys, which are usually stored in `/etc/krb5/krb5.keytab`. This is useful in the following scenarios:

- The client IP address is dynamically assigned and therefore does not host Kerberized services.
- Client has a static IP address, but does not want to host any Kerberized services.
- Client has a static IP address, but the local administrator does not currently have service keys available for the machine. It is expected that, at a later time, these keys will be installed on the machine.

-f [*fqdn_list*]

This option creates a service principal entry (`host/nfs/root`) associated with each of the listed *fqdn*'s, if required, and subsequently adds the entries to the local host's `keytab`.

fqdn_list is a comma-separated list of one or more fully qualified DNS domain names.

This option is especially useful in Kerberos realms having systems offering kerberized services, but situated in multiple different DNS domains.

-h *logical_host_name*

Specifies that the Kerberos client is a node in a cluster. The *logical_host_name* is the logical host name given to the cluster. The resulting `/etc/krb5/krb5.conf` and `/etc/krb5/krb5.keytab` files must be manually copied over to the other members of the cluster.

-m *master_kdc*

This option specifies the master KDC to be used by the Kerberos client. *master_kdc* is the host name of the master KDC for the client. If the `-m` option is not used, then it is assumed that the first KDC host name listed with the `-k` option is the master KDC.

-p [*profile*]

Specifies the profile to be used to enable the reading in of the values of all the parameters required for setup of the machine as a Kerberos client.

The profile should have entries in the format:

PARAM <*value*>

Valid *PARAM* entries are: REALM, KDC, ADMIN, FILEPATH, NFS, DNSLOOKUP, FQDN, NOKEY, NOSOL, LHN, KDCVENDOR, RMAP, MAS, and PAM.

These profile entries correspond to the **-R** [*realm*], **-k** [*kdc*], **-a** [*adminuser*], **-c** [*filepath*], **-n**, **-d** [*dnsarg*], **-f** [*fqdn_list*], **-K**, **-h** [*logical_host_name*], **-T** [*kdc_vendor*], **-D** [*domain_list*], **-m** [*master_kdc*], and **-s** [*pam_service*] command-line options, respectively. Any other *PARAM* entry is considered invalid and is ignored.

The NFS profile entry can have a value of 0 (do nothing) or 1 (operation is requested). Any other value is considered invalid and is ignored.

Keep in mind that the command line options override the *PARAM* values listed in the profile.

-s *pam_service:auth_type*[,...]

Specifies that the PAM service names, listed in *pam_service*, are authenticated through Kerberos. Using this option updates [pam.conf\(4\)](#) to include a separate authentication stack with [pam_krb5\(5\)](#). Examples of *pam_service* names are *sshd-kbdint*, *xscreensaver*, and so forth.

auth_type can be one of the following keywords:

first

Try authenticating through Kerberos first. If this fails try to authenticate through Unix.

only

Try to authenticate only through Kerberos.

optional

Try authenticating through Unix first. If this is successful try to authenticate through Kerberos.

示例

示例 1 Setting Up a Kerberos Client Using Command-Line Options

To setup a Kerberos client using the `clntconfig/admin` administrative principal for realm 'ABC.COM', `kdc 'example1.com'` and that also does kerberized NFS, enter:

```
# /usr/sbin/kclient -n -R ABC.COM -k example1.com -a clntconfig
```

Alternatively, to set up a Kerberos client using the `clntconfig/admin` administrative principal for the realm 'EAST.ABC.COM', `kdc 'example2.east.abc.com'` and that also needs service principal(s) created and/or added to the local keytab for multiple DNS domains, enter:

示例 1 Setting Up a Kerberos Client Using Command-Line Options (续)

```
# /usr/sbin/kclient -n -R EAST.ABC.COM -k example2.east.abc.com \
-f west.abc.com,central.abc.com -a clntconfig
```

Note that the krb5 administrative principal used by the administrator needs to have only add, inquire, change-pwd and modify privileges (for the principals in the KDC database) in order for the kclient utility to run. A sample [kadm5.acl\(4\)](#) entry is:

```
clntconfig/admin@ABC.COM acmi
```

示例 2 Setting Up a Kerberos Client Using the Profile Option

To setup a Kerberos client using the clntconfig/admin administrative principal for realm 'ABC.COM', kdc 'example1.com' and that also copies over the master krb5.conf from a specified location, enter:

```
# /usr/sbin/kclient -p /net/example1.com/export/profile.krb5
```

The contents of profile.krb5:

```
REALM ABC.COM
KDC example1
ADMIN clntconfig
FILEPATH /net/example1.com/export/krb5.conf
NFS 0
DNSLOOKUP none
```

示例 3 Setting Up a Kerberos Client That Has a Dynamic IP Address

In this example a Kerberos client is a DHCP client that has a dynamic IP address. This client does not wish to host any Kerberized services and therefore does not require a keytab (/etc/krb5/krb5.keytab) file.

For this type of client the administrator would issue the following command to configure this machine to be a Kerberos client of the ABC.COM realm with the KDC server kdc1.example.com:

```
# /usr/sbin/kclient -K -R EXAMPLE.COM -k kdc1.example.com
```

```
/etc/krb5/kadm5.acl
```

Kerberos access control list (ACL) file.

```
/etc/krb5/krb5.conf
```

Default location for the local host's configuration file.

```
/etc/krb5/krb5.keytab
```

Default location for the local host's keytab file.

```
/etc/nfssec.conf
```

File listing NFS security modes.

文件

/etc/resolv.conf
DNS resolver configuration file.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/security/kerberos-5
Interface Stability	Committed

另请参见 [encrypt\(1\)](#), [ksh\(1\)](#), [ldapdelete\(1\)](#), [ldapmodify\(1\)](#), [ldapsearch\(1\)](#), [dd\(1M\)](#), [smbadm\(1M\)](#), [kadm5.acl\(4\)](#), [krb5.conf\(4\)](#), [nfssec.conf\(4\)](#), [pam.conf\(4\)](#), [resolv.conf\(4\)](#), [attributes\(5\)](#), [pam_krb5\(5\)](#)

附注 fqdn stands for the Fully Qualified Domain Name of the local host. The `kclient` utility saves copies of both the [krb5.conf\(4\)](#) and [nfssec.conf\(4\)](#) files to files with corresponding names and `.sav` extensions. The optional copy of the [krb5.conf\(4\)](#) master file is neither encrypted nor integrity-protected and it takes place over regular NFS.

引用名	kdb5_ldap_util – Kerberos configuration utility
用法概要	<code>kdb5_ldap_util [-D <i>user_dn</i> [-w <i>passwd</i>]] [-H <i>ldap_uri</i>] <i>command</i> [<i>command_options</i>]</code>
描述	The <code>kdb5_ldap_util</code> utility allows an administrator to manage realms, Kerberos services, and ticket policies. The utility offers a set of general options, described under <code>OPTIONS</code> , and a set of commands, which, in turn, have their own options. Commands and their options are described in their own subsections, below.
选项	<code>kdb5_ldap_util</code> has a small set of general options that apply to the <code>kdb5_ldap_util</code> utility itself and a larger number of options that apply to specific commands. A number of these command-specific options apply to multiple commands and are described in their own section, below.
General Options	The following general options are supported: <code>-D <i>user_dn</i></code> Specifies the distinguished name (DN) of a user who has sufficient rights to perform the operation on the LDAP server. <code>-H <i>ldap_uri</i></code> Specifies the URI of the LDAP server. <code>-w <i>passwd</i></code> Specifies the password of <i>user_dn</i>. This option is not recommended.
Common Command-specific Options	The following options apply to a number of <code>kdb5_ldap_util</code> commands. <code>-subtrees <i>subtree_dn_list</i></code> Specifies the list of subtrees containing the principals of a realm. The list contains the DNs of the subtree objects separated by a colon. <code>-sscope <i>search_scope</i></code> Specifies the scope for searching the principals under a subtree. The possible values are 1 or one (one level), 2 or sub (subtrees). <code>-containerref <i>container_reference_dn</i></code> Specifies the DN of the container object in which the principals of a realm will be created. If the container reference is not configured for a realm, the principals will be created in the realm container. <code>-maxttl <i>max_ticket_life</i></code> Specifies maximum ticket life for principals in this realm. <code>-maxrenewlife <i>max_renewable_ticket_life</i></code> Specifies maximum renewable life of tickets for principals in this realm. <code>-r <i>realm</i></code> Specifies the Kerberos realm of the database; by default the realm returned by <code>krb5_default_local_realm(3)</code> is used.

kdb5_ldap_util Commands The kdb5_ldap_util utility comprises a set of commands, each with its own set of options. These commands are described in the following subsections.

The create Command The create command creates a realm in a directory. The command has the following syntax:

```
create \
[-subtrees subtree_dn_list]
[-sscope search_scope]
[-containerref container_reference_dn]
[-k mkeytype]
[-m|-P password] -sf stashfilename]
[-s]
[-r realm]
[-maxttl life max_ticket_life]
[-kdc dn kdc_service_list]
[-admin dn admin_service_list]
[-maxrenew life max_renewable_ticket_life]
[ticket_flags]
```

The create command has the following options:

- subtree *subtree_dn_list*
See “Common Command-specific Options,” above.
- sscope *search_scope*
See “Common Command-specific Options,” above.
- containerref *container_reference_dn*
See “Common Command-specific Options,” above.
- k *mkeytype*
Specifies the key type of the master key in the database; the default is that given in `kdc.conf(4)`.
- m
Specifies that the master database password should be read from the TTY rather than fetched from a file on the disk.
- P *password*
Specifies the master database password. This option is not recommended.
- sf *stashfilename*
Specifies the stash file of the master database password.
- s
Specifies that the stash file is to be created.
- maxttl life *max_ticket_life*
See “Common Command-specific Options,” above.
- maxrenew life *max_renewable_ticket_life*
See “Common Command-specific Options,” above.

-r realm

See “Common Command-specific Options,” above.

ticket_flags

Specifies the ticket flags. If this option is not specified, by default, none of the flags are set. This means all the ticket options will be allowed and no restriction will be set. See “Ticket Flags” for a list and descriptions of these flags.

The **modify** Command The **modify** command modifies the attributes of a realm. The command has the following syntax:

```
modify \
[-subtrees subtree_dn_list]
[-sscope search_scope]
[-containerref container_reference_dn]
[-r realm]
[-maxtktlife max_ticket_life]
[-maxrenewlife max_renewable_ticket_life]
[ticket_flags]
```

The **modify** command has the following options:

-subtree subtree_dn_list

See “Common Command-specific Options,” above.

-sscope search_scope

See “Common Command-specific Options,” above.

-containerref container_reference_dn

See “Common Command-specific Options,” above.

-maxtktlife max_ticket_life

See “Common Command-specific Options,” above.

-maxrenewlife max_renewable_ticket_life

See “Common Command-specific Options,” above.

-r realm

See “Common Command-specific Options,” above.

ticket_flags

Specifies the ticket flags. If this option is not specified, by default, none of the flags are set. This means all the ticket options will be allowed and no restriction will be set. See “Ticket Flags” for a list and descriptions of these flags.

The **view** Command The **view** command displays the attributes of a realm. The command has the following syntax:

```
view [-r realm]
```

The **view** command has the following option:

	<i>-r realm</i> See “Common Command-specific Options,” above.
The destroy Command	The <code>destroy</code> command destroys a realm, including the master key stash file. The command has the following syntax: <pre>destroy [-f] [-r realm]</pre> The <code>destroy</code> command has the following options: <i>-f</i> If specified, <code>destroy</code> does not prompt you for confirmation. <i>-r realm</i> See “Common Command-specific Options,” above.
The list Command	The <code>list</code> command displays the names of realms. The command has the following syntax: <pre>list</pre> The <code>list</code> command has no options.
The stashsrpw Command	The <code>stashsrpw</code> command enables you to store the password for service object in a file so that a KDC and Administration server can use it to authenticate to the LDAP server. The command has the following syntax: <pre>stashsrpw [-f filename] servicedn</pre> The <code>stashsrpw</code> command has the following option and argument: <i>-f filename</i> Specifies the complete path of the service password file. The default is: <pre>/var/krb5/service_passwd</pre> <i>servicedn</i> Specifies the distinguished name (DN) of the service object whose password is to be stored in file.
The create_policy Command	The <code>create_policy</code> command creates a ticket policy in a directory. The command has the following syntax: <pre>create_policy \ [-r realm] [-maxtktlife max_ticket_life] [-maxrenewlife max_renewable_ticket_life] [ticket_flags] policy_name</pre> The <code>create_policy</code> command has the following options: <i>-r realm</i> See “Common Command-specific Options,” above.

-maxttl *max_ticket_life*

See “Common Command-specific Options,” above.

-maxrenewlife *max_renewable_ticket_life*

See “Common Command-specific Options,” above.

ticket_flags

Specifies the ticket flags. If this option is not specified, by default, none of the flags are set. This means all the ticket options will be allowed and no restriction will be set. See “Ticket Flags” for a list and descriptions of these flags.

policy_name

Specifies the name of the ticket policy.

The *modify_policy*
Command

The *modify_policy* command modifies the attributes of a ticket policy. The command has the following syntax:

```
modify_policy \  
[-r realm]  
[-maxttl max_ticket_life]  
[-maxrenewlife max_renewable_ticket_life]  
[ticket_flags]  
policy_name
```

The *modify_policy* command has the same options and argument as those for the *create_policy* command.

The *view_policy*
Command

The *view_policy* command displays the attributes of a ticket policy. The command has the following syntax:

```
view_policy [-r realm] policy_name
```

The *view_policy* command has the following options:

-r realm

See “Common Command-specific Options,” above.

policy_name

Specifies the name of the ticket policy.

The *destroy_policy*
Command

The *destroy_policy* command destroys an existing ticket policy. The command has the following syntax:

```
destroy_policy [-r realm] [-force] policy_name
```

The *destroy_policy* command has the following options:

-r realm

See “Common Command-specific Options,” above.

-force

Forces the deletion of the policy object. If not specified, you will be prompted for confirmation before the policy is deleted. Enter yes to confirm the deletion.

policy_name

Specifies the name of the ticket policy.

The `list_policy`
Command

The `list_policy` command lists the ticket policies in the default or a specified realm. The command has the following syntax:

```
list_policy [-r realm]
```

The `list_policy` command has the following option:

-r *realm*

See “Common Command-specific Options,” above.

Ticket Flags

A number of `kdb5_ldap_util` commands have `ticket_flag` options. These flags are described as follows:

{-|+}allow_dup_skey

-allow_dup_skey disables user-to-user authentication for principals by prohibiting principals from obtaining a session key for another user. This setting sets the `KRB5_KDB_DISALLOW_DUP_SKEY` flag. +allow_dup_skey clears this flag.

{-|+}allow_forwardable

-allow_forwardable prohibits principals from obtaining forwardable tickets. This setting sets the `KRB5_KDB_DISALLOW_FORWARDABLE` flag. +allow_forwardable clears this flag.

{-|+}allow_postdated

-allow_postdated prohibits principals from obtaining postdated tickets. This setting sets the `KRB5_KDB_DISALLOW_POSTDATED` flag. +allow_postdated clears this flag.

{-|+}allow_proxiabile

-allow_proxiabile prohibits principals from obtaining proxiabile tickets. This setting sets the `KRB5_KDB_DISALLOW_PROXIABLE` flag. +allow_proxiabile clears this flag.

{-|+}allow_renewable

-allow_renewable prohibits principals from obtaining renewable tickets. This setting sets the `KRB5_KDB_DISALLOW_RENEWABLE` flag. +allow_renewable clears this flag.

{-|+}allow_svr

-allow_svr prohibits the issuance of service tickets for principals. This setting sets the `KRB5_KDB_DISALLOW_SVR` flag. +allow_svr clears this flag.

{-|+}allow_tgs_req

-allow_tgs_req specifies that a Ticket-Granting Service (TGS) request for a service ticket for principals is not permitted. This option is useless for most purposes. +allow_tgs_req clears this flag. The default is +allow_tgs_req. In effect, -allow_tgs_req sets the `KRB5_KDB_DISALLOW_TGT_BASED` flag on principals in the database.

{-|+}allow_tix

-allow_tix forbids the issuance of any tickets for principals. +allow_tix clears this flag. The default is +allow_tix. In effect, -allow_tix sets the KRB5_KDB_DISALLOW_ALL_TIX flag on principals in the database.

{-|+}needchange

+needchange sets a flag in the attributes field to force a password change; -needchange clears that flag. The default is -needchange. In effect, +needchange sets the KRB5_KDB_REQUIRES_PWCHANGE flag on principals in the database.

{-|+}password_changing_service

+password_changing_service sets a flag in the attributes field marking a principal as a password-change-service principal (a designation that is most often not useful). -password_changing_service clears the flag. That this flag has a long name is intentional. The default is -password_changing_service. In effect, +password_changing_service sets the KRB5_KDB_PWCHANGE_SERVICE flag on principals in the database.

{-|+}requires_hwauth

+requires_hwauth requires principals to preauthenticate using a hardware device before being allowed to [kinit\(1\)](#). This setting sets the KRB5_KDB_REQUIRES_HW_AUTH flag. -requires_hwauth clears this flag.

{-|+}requires_preauth

+requires_preauth requires principals to preauthenticate before being allowed to [kinit\(1\)](#). This setting sets the KRB5_KDB_REQUIRES_PRE_AUTH flag. -requires_preauth clears this flag.

示例

示例1 Using create

The following is an example of the use of the create command.

```
# kdb5_ldap_util -D cn=admin,o=org -H ldaps://ldap-server1.mit.edu \
create -subtrees o=org -sscope SUB -r ATHENA.MIT.EDU
Password for "cn=admin,o=org": password entered
Initializing database for realm 'ATHENA.MIT.EDU'
You will be prompted for the database Master Password.
It is important that you NOT FORGET this password.
Enter KDC database master key: master key entered
Re-enter KDC database master key to verify: master key re-enteredjjjjjj
```

示例2 Using modify

The following is an example of the use of the modify command.

```
# kdb5_ldap_util -D cn=admin,o=org -H ldaps://ldap-server1.mit.edu \
modify +requires_preauth -r ATHENA.MIT.EDU
Password for "cn=admin,o=org": password entered
Password for "cn=admin,o=org": password entered
```

示例3 Using view

The following is an example of the use of the view command.

```
# kdb5_ldap_util -D cn=admin,o=org -H ldaps://ldap-server1.mit.edu \
view -r ATHENA.MIT.EDU
    Password for "cn=admin,o=org":
        Realm Name: ATHENA.MIT.EDU
        Subtree: ou=users,o=org
        Subtree: ou=servers,o=org
        SearchScope: ONE
        Maximum ticket life: 0 days 01:00:00
        Maximum renewable life: 0 days 10:00:00
        Ticket flags: DISALLOW_FORWARDABLE REQUIRES_PWCHANGE
```

示例4 Using destroy

The following is an example of the use of the destroy command.

```
# kdb5_ldap_util -D cn=admin,o=org -H ldaps://ldap-server1.mit.edu \
destroy -r ATHENA.MIT.EDU
Password for "cn=admin,o=org": password entered
Deleting KDC database of 'ATHENA.MIT.EDU', are you sure?
(type 'yes' to confirm)? yes
OK, deleting database of 'ATHENA.MIT.EDU'...
```

示例5 Using list

The following is an example of the use of the list command.

```
# kdb5_ldap_util -D cn=admin,o=org -H ldaps://ldap-server1.mit.edu list
Password for "cn=admin,o=org": password entered
Re-enter Password for "cn=admin,o=org": password re-entered
ATHENA.MIT.EDU
OPENLDAP.MIT.EDU
MEDIA-LAB.MIT.EDU
```

示例6 Using stashsrwpw

The following is an example of the use of the stashsrwpw command.

```
# kdb5_ldap_util stashsrwpw -f \
/home/andrew/conf_keyfile cn=service-kdc,o=org
Password for "cn=service-kdc,o=org": password entered
Re-enter password for "cn=service-kdc,o=org": password re-entered
```

示例7 Using create_policy

The following is an example of the use of the create_policy command.

```
# kdb5_ldap_util -D cn=admin,o=org -H ldaps://ldap-server1.mit.edu \
create_policy -r ATHENA.MIT.EDU \
-maxtktlife "1 day" -maxrenewlife "1 week" \
```

示例7 Using create_policy (续)

```
-allow_postdated +needchange -allow_forwardable tktpolicy
Password for "cn=admin,o=org": password entered
```

示例8 Using modify_policy

The following is an example of the use of the modify_policy command.

```
# kdb5_ldap_util -D cn=admin,o=org -H ldaps://ldap-server1.mit.edu \
modify_policy -r ATHENA.MIT.EDU \
-maxtktlife "60 minutes" -maxrenewlife "10 hours" \
+allow_postdated -requires_preauth tktpolicy
Password for "cn=admin,o=org": password entered
```

示例9 Using view_policy

The following is an example of the use of the view_policy command.

```
# kdb5_ldap_util -D cn=admin,o=org -H ldaps://ldap-server1.mit.edu \
view_policy -r ATHENA.MIT.EDU tktpolicy
Password for "cn=admin,o=org": password entered
Ticket policy: tktpolicy
Maximum ticket life: 0 days 01:00:00
Maximum renewable life: 0 days 10:00:00
Ticket flags: DISALLOW_FORWARDABLE REQUIRES_PWCHANGE
```

示例10 Using destroy_policy

The following is an example of the use of the destroy_policy command.

```
# kdb5_ldap_util -D cn=admin,o=org -H ldaps://ldap-server1.mit.edu \
destroy_policy -r ATHENA.MIT.EDU tktpolicy
Password for "cn=admin,o=org": password entered
This will delete the policy object 'tktpolicy', are you sure?
(type 'yes' to confirm)? yes
** policy object 'tktpolicy' deleted.
```

示例11 Using list_policy

The following is an example of the use of the list_policy command.

```
# kdb5_ldap_util -D cn=admin,o=org -H ldaps://ldap-server1.mit.edu \
list_policy -r ATHENA.MIT.EDU
Password for "cn=admin,o=org": password entered
tktpolicy
tmppolicy
userpolicy
```

示例12 Using setsrvpw

The following is an example of the use of the setsrvpw command.

示例 12 Using setsrvpw (续)

```
# kdb5_ldap_util setsrvpw -D cn=admin,o=org setsrvpw \
-fileonly -f /home/andrew/conf_keyfile cn=service-kdc,o=org
Password for "cn=admin,o=org": password entered
Password for "cn=service-kdc,o=org": password entered
Re-enter password for "cn=service-kdc,o=org": password re-entered
```

示例 13 Using create_service

The following is an example of the use of the create_service command.

```
# kdb5_ldap_util -D cn=admin,o=org create_service \
-kdc -randpw -f /home/andrew/conf_keyfile cn=service-kdc,o=org
Password for "cn=admin,o=org": password entered
File does not exist. Creating the file /home/andrew/conf_keyfile...
```

示例 14 Using modify_service

The following is an example of the use of the modify_service command.

```
# kdb5_ldap_util -D cn=admin,o=org modify_service \
-realm ATHENA.MIT.EDU cn=service-kdc,o=org
Password for "cn=admin,o=org": password entered
Changing rights for the service object. Please wait ... done
```

示例 15 Using view_service

The following is an example of the use of the view_service command.

```
# kdb5_ldap_util -D cn=admin,o=org view_service \
cn=service-kdc,o=org
Password for "cn=admin,o=org": password entered
      Service dn: cn=service-kdc,o=org
      Service type: kdc
      Service host list:
      Realm DN list: cn=ATHENA.MIT.EDU,cn=Kerberos,cn=Security
```

示例 16 Using destroy_service

The following is an example of the use of the destroy_service command.

```
# kdb5_ldap_util -D cn=admin,o=org destroy_service \
cn=service-kdc,o=org
Password for "cn=admin,o=org": password entered
This will delete the service object 'cn=service-kdc,o=org', are you sure?
(type 'yes' to confirm)? yes
** service object 'cn=service-kdc,o=org' deleted.
```

示例 17 Using list_service

The following is an example of the use of the list_service command.

示例 17 Using list_service (续)

```
# kdb5_ldap_util -D cn=admin,o=org list_service
Password for "cn=admin,o=org": password entered
cn=service-kdc,o=org
cn=service-adm,o=org
cn=service-pwd,o=org
```

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/security/kerberos-5
Interface Stability	Volatile

另请参见 [kinit\(1\)](#), [kadmin\(1M\)](#), [kdc.conf\(4\)](#), [attributes\(5\)](#)

引用名	kdb5_util – Kerberos Database maintenance utility
用法概要	<pre>/usr/sbin/kdb5_util [-d <i>dbname</i>] [-k <i>mkeytype</i>] [-kv <i>mkeyVNO</i>] [-m] [-M <i>mkeyname</i>] [-P <i>password</i>] [-r <i>realm</i>] [-sf <i>stashfilename</i>] [-x <i>db_args</i>]... <i>cmd</i></pre>
描述	The kdb5_util utility enables you to create, dump, load, and destroy the Kerberos V5 database. You can also use kdb5_util to create a stash file containing the Kerberos database master key.
选项	The following options are supported: -d <i>dbname</i> Specify the database name. .db is appended to whatever name is specified. You can specify an absolute path. If you do not specify the -d option, the default database name is /var/krb5/principal. -k <i>mkeytype</i> Specify the master key type. Valid values are des3-cbc-sha1, des-cbc-crc, des-cbc-md5, des-cbc-raw, arcfour-hmac-md5, arcfour-hmac-md5-exp, aes128-cts-hmac-sha1-96, and aes256-cts-hmac-sha1-96. -m Enter the master key manually. -M <i>mkeyname</i> Specify the master key name. -P <i>password</i> Use the specified <i>password</i> instead of the stash file. -r <i>realm</i> Use <i>realm</i> as the default database realm. -sf <i>stashfile_name</i> Specifies the stash file of the master database password. -x <i>db_args</i> Pass database-specific arguments to kadmin. Supported arguments are for LDAP and the Berkeley-db2 plug-in. These arguments are: - binddn=<i>binddn</i> LDAP simple bind DN for authorization on the directory server. Overrides the ldap_kadmind_dn parameter setting in krb5.conf(4). - bindpwd=<i>bindpwd</i> Bind password. - dbname=<i>name</i> For the Berkeley-db2 plug-in, specifies a name for the Kerberos database.

nconns=*num*

Maximum number of server connections.

port=*num*

Directory server connection port.

操作数

The following operands are supported:

cmd

Specifies whether to create, destroy, dump, or load the database, or to create a stash file.

You can specify the following commands:

create -s

Creates the database specified by the -d option. You will be prompted for the database master password. If you specify -s, a stash file is created as specified by the -f option. If you did not specify -f, the default stash file name is /var/krb5/.k5.realm. If you use the -f, -k, or -M options when you create a database, then you must use the same options when modifying or destroying the database.

destroy

Destroys the database specified by the -d option.

stash

Creates a stash file. If -f was not specified, the default stash file name is /var/krb5/.k5.realm. You will be prompted for the master database password. This command is useful when you want to generate the stash file from the password.

dump [-old] [-b6] [-b7] [-ov] [-r13] [-verbose] [-mkey_convert] [-new_mkey_file *mkey_file*] [-rev] [-recurse] [*filename* [*principals...*]]

Dumps the current Kerberos and KADM5 database into an ASCII file. By default, the database is dumped in current format, “kdb5_util load_dumpversion 6”. If *filename* is not specified or is the string “-”, the dump is sent to standard output. Options are as follows:

-old

Causes the dump to be in the Kerberos 5 Beta 5 and earlier dump format (“kdb5_edit load_dump version 2.0”).

-b6

Causes the dump to be in the Kerberos 5 Beta 6 format (“kdb5_edit load_dump version 3.0”).

-b7

Causes the dump to be in the Kerberos 5 Beta 7 format (“kdb5_util load_dump version 4”). This was the dump format produced on releases prior to 1.2.2.

-ov

Causes the dump to be in ovsec_adm_export format.

-r13

Causes the dump to be in the Kerberos 5 1.3 format (“kdb5_util load_dump version 5”). This was the dump format produced on releases prior to 1.8.

-verbose

Causes the name of each principal and policy to be displayed as it is dumped.

-mkey_convert

Prompts for a new master key. This new master key will be used to re-encrypt the key data in the dumpfile. The key data in the database will not be changed.

-new_mkey_file *mkey_file*

The filename of a stash file. The master key in this stash file will be used to re-encrypt the key data in the dumpfile. The key data in the database will not be changed.

-rev

Dumps in reverse order. This might recover principals that do not dump normally, in cases where database corruption has occurred.

-recurse

Causes the dump to walk the database recursively (bt ree only). This might recover principals that do not dump normally, in cases where database corruption has occurred. In cases of such corruption, this option will probably retrieve more principals than will the -rev option.

load [-old] [-b6] [-b7] [-ov] [-r13] [-hash] [-verbose] [-update] *filename dbname* [*admin_dbname*]

Loads a database dump from *filename* into *dbname*. Unless the -old or -b6 option is specified, the format of the dump file is detected automatically and handled appropriately. Unless the -update option is specified, load creates a new database containing only the principals in the dump file, overwriting the contents of any existing database. The -old option requires the database to be in the Kerberos 5 Beta 5 or earlier format (“kdb5_edit load_dump version 2.0”).

-b6

Requires the database to be in the Kerberos 5 Beta 6 format (“kdb5_edit load_dump version 3.0”).

-b7

Requires the database to be in the Kerberos 5 Beta 7 format (“kdb5_util load_dump version 4”).

-ov

Requires the database to be in ovsec_adm_import format. Must be used with the -update option.

-hash

Requires the database to be stored as a hash. If this option is not specified, the database will be stored as a bt ree. This option is not recommended, as databases stored in hash format are known to corrupt data and lose principals.

-r13

Causes the dump to be in the Kerberos 5 1.3 format (“kdb5_util load_dump version 5”). This was the dump format produced on releases prior to 1.8.

-verbose

Causes the name of each principal and policy to be displayed as it is dumped.

-update

Records from the dump file are added to or updated in the existing database.

Otherwise, a new database is created containing only what is in the dump file and the old one is destroyed upon successful completion.

filename

Required argument that specifies a path to a file containing database dump.

dbname

Required argument that overrides the value specified on the command line or overrides the default.

admin_dbname

Optional argument that is derived from *dbname* if not specified.

add_mkey [-e *etype*] [-s]

Adds a new master key to the K/M (master key) principal. Existing master keys will remain. The **-e *etype*** option allows specification of the enctype of the new master key. The **-s** option stashes the new master key in a local stash file which will be created if it does not already exist.

use_mkey *mkeyVNO* [*time*]

Sets the activation time of the master key specified by *mkeyVNO*. Once a master key is active (that is, its activation time has been reached) it will then be used to encrypt principal keys either when the principal keys change, are newly created, or when the **update_princ_encryption** command is run. If the *time* argument is provided, that will be the activation time; otherwise the current time is used by default. The format of the optional time argument is that specified in the Time Formats section of the [kadmin\(1M\)](#) man page.

list_mkeys

List all master keys from most recent to earliest in K/M principal. The output will show the KVNO, enctype and salt for each mkey, similar to **kadmin getprinc** output. An asterisk (*) following an mkey denotes the currently active master key.

purge_mkeys [-f] [-n] [-v]

Delete master keys from the K/M principal that are not used to protect any principals. This command can be used to remove old master keys from a K/M principal once all principal keys are protected by a newer master key.

-f

Does not prompt user.

- n
Does a dry run, shows master keys that would be purged, but does not actually purge any keys.
- v
Verbose output.

`update_princ_encryption [-f] [-n] [-v] [princ-pattern]`
Update all principal records (or only those matching the *princ-pattern* glob pattern) to re-encrypt the key data using the active database master key, if they are encrypted using older versions, and give a count at the end of the number of principals updated. If the -f option is not given, ask for confirmation before starting to make changes. The -v option causes each principal processed (each one matching the pattern) to be listed, and an indication given as to whether it needed updating or not. The -n option causes the actions not to be taken, only the normal or verbose status messages displayed; this implies -f, since no database changes will be performed and thus there is little reason to seek confirmation.

示例

示例 1 Creating File that Contains Information about Two Principals
The following example creates a file named `slavedata` that contains the information about two principals, `jdb@ACME.COM` and `pak@ACME.COM`.

```
# /usr/krb5/bin/kdb5_util dump -verbose slavedata
jdb@ACME.COM pak@ACME.COM
```

文件

- `/var/krb5/principal`
Kerberos principal database.
- `/var/krb5/principal.kadm5`
Kerberos administrative database. Contains policy information.
- `/var/krb5/principal.kadm5.lock`
Lock file for the Kerberos administrative database. This file works backwards from most other lock files (that is, `kadmin` exits with an error if this file does not exist).
- `/var/krb5/principal.ulong`
The update log file for incremental propagation.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/security/kerberos-5
Interface Stability	Committed

另请参见

[kpasswd\(1\)](#), [gkadmin\(1M\)](#), [kadmin\(1M\)](#), [kadmind\(1M\)](#), [kadmin.local\(1M\)](#), [kdb5_ldap_util\(1M\)](#), [kproplog\(1M\)](#), [kadm5.acl\(4\)](#), [kdc.conf\(4\)](#), [attributes\(5\)](#), [kerberos\(5\)](#)

附注 The global `-f` is made obsolete with the `-sf` argument for specifying a non-default stash file location. The global `-f` argument might be removed in a future release of the Solaris operating system. Use caution in specifying `-f` as it has different semantics in subcommands as distinguished from its use as a global argument.

引用名	kdcmgr – set up a Kerberos Key Distribution Center (KDC)
用法概要	<pre>/usr/sbin/kdcmgr [-a <i>admprincipal</i>] [-e <i>enctype</i>] [-h] [-p <i>pwfile</i>] [-r <i>realm</i>] <i>subcommand</i></pre>
描述	Use the kdcmgr utility to do the following: ■ Configure a master Key Distribution Center (KDC) server. ■ Configure a slave KDC. This assumes that a master KDC has already been configured. The default propagation method configured is incremental propagation. See kpropd(1M). ■ Specify a list of slave KDCs to configure service principals and create access control list for those slaves on the master KDC. If you specify no options, kdcmgr prompts you for required information, including a password to generate the master key and a password for the administrative principal. When you specify sufficient options, you are still prompted for these passwords, unless you specified the -p <i>pwfile</i> option. The kdcmgr utility must be run as superuser or by someone who has the Primary Administrator role. The command must be run on the server from which it is invoked. Note that kdcmgr requires the user to enter sensitive information, such as the password used to generate the database's master key and the password for the administrative principal. Great care must be taken to ensure that the connection to the server is secured over the network, by using a protocol such as ssh(1). You must also exercise great care when selecting the administrative and master key passwords. They should be derived from non-dictionary words and a long string of characters consisting of all of the following character classes: ■ special characters (for example, !@#\$\$%^&*) ■ numerals (0-9) ■ uppercase letters ■ lowercase letters
选项	The following options are supported: -a <i>admprincipal</i> When creating a master KDC, specifies the administrative principal, <i>admprincipal</i>, that will be created. When creating a slave KDC, <i>admprincipal</i> is used to authenticate as the administrative principal. If you omit -a, the suggested default administrative principal name is the output of logname(1) appended by /admin.

-e *enctype*

Specifies the encryption type to be used when creating the key for the master key, which is used to encrypt all principal keys in the database. The set of valid encryption types used here are described in [krb5.conf\(4\)](#) under the `permitted_enctypes` option. Note that the encryption type specified here must be supported on all KDCs or else they will not be able to decrypt any of the principal keys. Solaris 9 and earlier releases support only the `des-cbc-crc` encryption type for the master key. Therefore, if any of the master or slave KDCs are of these older releases, then `-e des-cbc-crc` would need to be specified on all KDCs configured with `kdcmgr`.

The default encryption type is `aes256-cts-hmac-sha1-96`.

-h

Displays usage information for `kdcmgr`.

-p *pwfile*

Provides the location of the password file that contains the password used to create the administrative principal and/or master key.

Warning: This option should be used with great care. Make sure that this *pwfile* is accessible only by a privileged user and on a local file system. Once the KDC has been configured, you should remove *pwfile*.

-r *realm*

Set the default realm for this server.

If the `-r` option is not specified, `kdcmgr` attempts to obtain the machine's local domain name by submitting the canonical form of the machine's host name to DNS and using the return value to derive the domain name. If successful, the domain name is converted to uppercase and proposed as the default realm name.

Subcommands

The following subcommands are supported:

`create [ master ]`

`create [ -m masterkdc ] slave`

Creates a KDC. If no option is specified, an attempt to create a master KDC is made.

`create [ master ]`

Create a master KDC. Upon successful configuration the [krb5kdc\(1M\)](#) and [kadmind\(1M\)](#) are enabled on the machine.

`create [ -m masterkdc ] slave`

Configures a slave KDC. After configuration, the [krb5kdc\(1M\)](#) and [kpropd\(1M\)](#) services are enabled on the machine.

masterkdc specifies the master KDC to authenticate and with which to perform administrative tasks. If the `-m` option is not specified, you are prompted for a master KDC host name.

destroy

Remove all Kerberos configuration and database files associated with the KDC server. A confirmation is required before these files are deleted.

status

Determines the role of the KDC, master or slave, and outputs this and the state of such associated processes as:

- [krb5kdc\(1M\)](#)
- [kadmind\(1M\)](#)
- [kpropd\(1M\)](#)

The subcommand also displays information on incremental propagation if the configuration has this feature enabled, as well as any issues with dependent files.

示例**示例 1** Setting up a Master KDC

The following command configures a master KDC with the administrative principal `user1/admin` and with the realm name `EXAMPLE.COM`:

```
$ kdcmgr -a user1/admin -r EXAMPLE.COM create
```

Note that a password will be required to assign to the newly created `user1/admin` principal. The password for the master key will also need to be provided.

示例 2 Setting up a Slave KDC

The following command configures a slave KDC, authenticates with the administrative principal `user1/admin`, specifies `kdc1` as the master, and uses the `EXAMPLE.COM` realm name:

```
$ kdcmgr -a user1/admin -r EXAMPLE.COM create -m kdc1 slave
```

Note that you must enter the correct password for `user1/admin` and that the master KDC must already have been created before entering this command. The correct password for the master key is also required.

文件

`/etc/krb5/krb5.conf`

Main Kerberos configuration file.

`/etc/krb5/kdc.conf`

KDC configuration, used by both master and slave servers.

`/etc/krb5/krb5.keytab`

Default location of the local host's service keys.

`/etc/krb5/kadm5.acl`

Kerberos administrative access control list (ACL).

`/etc/krb5/kadm5.keytab`

Service keys specific to [kadmind\(1M\)](#).

- `/var/krb5/principal`
Kerberos principal database.
- `/var/krb5/principal.kadm5`
Kerberos policy database.
- `/etc/krb5/kpropd.acl`
Used by slaves to indicate from which server to receive updates.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/security/kerberos-5
Interface Stability	See below

The command line interface (CLI) is Uncommitted. The CLI output is Not an Interface.

另请参见

[logname\(1\)](#), [ssh\(1\)](#), [kadmin\(1M\)](#), [kadmind\(1M\)](#), [kdb5_util\(1M\)](#), [kdb5_ldap_util\(1M\)](#), [kpropd\(1M\)](#), [krb5kdc\(1M\)](#), [ping\(1M\)](#), [svcadm\(1M\)](#), [kdc.conf\(4\)](#), [krb5.conf\(4\)](#), [attributes\(5\)](#)

引用名	kernel – UNIX system executable file containing basic operating system services
用法概要	kernel-name [-asrvx] [-m <i>smf_options</i>] [-i <i>altinit</i>]
描述	The operating system image, or kernel, is the collection of software comprising the image files (unix and genunix) and the modules loaded at any instant in time. The system will not function without a kernel to control it. The kernel is loaded by the boot(1M) command in a machine-specific way. The kernel may be loaded from disk, CD-ROM, or DVD (diskfull boot) or over the network (diskless boot). In either case, the directories under /platform and /kernel must be readable and must contain executable code which is able to perform the required kernel service. If the -a flag is given, the user is able to supply different pathnames for the default locations of the kernel and modules. See boot(1M) for more information on loading a specific kernel. The moddir variable contains a list of module directories separated by whitespace. moddir can be set in the /etc/system file. The minimal default is: <pre>/platform/platform-name/kernel /kernel /usr/kernel</pre> This default can be supplemented by a specific platform. It is common for many SPARC systems to override the default path with: <pre>/platform/platform-name/kernel:/platform/hardware-class-name\ /kernel:/kernel:/usr/kernel</pre> where <i>platform-name</i> can be found using the -i option of uname(1), and <i>hardware-class-name</i> can be found using the -m option of uname(1). The kernel configuration can be controlled using the /etc/system file (see system(4)). genunix is the platform-independent component of the base kernel.
选项	The following options are supported: -a Asks the user for configuration information, such as where to find the system file, where to mount root, and even override the name of the kernel itself. Default responses will be contained in square brackets ([]), and the user may simply enter RETURN to use the default response (note that RETURN is labeled ENTER on some keyboards). To help repair a damaged /etc/system file, enter /dev/null at the prompt that asks for the pathname of the system configuration file. See system(4). -i <i>altinit</i> Select an alternative executable to be the primordial process. <i>altinit</i> must be a valid path to an executable. The default primordial process is init(1M). -m <i>smf_options</i> The <i>smf_options</i> include two categories of options to control booting behavior of the service management facility: recovery options and messages options.

Message options determine the type and amount of messages that [smf\(5\)](#) displays during boot. Service options determine the services which are used to boot the system.

Recovery options

debug

Prints standard per-service output and all `svc.startd` messages to log.

milestone=[milestone]

Boot with some SMF services temporarily disabled, as indicated by *milestone*. *milestone* can be “none”, “single-user”, “multi-user”, “multi-user-server”, or “all”. See the `milestone` subcommand of [svcadm\(1M\)](#).

Messages options

quiet

Prints standard per-service output and error messages requiring administrative intervention.

verbose

Prints standard per-service output with more informational messages.

-r

Reconfiguration boot. The system will probe all attached hardware devices and configure the logical namespace in `/dev`. See [add_drv\(1M\)](#) and [rem_drv\(1M\)](#) for additional information about maintaining device drivers.

-s

Boots only to init level 's'. See [init\(1M\)](#).

-v

Boots with verbose messages enabled. If this flag is not given, the messages are still printed, but the output is directed to the system logfile. See [syslogd\(1M\)](#).

-x

Does not boot in clustered mode. This option only has an effect when a version of Sun Cluster software that supports this option has been installed.

示例 文件

See [boot\(1M\)](#) for examples and instructions on how to boot.

`/kernel`

Contains kernel components common to all platforms within a particular instruction set that are needed for booting the system. of the core image file.

`/platform/platform-name/kernel`

The platform-specific kernel components.

`/platform/hardware-class-name/kernel`

The kernel components specific to this hardware class.

`/usr/kernel`
Contains kernel components common to all platforms within a particular instruction set.

The directories in this section can potentially contain the following subdirectories:

- `drv`
Loadable device drivers
- `exec`
The modules that execute programs stored in various file formats.
- `fs`
File system modules
- `misc`
Miscellaneous system-related modules
- `sched`
Operating system schedulers
- `strmod`
System V STREAMS loadable modules
- `sys`
Loadable system calls

SPARC

- `cpu`
Processor specific modules
- `tod`
Time-Of-Day hardware interface modules

As only 64-bit SPARC platforms are supported, all SPARC executable modules are contained within `sparcv9` directories in the directories listed above.

x86

- `mach`
x86 hardware support

Modules comprising the 32-bit x86 kernel are contained in the above directories, with the 64-bit x86 kernel components contained within `amd64` subdirectories.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/library/processor, system/library/processorx

另请参见

[kmdb\(1\)](#), [uname\(1\)](#), [isainfo\(1\)](#), [add_drv\(1M\)](#), [boot\(1M\)](#), [init\(1M\)](#), [rem_drv\(1M\)](#), [savecore\(1M\)](#), [svc.startd\(1M\)](#), [svcadm\(1M\)](#), [syslogd\(1M\)](#), [system\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [devfs\(7FS\)](#)

SPARC Only

[monitor\(1M\)](#)

诊断

The kernel gives various warnings and error messages. If the kernel detects an unrecoverable fault, it will panic or halt.

附注

Reconfiguration boot will, by design, not remove /dev entries for some classes of devices that have been physically removed from the system.

引用名	keyserv – server for storing private encryption keys																
用法概要	keyserv [-c] [-d -e] [-D] [-n] [-s <i>sizespec</i>]																
描述	keyserv is a daemon that is used for storing the private encryption keys of each user logged into the system. These encryption keys are used for accessing secure network services such as secure NFS. Normally, root's key is read from the file <code>/etc/.rootkey</code> when the daemon is started. This is useful during power-fail reboots when no one is around to type a password. keyserv does not start up if the system does not have a secure rpc domain configured. Set up the domain name by using the <code>/usr/bin/domainname</code> command. Usually the <code>svc:/system/identity:domain</code> service reads the domain from <code>/etc/defaultdomain</code>. Invoking the <code>domainname</code> command without arguments tells you if you have a domain set up. The <code>/etc/default/keyserv</code> file contains the following default parameter settings. See 文件. <table><tr><td>ENABLE_NOBODY_KEYS</td><td>Specifies whether default keys for nobody are used.</td></tr><tr><td colspan="2">ENABLE_NOBODY_KEYS=NO is equivalent to the <code>-d</code> command-line option. The default value for ENABLE_NOBODY_KEYS is YES.</td></tr></table>	ENABLE_NOBODY_KEYS	Specifies whether default keys for nobody are used.	ENABLE_NOBODY_KEYS=NO is equivalent to the <code>-d</code> command-line option. The default value for ENABLE_NOBODY_KEYS is YES.													
ENABLE_NOBODY_KEYS	Specifies whether default keys for nobody are used.																
ENABLE_NOBODY_KEYS=NO is equivalent to the <code>-d</code> command-line option. The default value for ENABLE_NOBODY_KEYS is YES.																	
选项	The following options are supported: <table><tr><td>-c</td><td>Do not use disk caches. This option overrides any <code>-s</code> option.</td></tr><tr><td>-D</td><td>Run in debugging mode and log all requests to keyserv.</td></tr><tr><td>-d</td><td>Disable the use of default keys for nobody. See 文件.</td></tr><tr><td>-e</td><td>Enable the use of default keys for nobody. This is the default behavior. See 文件.</td></tr><tr><td>-n</td><td>Root's secret key is not read from <code>/etc/.rootkey</code>. Instead, keyserv prompts the user for the password to decrypt root's key stored in the <code>publickey</code> database and then stores the decrypted key in <code>/etc/.rootkey</code> for future use. This option is useful if the <code>/etc/.rootkey</code> file ever gets out of date or corrupted.</td></tr><tr><td>-s <i>sizespec</i></td><td>Specify the size of the extended Diffie-Hellman common key disk caches. The <i>sizespec</i> can be one of the following forms:</td></tr><tr><td><i>mechtype</i>=<i>size</i></td><td><i>size</i> is an integer specifying the maximum number of entries in the cache, or an integer immediately followed by the letter <i>M</i>, denoting the maximum size in MB.</td></tr><tr><td><i>size</i></td><td>This form of <i>sizespec</i> applies to all caches.</td></tr></table> Note that the <code>des</code> mechanism, <code>AUTH_DES</code>, does not use a disk cache.	-c	Do not use disk caches. This option overrides any <code>-s</code> option.	-D	Run in debugging mode and log all requests to keyserv.	-d	Disable the use of default keys for nobody. See 文件 .	-e	Enable the use of default keys for nobody. This is the default behavior. See 文件 .	-n	Root's secret key is not read from <code>/etc/.rootkey</code> . Instead, keyserv prompts the user for the password to decrypt root's key stored in the <code>publickey</code> database and then stores the decrypted key in <code>/etc/.rootkey</code> for future use. This option is useful if the <code>/etc/.rootkey</code> file ever gets out of date or corrupted.	-s <i>sizespec</i>	Specify the size of the extended Diffie-Hellman common key disk caches. The <i>sizespec</i> can be one of the following forms:	<i>mechtype</i> = <i>size</i>	<i>size</i> is an integer specifying the maximum number of entries in the cache, or an integer immediately followed by the letter <i>M</i> , denoting the maximum size in MB.	<i>size</i>	This form of <i>sizespec</i> applies to all caches.
-c	Do not use disk caches. This option overrides any <code>-s</code> option.																
-D	Run in debugging mode and log all requests to keyserv.																
-d	Disable the use of default keys for nobody. See 文件 .																
-e	Enable the use of default keys for nobody. This is the default behavior. See 文件 .																
-n	Root's secret key is not read from <code>/etc/.rootkey</code> . Instead, keyserv prompts the user for the password to decrypt root's key stored in the <code>publickey</code> database and then stores the decrypted key in <code>/etc/.rootkey</code> for future use. This option is useful if the <code>/etc/.rootkey</code> file ever gets out of date or corrupted.																
-s <i>sizespec</i>	Specify the size of the extended Diffie-Hellman common key disk caches. The <i>sizespec</i> can be one of the following forms:																
<i>mechtype</i> = <i>size</i>	<i>size</i> is an integer specifying the maximum number of entries in the cache, or an integer immediately followed by the letter <i>M</i> , denoting the maximum size in MB.																
<i>size</i>	This form of <i>sizespec</i> applies to all caches.																

文件 `/etc/.rootkey`
`/etc/default/keyserv` Contains default settings. You can use command-line options to override these settings.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [keylogin\(1\)](#), [svcs\(1\)](#), [keylogout\(1\)](#), [svcadm\(1M\)](#), [publickey\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The keyserv service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/network/rpc/keyserv:default`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名 killall – kill all active processes

用法概要 /usr/sbin/killall [*signal*]

描述 killall is used by [shutdown\(1M\)](#) to kill all active processes not directly related to the shutdown procedure.

killall terminates all processes with open files so that the mounted file systems will be unbusied and can be unmounted.

killall sends *signal* (see [kill\(1\)](#)) to the active processes. If no *signal* is specified, a default of 15 is used.

The killall command can be run only by the super-user.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [kill\(1\)](#), [ps\(1\)](#), [fuser\(1M\)](#), [shutdown\(1M\)](#), [signal\(3C\)](#), [attributes\(5\)](#)

引用名 kmscfg – 配置 PKCS#11 KMS 提供者

用法概要

```
kmscfg
kmscfg -p[rofile] Profile_Name
kmscfg -a[gent] Agent_ID
kmscfg -i[paddr] Agent_Address
kmscfg -t[imeout] Transaction_Timeout
kmscfg -f[ailover] Failover_Limit
kmscfg -d[iscovery] Discovery_Freq
```

描述

kmscfg 命令用于初始化 PKCS#11 KMS 提供者 (pkcs11_kms)，以便与 Solaris 加密框架一起使用。要使 KMS 提供者与 Oracle Key Manager (OKM) 进行通信，必须提供一些配置信息。该配置数据包含要使用的配置文件的名称、OKM 代理的名称、OKM 设备 (KMA) 的 IP 地址及其他一些参数（请参见SYNOPSIS部分）。

缺省情况下，**kmscfg** 将配置信息存储在 `/var/user/kms/$USERNAME` 中。如果该目录不存在，则将创建该目录。如果已检测到该配置，则将为用户提供覆盖现有数据的选项。可使用 `KMSTOKEN_DIR` 环境变量覆盖缺省位置，必须在调用 **kmscfg** 之前设置该变量。

在运行 **kmscfg** 之前，OKM 管理员必须已对设备本身执行了所需的初始化和配置步骤才能设置 PKCS11 KMS 使用者将使用的各个配置文件和代理。有关配置这些配置文件的说明可以在 Oracle Web 站点 (<http://docs.oracle.com>) 上的《Oracle Key Manager Administration Guide》中找到。

一旦配置了 KMA，管理员必须提供必需的标识信息（配置文件名称、代理 ID、IP 地址）才能在 Oracle Solaris 客户机上运行 **kmscfg** 和初始化提供者。

选项

支持下面列出的选项。请注意，如果未在命令行上指定配置文件、代理 ID 或 KMA 地址，**kmscfg** 会提示您提供这些项。

-a *Agent_ID*

在 OKM 上配置的用户代理 ID，用于正在配置的 KMS 令牌。配置文件和代理 ID 通常是相同的，例如，MyAgent。

-d *Discovery_Freq*

客户机将尝试搜索 OKM 群集中其他 KMA 的可用性的频率（秒）。如果未指定，则 *Discovery_Freq* 缺省为 10。

-f *Failover_Limit*

在客户机放弃之前，与 KMA 的通信可以失败的次数。如果未指定，则 *Failover_Limit* 缺省为 3。

- i *Agent_Addr*
KMA的地址。此地址可以是 IPv4 地址 (*xxx.xxx.xxx.xxx*) 或 IPv6 地址。只要客户机上配置的名称服务可以解析全限定主机名，则还可以使用全限定主机名。如果使用的是 OKM 群集，则可以指定群集中任何成员的地址。
- p *Profile_Name*
要用于正在配置的 KMS 令牌的配置文件的名称。通常，配置文件名称和代理 ID 是相同的（尽管并不要求它们相同）。
- t *Transaction_Timeout*
各个 KMS 命令的超时周期（以秒为单位）。如果未指定，则该值缺省为 10。

退出状态 在完成所请求的操作之后，`kmscfg` 以下列状态值之一退出。

- 0
成功终止。
- 1
失败。无法完成所请求的操作。

文件 `/var/kms/$USERNAME`
缺省的 KMS 令牌配置目录。

`${KMSTOKEN_DIR}`
备用的 KMS 令牌配置目录。

属性 有关以下属性的说明，请参见 `attributes(5)`：

属性类型	属性值
可用性	<code>/system/library/security/crypto/pkcs11_kms</code>
接口稳定性	Volatile（可变）

另请参见 `pktool(1)`、`attributes(5)`、`pkcs11_kms(5)`

《Oracle Key Manager Administration Guide》(<http://docs.oracle.com>)

附注 PKCS#11 客户机要求在 OKM 上安装 2.4 版 Oracle Key Manager 软件。

如果 PKCS#11 客户机要对多个系统使用同一个代理 ID，则创建该代理时时不应设置 "One Time Passphrase"（一次性口令短语）标志。在一些成员运行的 OKM 软件版本低于 2.4 的 OKM 群集中，该选项将不可用。有关创建代理的帮助信息，请参考《Oracle Key Manager Administration Guide》。

OKM 代理在用于在 PKCS#11 客户机中创建密钥之前，必须分配有缺省密钥组。如果没有为该代理分配缺省密钥组，操作将失败并显示 `CKR_PIN_INCORRECT` 错误。有关为代理分配密钥组的帮助信息，请参考《Oracle Key Manager Administration Guide》。

引用名	kprop – Kerberos database propagation program										
用法概要	<pre>/usr/lib/krb5/kprop [-d] [-f file] [-p port-number] [-r realm] [-s keytab] [host]</pre>										
描述	kprop is a command-line utility used for propagating a Kerberos database from a master KDC to a slave KDC. This command must be run on the master KDC. See the Solaris System Administration Guide, Vol. 6 on how to set up periodic propagation between the master KDC and slave KDCs. To propagate a Kerberos database, the following conditions must be met: ■ The slave KDCs must have an <code>/etc/krb5/kpropd.acf</code> file that contains the principals for the master KDC and all the slave KDCs. ■ A keytab containing a host principal entry must exist on each slave KDC. ■ The database to be propagated must be dumped to a file using kdb5_util(1M).										
选项	The following options are supported: <table><tr><td><code>-d</code></td><td>Enable debug mode. Default is debug mode disabled.</td></tr><tr><td><code>-f file</code></td><td>File to be sent to the slave KDC. Default is the <code>/var/krb5/slave_datatrans</code> file.</td></tr><tr><td><code>-p port-number</code></td><td>Propagate <i>port-number</i>. Default is port 754.</td></tr><tr><td><code>-r realm</code></td><td>Realm where propagation will occur. Default <i>realm</i> is the local realm.</td></tr><tr><td><code>-s keytab</code></td><td>Location of the keytab. Default location is <code>/etc/krb5/krb5.keytab</code>.</td></tr></table>	<code>-d</code>	Enable debug mode. Default is debug mode disabled.	<code>-f file</code>	File to be sent to the slave KDC. Default is the <code>/var/krb5/slave_datatrans</code> file.	<code>-p port-number</code>	Propagate <i>port-number</i> . Default is port 754.	<code>-r realm</code>	Realm where propagation will occur. Default <i>realm</i> is the local realm.	<code>-s keytab</code>	Location of the keytab. Default location is <code>/etc/krb5/krb5.keytab</code> .
<code>-d</code>	Enable debug mode. Default is debug mode disabled.										
<code>-f file</code>	File to be sent to the slave KDC. Default is the <code>/var/krb5/slave_datatrans</code> file.										
<code>-p port-number</code>	Propagate <i>port-number</i> . Default is port 754.										
<code>-r realm</code>	Realm where propagation will occur. Default <i>realm</i> is the local realm.										
<code>-s keytab</code>	Location of the keytab. Default location is <code>/etc/krb5/krb5.keytab</code> .										
操作数	The following operands are supported: <table><tr><td><i>host</i></td><td>Name of the slave KDC.</td></tr></table>	<i>host</i>	Name of the slave KDC.								
<i>host</i>	Name of the slave KDC.										
示例	示例 1 Propagating the Kerberos Database The following example propagates the Kerberos database from the <code>/tmp/slave_data</code> file to the slave KDC <code>london</code>. The machine <code>london</code> must have a host principal keytab entry and the <code>kpropd.acf</code> file must contain an entry for the all the KDCs. <pre># kprop -f /tmp/slave_data london</pre>										
文件	<table><tr><td><code>/etc/krb5/kpropd.acf</code></td><td>List of principals of all the KDCs; resides on each slave KDC.</td></tr><tr><td><code>/etc/krb5/krb5.keytab</code></td><td>Keytab for Kerberos clients.</td></tr><tr><td><code>/var/krb5/slave_datatrans</code></td><td>Kerberos database propagated to the KDC slaves.</td></tr></table>	<code>/etc/krb5/kpropd.acf</code>	List of principals of all the KDCs; resides on each slave KDC.	<code>/etc/krb5/krb5.keytab</code>	Keytab for Kerberos clients.	<code>/var/krb5/slave_datatrans</code>	Kerberos database propagated to the KDC slaves.				
<code>/etc/krb5/kpropd.acf</code>	List of principals of all the KDCs; resides on each slave KDC.										
<code>/etc/krb5/krb5.keytab</code>	Keytab for Kerberos clients.										
<code>/var/krb5/slave_datatrans</code>	Kerberos database propagated to the KDC slaves.										

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/security/kerberos-5

另请参见

[kpasswd\(1\)](#), [svcs\(1\)](#), [gkadmin\(1M\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [kadmind\(1M\)](#), [kadmin.local\(1M\)](#), [kdb5_util\(1M\)](#), [svcadm\(1M\)](#), [kadm5.acl\(4\)](#), [kdc.conf\(4\)](#), [attributes\(5\)](#), [kerberos\(5\)](#), [smf\(5\)](#)

《Oracle Solaris 管理：安全服务》

引用名	kpropd – Kerberos propagation daemon for slave KDCs												
用法概要	<pre>/usr/lib/krb5/kpropd [-d] [-f temp_dbfile] [-F dbfile] [-p kdb_util] [-P port_number] [-r realm] [-s srv_tabfile] [-S] [-a acl_file]</pre>												
描述	The kpropd command runs on the slave KDC server. It listens for update requests made by kprop(1M) from the master KDC and periodically requests incremental updates from the master KDC. When the slave receives a kprop request from the master, kpropd copies principal data to a temporary text file. Next, kpropd invokes kdb5_util(1M) (unless a different database utility is selected) to load the text file in database format. When the slave periodically requests incremental updates, kpropd update its <code>principal.ulong</code> file with any updates from the master. kproplog(1M) can be used to view a summary of the update entry log on the slave KDC. kpropd is not configured for incremental database propagation by default. These settings can be changed in the kdc.conf(4) file: <table><tr><td><code>sunw_dbprop_enable = [true false]</code></td><td>Enables or disables incremental database propagation. Default is <code>false</code>.</td></tr><tr><td><code>sunw_dbprop_slave_poll = N[s, m, h]</code></td><td>Specifies how often the slave KDC polls for any updates that the master might have. Default is <code>2m</code> (two minutes).</td></tr></table> The <code>kprop/<hostname>@<REALM></code> principal must exist in the slave's <code>keytab</code> file to enable the master to authenticate incremental propagation requests from the slave. In this syntax, <code><hostname></code> is the slave KDC's host name and <code><REALM></code> is the realm in which the slave KDC resides.	<code>sunw_dbprop_enable = [true false]</code>	Enables or disables incremental database propagation. Default is <code>false</code> .	<code>sunw_dbprop_slave_poll = N[s, m, h]</code>	Specifies how often the slave KDC polls for any updates that the master might have. Default is <code>2m</code> (two minutes).								
<code>sunw_dbprop_enable = [true false]</code>	Enables or disables incremental database propagation. Default is <code>false</code> .												
<code>sunw_dbprop_slave_poll = N[s, m, h]</code>	Specifies how often the slave KDC polls for any updates that the master might have. Default is <code>2m</code> (two minutes).												
选项	The following options are supported: <table><tr><td><code>-d</code></td><td>Enable debug mode. Default is debug mode disabled.</td></tr><tr><td><code>-f temp_dbfile</code></td><td>The location of the slave's temporary principal database file. Default is <code>/var/krb5/from_master</code>.</td></tr><tr><td><code>-F dbfile</code></td><td>The location of the slave's principal database file. Default is <code>/var/krb5/principal</code>.</td></tr><tr><td><code>-p kdb_util</code></td><td>The location of the Kerberos database utility used for loading principal databases. Default is <code>/usr/sbin/kdb5_util</code>.</td></tr><tr><td><code>-P port_number</code></td><td>Specifies the port number on which kpropd will listen. Default is 754 (service name: <code>krb5_prop</code>).</td></tr><tr><td><code>-r realm</code></td><td>Specifies from which Kerberos realm kpropd will receive information. Default is specified in <code>/etc/krb5/krb5.conf</code>.</td></tr></table>	<code>-d</code>	Enable debug mode. Default is debug mode disabled.	<code>-f temp_dbfile</code>	The location of the slave's temporary principal database file. Default is <code>/var/krb5/from_master</code> .	<code>-F dbfile</code>	The location of the slave's principal database file. Default is <code>/var/krb5/principal</code> .	<code>-p kdb_util</code>	The location of the Kerberos database utility used for loading principal databases. Default is <code>/usr/sbin/kdb5_util</code> .	<code>-P port_number</code>	Specifies the port number on which kpropd will listen. Default is 754 (service name: <code>krb5_prop</code>).	<code>-r realm</code>	Specifies from which Kerberos realm kpropd will receive information. Default is specified in <code>/etc/krb5/krb5.conf</code> .
<code>-d</code>	Enable debug mode. Default is debug mode disabled.												
<code>-f temp_dbfile</code>	The location of the slave's temporary principal database file. Default is <code>/var/krb5/from_master</code> .												
<code>-F dbfile</code>	The location of the slave's principal database file. Default is <code>/var/krb5/principal</code> .												
<code>-p kdb_util</code>	The location of the Kerberos database utility used for loading principal databases. Default is <code>/usr/sbin/kdb5_util</code> .												
<code>-P port_number</code>	Specifies the port number on which kpropd will listen. Default is 754 (service name: <code>krb5_prop</code>).												
<code>-r realm</code>	Specifies from which Kerberos realm kpropd will receive information. Default is specified in <code>/etc/krb5/krb5.conf</code> .												

- s *srv_tabfile*

The location of the service table file used to authenticate the kpropd daemon.
- S

Run the daemon in standalone mode, instead of having inetd listen for requests. Default is non-standalone mode.
- a *acl_file*

The location of the kpropd's access control list to verify if this server can run the kpropd daemon. The file contains a list of principal name(s) that will be receiving updates. Default is /etc/krb5/kpropd.acl.
- 文件

/var/krb5/principal

Kerberos principal database.

/var/krb5/principal.u^log

The update log file.

/etc/krb5/kdc.conf

KDC configuration information.

/etc/krb5/kpropd.acl

List of principals of all the KDCs; resides on each slave KDC.

/var/krb5/from_master

Temporary file used by kpropd before loading this to the principal database.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/security/kerberos-5
Interface Stability	Committed

另请参见

[kdb5_util\(1M\)](#), [kprop\(1M\)](#), [kproplog\(1M\)](#), [kdc.conf\(4\)](#), [krb5.conf\(4\)](#), [attributes\(5\)](#), [kerberos\(5\)](#)

附注

The kprop service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/network/security/krb5_prop:default

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名 kproplog – display the contents of the Kerberos principal update log

用法概要 /usr/sbin/kproplog [-h | -e *num*]

描述 The kproplog displays the contents of the Kerberos principal update log to standard output. This command can be used to keep track of the incremental updates to the principal database, which is enabled by default. The /var/krb5/principal.u_{log} file contains the update log maintained by the [kadmind\(1M\)](#) process on the master KDC server and the [kpropd\(1M\)](#) process on the slave KDC servers. When updates occur, they are logged to this file. Subsequently any KDC slave configured for incremental updates will request the current data from the master KDC and update their principal.u_{log} file with any updates returned.

The kproplog command can only be run on a KDC server by someone with privileges comparable to the superuser. It will display update entries for that server only.

If no options are specified, the summary of the update log is displayed. If invoked on the master, all of the update entries are also displayed. When invoked on a slave KDC server, only a summary of the updates are displayed, which includes the serial number of the last update received and the associated time stamp of the last update.

选项 The following options are supported:

- h Display a summary of the update log. This information includes the database version number, state of the database, the number of updates in the log, the time stamp of the first and last update, and the version number of the first and last update entry.
- e *num* Display the last *num* update entries in the log. This is useful when debugging synchronization between KDC servers.
- v Display individual attributes per update. If more than one -v is specified then very verbose output is displayed. An example of the output generated for one entry:

```
Update Entry
  Update serial # : 4
  Update operation : Add
  Update principal : test@EXAMPLE.COM
  Update size : 424
  Update committed : True
  Update time stamp : Fri Feb 20 23:37:42 2004
  Attributes changed : 6
    Principal
    Key data
    Password last changed
    Modifying principal
    Modification time
    TL data
```

文件 /var/krb5/principal.u`log` The update log file for incremental propagation.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/security/kerberos-5
Interface Stability	Committed

另请参见 [kpasswd\(1\)](#), [gkadmin\(1M\)](#), [kadmin\(1M\)](#), [kadmin\(1M\)](#), [kdb5_util\(1M\)](#), [kprop\(1M\)](#), [kpropd\(1M\)](#), [kadm5.acl\(4\)](#), [kdc.conf\(4\)](#), [attributes\(5\)](#), [kerberos\(5\)](#)

引用名	krb5kdc – KDC daemon
用法概要	<pre>/usr/lib/krb5/krb5kdc [-d <i>dbpath</i>] [-r <i>realm</i>] [-m] [-k <i>masterenctype</i>] [-M <i>masterkeyname</i>] [-p <i>port</i>] [-n] [-x <i>db_args</i>] [-P <i>pid_file</i>]</pre>
描述	krb5kdc is the daemon that runs on the master and slave KDCs to process the Kerberos tickets. For Kerberos to function properly, krb5kdc must be running on at least one KDC that the Kerberos clients can access. Prior to running krb5kdc, you must initialize the Kerberos database using kdb5_util(1M). See the 《Oracle Solaris 管理：安全服务》 for information regarding how to set up KDCs and initialize the Kerberos database.
选项	The following options are supported: -d <i>dbpath</i> Specify the path to the database; default value is <code>/var/krb5</code>. -k <i>masterenctype</i> Specify the encryption type for encrypting the database. The default value is <code>des-cbc-crc</code>. <code>des3-cbc-sha1</code>, <code>arcfour-hmac-md5</code>, <code>arcfour-hmac-md5-exp</code>, <code>aes128-cts-hmac-sha1-96</code>, and <code>aes256-cts-hmac-sha1-96</code> are also valid. -m Specify that the master key for the database is to be entered manually. -M <i>masterkeyname</i> Specify the principal to retrieve the master Key for the database. -n Specify that krb5kdc should not detach from the terminal. -p <i>port</i> Specify the port that will be used by the KDC to listen for incoming requests. -P <i>pid_file</i> Tells the KDC to write its PID (followed by a newline) into <i>pid_file</i> after it starts up. This can be used to identify whether the KDC is still running and to allow <code>init</code> scripts to stop the correct process. -r <i>realm</i> Specify the realm name; default is the local realm name. -x <i>db_args</i> Pass database-specific arguments to <code>kadmin</code>. Supported arguments are for the LDAP plug-in. These arguments are: <code>binddn=binddn</code> Specifies the DN of the object used by the KDC server to bind to the LDAP server. This object should have the rights to read the realm container, principal container and the subtree that is referenced by the realm. Overrides the <code>ldap_kdc_dn</code> parameter setting in krb5.conf(4).

`bindpwd=bindpwd`
Specifies the password for the above-mentioned `binddn`. It is recommended not to use this option. Instead, the password can be stashed using the `stashrvpw` command of [kdb5_util\(1M\)](#).

`nconns=num`
Specifies the number of connections to be maintained per LDAP server.

`host=ldapuri`
Specifies, by an LDAP URI, the LDAP server to which to connect.

文件

`/var/krb5/principal.db`
Kerberos principal database.

`/var/krb5/principal.kadm5`
Kerberos administrative database. This file contains policy information.

`/var/krb5/principal.kadm5.lock`
Kerberos administrative database lock file. This file works backwards from most other lock files (that is, `kadmin` will exit with an error if this file does *not* exist).

`/etc/krb5/kdc.conf`
KDC configuration file. This file is read at startup.

`/etc/krb5/kpropd.acl`
File that defines the access control list for propagating the Kerberos database using `kprop`.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/security/kerberos-5

另请参见

[kill\(1\)](#), [kpasswd\(1\)](#), [gkadmin\(1M\)](#), [kadmind\(1M\)](#), [kadmin.local\(1M\)](#), [kdb5_util\(1M\)](#), [kdb5_ldap_util\(1M\)](#), [logadm\(1M\)](#), [krb5.conf\(4\)](#), [attributes\(5\)](#), [krb5envvar\(5\)](#), [kerberos\(5\)](#),

《Oracle Solaris 管理：安全服务》

附注

The following signal has the specified effect when sent to the server process using the [kill\(1\)](#) command:

SIGHUP
`krb5kdc` closes and re-opens log files that it directly opens. This can be useful for external log-rotation utilities such as [logadm\(1M\)](#). If this method is used for log file rotation, set the [krb5.conf\(4\)](#) `kdc_rotate` period relation to `never`.

引用名

ksslcfg – enable and configure SMF instance of Kernel SSL

用法概要

```
ksslcfg create -f pkcs11 -T token_label -C certificate_label
    [-d softtoken_directory]
    -p password_file [-u username]
    [-h ca_certchain_file] [-c ciphersuites]
    [-t ssl_session_cache_timeout]
    [-z ssl_session_cache_size] [-v] -x proxy_port [host] ssl_port

ksslcfg create -f pkcs12 -i cert_and_key_pk12file
    -p password_file [-u username]
    [-c ciphersuites] [-t ssl_session_cache_timeout]
    [-z ssl_session_cache_size] [-v] -x proxy_port [host] ssl_port

ksslcfg create -f pem -i cert_and_key_pemfile
    -p password_file [-u username]
    [-c ciphersuites] [-t ssl_session_cache_timeout]
    [-z ssl_session_cache_size] [-v] -x proxy_port [host] ssl_port

ksslcfg delete [-v] [host] ssl_port

ksslcfg -V

ksslcfg -?
```

描述

ksslcfg manages [smf\(5\)](#) instances for the Kernel SSL proxy module. An SSL-enabled web server can use the services of its Kernel SSL proxy to improve the performance of the HTTPS packets processing. It does so by creating an instance of the Kernel SSL service, specifying the SSL proxy port and parameters, and by listening on the proxy port.

The create subcommand creates an instance and enables the service for the given address and SSL port.

The delete subcommand disables the service for the given address and port, if it is enabled, and deletes the instance from the SMF repository.

ksslcfg can be run as root or by other users assigned to the Network Security profile. See [rbac\(5\)](#) and [user_attr\(4\)](#).

After ksslcfg successfully configures the service in the kernel, the proxy application must be started, or restarted if it is already running.

You must run ksslcfg to configure your Kernel SSL proxy before you start your application.

ksslcfg allows you to specify an *ssl_port* operand, described under OPERANDS, and, with the -x option, a *proxy_port* value. When specified for use with the Kernel SSL proxy, these values cannot also be configured for the Solaris Network Cache and Acceleration (NCA) feature.

The Fault Managed Resource Identifier (FMRI) for the kernel SSL proxy instances is `svc://network/ssl/proxy`. ksslcfg creates an instance of that service unique to the

combination of host and SSL port. Instance FMRI for particular proxy entries can be found with [svcs\(1\)](#) and used for dependencies of other services. The state of the service instance tracks in-kernel configuration only. It does not reflect the presence or state of the application listening on the proxy port.

选项

The following options are supported:

-c *ciphersuites*

Set of ciphers a client is allowed to negotiate in a sorted order. The supported SSL version3 and TLS ciphers are listed below. Note that the names are case-insensitive.

```
rsa_rc4_128_sha
rsa_rc4_128_md5
rsa_aes_256_cbc_sha
rsa_aes_128_cbc_sha
rsa_3des_ede_cbc_sha
rsa_des_cbc_sha
```

-f *key_format*

Uses the certificate/key format specified in *key_format*. The supported options are pkcs11, pkcs12, and pem.

-i *key_and_certificate_file*

When pkcs12 or pem is specified with the -f option, reads a key and a certificate of the web server from *key_and_certificate_file*. This file can also contain any intermediate CA certificates that form the certificate chain to the root CA for the server certificate. These certificates must follow the server certificate in the file and the order must be bottom up: lowest level CA certificate followed by the next higher level CA certificate, and so on.

-C *certificate_label*

PKCS#11 can store multiple certificates in single token. This option enables you to specify a single certificate, identified by *certificate_label*. This label must match the CKA_LABEL on the certificate object in the token specified by -T. This option is to be used only with -f pkcs11.

-d *softtoken_directory*

This option is applicable only with the pkcs11 key format, when the token label is the Sun Software PKCS#11 softtoken. Use this option to override the default location of the PKCS#11 softtoken directory (\$HOME/.sunw). See [pkcs11_softtoken\(5\)](#).

-h *ca_certchain_file*

When pkcs11 is specified with the -f option, reads a set of intermediate CA certificates that form the certificate chain to the root CA for the server certificate (specified with the -C option), from *ca_certchain_file*. The file must be in PEM format.

-p *password_file*

Obtains the password used to encrypt the private key from *password_file*. When using the pkcs11 option (see -f, above), the password is used to authenticate the user to the PKCS#11 token.

-t *ssl_session_cache_timeout*

The timeout value, in seconds, for an SSL session. It corresponds to `SSL3SessionTimeout` of the Sun ONE web server configuration or `SSLSessionCacheTimeout` of `mod_ssl`.

-T *token_label*

When `pkcs11` is specified with `-f`, uses the PKCS#11 token specified in *token_label*. Use `cryptoadm list -v` to display all PKCS#11 tokens available.

-u *username*

The username of the user who owns the password file. If omitted, the system will try to read the password file as root.

-v

Verbose mode.

-V

Displays the version.

-x *proxy_port*

The SSL proxy port. The port number is designated exclusively for clear-text HTTP communication between the web server and the kernel SSL proxy module. No external HTTP packets are delivered to this port.

-z *ssl_session_cache_size*

The maximum number of SSL sessions that can be cached. It corresponds to `SSLCacheEntries` of the Sun ONE web server configuration. When this option is not specified, the default is 5000 entries.

-?

Displays the usage of the command.

操作数**[*host*] [*ssl_port*]**

The address and the port of the web server for which the kernel SSL entry is created. If *host* is omitted, the entry will be used for all requests that arrived at the *ssl_port*, regardless of the destination address. Both a host name and an IP address are acceptable forms for *host*. *ssl_port* is required. Typically, this has a value of 443.

示例**示例 1 Create and Enable a Kernel SSL Instance**

The following command creates and enables a Kernel SSL instance using a certificate and a key in PKCS#11 format.

```
# ksslcfg create -f pkcs11 -T "Sun Software PKCS#11 softtoken" \
-C "Server-Cert" -p /some/directory/password -u websrvd \
-x 8080 www.mysite.com 443

% svcs svc:/network/ssl/proxy
STATE      STIME     FMRI
online     Sep_27    svc:/network/ssl/proxy:kssl-www-mysite-com-443
```

示例 2 Create and Enable a Default Instance for All Addresses

The following command creates and enables a default instance for all addresses from a certificate and key in a pkcs#12 file.

```
# ksslcfg create -x 8888 -f pkcs12 -i /some/directory/keypair.p12 \
  -p /some/directory/password -u webserverd 443
```

示例 3 Create and Enable an Instance with Specific Cipher Suites

The following command creates and enables an instance with specific cipher suites.

```
# ksslcfg create -x 8080 -f pem \
  -i /some/directory/keypair.pem -p /some/directory/password \
  -c "rsa_rc4_128_md5,rsa_rc4_128_sha" \
  209.249.116.195 443
```

示例 4 Disable and Delete an Instance

The following command disables and deletes an instance.

```
# ksslcfg delete www.mysite.com 443
```

示例 5 Establishing Dependency of Proxy Application

The sequence of commands shown below establishes a dependency of a proxy application on a KSSL instance. Note that the proxy application should only be started after the SSL kernel proxy instance has been started.

The following commands establish the dependency of an Apache 2.2 web server. KSSL has been configured to listen on SSL port 443 and a wildcard address.

```
# svccfg -s svc:/network/http:apache22
svc:/network/http:apache22> addpg kssl dependency
svc:/network/http:apache22> setprop kssl/entities = fmri:svc:/network/\
ssl/proxy:kssl-INADDR_ANY-443
svc:/network/http:apache22> setprop kssl/grouping = astring: require_all
svc:/network/http:apache22> setprop kssl/restart_on = astring: refresh
svc:/network/http:apache22> setprop kssl/type = astring: service
svc:/network/http:apache22> end
```

Following these commands, enable the web server:

```
# svcadm enable svc:/network/http:apache22
```

If the web server was already running, restart it:

```
# svcadm refresh svc:/network/http:apache22
# svcadm restart svc:/network/http:apache22
```

退出状态

0
Successful completion.

>0
An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	See below.

Command line options and the utility name are Committed. The command output, the FMRI service name (svc://network/ssl/proxy), and the FMRI instance's name format are Uncommitted

另请参见

[svcprop\(1\)](#), [svcs\(1\)](#), [cryptoadm\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [user_attr\(4\)](#), [attributes\(5\)](#), [kssl\(5\)](#), [pkcs11_softtoken\(5\)](#), [rbac\(5\)](#), [smf\(5\)](#), [nca\(7d\)](#)

Chown, P., Advanced Encryption Standard (AES) Ciphersuites for Transport Layer Security (TLS), RFC 3268, June 2002.

附注

ksslcfg create without an host argument creates an INADDR_ANY smf instance. ksslcfg delete without an host argument deletes only the INADDR_ANY instance. ksslcfg delete needs a host argument to delete any non-INADDR_ANY instance.

On a system with [zones\(5\)](#) installed, the ksslcfg command can be used only in the global zone at this time.

引用名	kstat – display kernel statistics
用法概要	<pre>kstat [-lpq] [-T u d] [-c class] [-m module] [-i instance] [-n name] [-s statistic] [interval [count]] kstat [-lpq] [-T u d] [-c class] [module:instance:name:statistic]... [interval [count]]</pre>
描述	The kstat utility examines the available kernel statistics, or kstats, on the system and reports those statistics which match the criteria specified on the command line. Each matching statistic is printed with its module, instance, and name fields, as well as its actual value. Kernel statistics may be published by various kernel subsystems, such as drivers or loadable modules; each kstat has a module field that denotes its publisher. Since each module might have countable entities (such as multiple disks associated with the sd(7D) driver) for which it wishes to report statistics, the kstat also has an instance field to index the statistics for each entity; kstat instances are numbered starting from zero. Finally, the kstat is given a name unique within its module. Each kstat may be a special kstat type, an array of name-value pairs, or raw data. In the name-value case, each reported value is given a label, which we refer to as the statistic. Known raw and special kstats are given statistic labels for each of their values by kstat; thus, all published values can be referenced as <i>module:instance:name:statistic</i>. When invoked without any module operands or options, kstat will match all defined statistics on the system. Example invocations are provided below. All times are displayed as fractional seconds since system boot.
选项	The tests specified by the following options are logically ANDed, and all matching kstats will be selected. A regular expression containing shell metacharacters must be protected from the shell by enclosing it with the appropriate quotes. The argument for the -c, -i, -m, -n, and -s options may be specified as a shell glob pattern, or a Perl regular expression enclosed in '/' characters. -c <i>class</i> Displays only kstats that match the specified class. <i>class</i> is a kernel-defined string which classifies the “type” of the kstat. -i <i>instance</i> Displays only kstats that match the specified instance. -l Lists matching kstat names without displaying values. -m <i>module</i> Displays only kstats that match the specified module. -n <i>name</i> Displays only kstats that match the specified name. -p Displays output in parseable format. All example output in this document is given in this format. If this option is not specified, kstat produces output in a human-readable, table format.

- q Displays no output, but return appropriate exit status for matches against given criteria.
- s *statistic* Displays only kstats that match the specified statistic.
- T d | u Displays a time stamp before each statistics block, either in [date\(1\)](#) format (d) or as an alphanumeric representation of the value returned by [time\(2\)](#) (u).

操作数

The following operands are supported:

- module:instance:name:statistic* Alternate method of specifying module, instance, name, and statistic as described above. Each of the module, instance, name, or statistic specifiers may be a shell glob pattern or a Perl regular expression enclosed by '/' characters. It is possible to use both specifier types within a single operand. Leaving a specifier empty is equivalent to using the '*' glob pattern for that specifier.
- interval* The number of seconds between reports.
- count* The number of reports to be printed.

示例

In the following examples, all the command lines in a block produce the same output, as shown immediately below. The exact statistics and values will of course vary from machine to machine.

示例1 Using the kstat Command

```
example$ kstat -p -m unix -i 0 -n system_misc -s 'avenrun*'
example$ kstat -p -s 'avenrun*'
example$ kstat -p 'unix:0:system_misc:avenrun*'
example$ kstat -p '::::avenrun*'
example$ kstat -p '::::^avenrun_\d+min$/'

unix:0:system_misc:avenrun_15min      3
unix:0:system_misc:avenrun_1min      4
unix:0:system_misc:avenrun_5min      2
```

示例2 Using the kstat Command

```
example$ kstat -p -m cpu_stat -s 'intr*'
example$ kstat -p cpu_stat::::^intr/

cpu_stat:0:cpu_stat0:intr      29682330
cpu_stat:0:cpu_stat0:intrblk    87
cpu_stat:0:cpu_stat0:intrthread 15054222
cpu_stat:1:cpu_stat1:intr      426073
cpu_stat:1:cpu_stat1:intrblk    51
cpu_stat:1:cpu_stat1:intrthread 289668
```


示例2 Using the kstat Command (续)

```

cpu_stat:2:cpu_stat2:intr      134160
cpu_stat:2:cpu_stat2:intrblk    0
cpu_stat:2:cpu_stat2:intrthread 131
cpu_stat:3:cpu_stat3:intr      196566
cpu_stat:3:cpu_stat3:intrblk    30
cpu_stat:3:cpu_stat3:intrthread 59626

```

示例3 Using the kstat Command

```

example$ kstat -p :::state '::::avenrun*'
example$ kstat -p :::state :::/^avenrun/

```

```

cpu_info:0:cpu_info0:state      on-line
cpu_info:1:cpu_info1:state      on-line
cpu_info:2:cpu_info2:state      on-line
cpu_info:3:cpu_info3:state      on-line
unix:0:system_misc:avenrun_15min      4
unix:0:system_misc:avenrun_1min 10
unix:0:system_misc:avenrun_5min 3

```

示例4 Using the kstat Command

```

example$ kstat -p 'unix:0:system_misc:avenrun*' 1 3
unix:0:system_misc:avenrun_15min      15
unix:0:system_misc:avenrun_1min 11
unix:0:system_misc:avenrun_5min 21

unix:0:system_misc:avenrun_15min      15
unix:0:system_misc:avenrun_1min 11
unix:0:system_misc:avenrun_5min 21

unix:0:system_misc:avenrun_15min      15
unix:0:system_misc:avenrun_1min 11
unix:0:system_misc:avenrun_5min 21

```

示例5 Using the kstat Command

```

example$ kstat -p -T d 'unix:0:system_misc:avenrun*' 5 2
Thu Jul 22 19:39:50 1999
unix:0:system_misc:avenrun_15min      12
unix:0:system_misc:avenrun_1min 0
unix:0:system_misc:avenrun_5min 11

Thu Jul 22 19:39:55 1999
unix:0:system_misc:avenrun_15min      12
unix:0:system_misc:avenrun_1min 0
unix:0:system_misc:avenrun_5min 11

```

```
示例 6 Using the kstat Command

example$ kstat -p -T u 'unix:0:system_misc:avenrun*'
932668656
unix:0:system_misc:avenrun_15min      14
unix:0:system_misc:avenrun_1min      5
unix:0:system_misc:avenrun_5min     18
```

退出状态

- The following exit values are returned:
- 0 One or more statistics were matched.
 - 1 No statistics were matched.
 - 2 Invalid command line options were specified.
 - 3 A fatal error occurred.

文件

/dev/kstat kernel statistics driver

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[date\(1\)](#), [sh\(1\)](#), [time\(2\)](#), [gmatch\(3GEN\)](#), [kstat\(3KSTAT\)](#), [attributes\(5\)](#), [kstat\(7D\)](#), [sd\(7D\)](#), [kstat\(9S\)](#)

附注

If the pattern argument contains glob or Perl RE metacharacters which are also shell metacharacters, it will be necessary to enclose the pattern with appropriate shell quotes.

引用名	ktkt_warnd – Kerberos warning daemon
用法概要	/usr/lib/krb5/ktkt_warnd
描述	ktkt_warnd is a daemon on Kerberos clients that can warn users when their Kerberos tickets are about to expire or renew the tickets before they expire. It is invoked by inetd when a ticket-granting ticket (TGT) is obtained for the first time, such as after using the kinit command. ktkt_warnd can be configured through the /etc/krb5/warn.conf file on the client. In warn.conf, you can specify that you be supplied notice, through syslog or terminal or mail, of ticket expiration or to renew the TGT.
文件	/etc/krb5/warn.conf Kerberos warning configuration file
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

另请参见 [svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [svcadm\(1M\)](#), [warn.conf\(4\)](#), [attributes\(5\)](#), [kerberos\(5\)](#), [smf\(5\)](#)

附注

The ktkt_warnd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/security/ktkt_warn:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名 labeld – Trusted Extensions 标签守护进程

用法概要 /usr/lib/labeld

描述 label 守护进程是当启用了 Trusted Extensions 时在引导时启动的一个系统守护进程。此守护进程会自动启动，不应直接调用。它不构成编程接口的一部分。

要启用 Trusted Extensions，请在命令行上输入以下内容：

```
% svcadm enable svc:/system/labeld
```

因为 Trusted Extensions 会影响区域的初始化和运行，因此在启用 Trusted Extensions 之前必须删除所有区域。启用之后，应重新引导系统以使 Trusted Extensions 正确初始化。

只有具有 solaris.smf.manage.labels 授权的用户或角色才能启用或禁用 Trusted Extensions。（例如，具有 "Information Security"（信息安全）或 "Object Label Management"（对象标签管理）权限配置文件的用户或角色，或者超级用户。）

在使用 Trusted Extensions 功能之前需要执行其他配置步骤。有关更多信息，请参见《Solaris Trusted Extensions Installation and Configuration Guide》。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/trusted
接口稳定性	Uncommitted（未确定）

另请参见 [svcs\(1\)](#)、[svcadm\(1M\)](#)、[syslogd\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

Solaris Trusted Extensions 管理员操作过程

附注 [labels\(5\)](#) 服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

```
svc:/system/labeld:default
```

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

引用名	labelit – list or provide labels for file systems
用法概要	<code>labelit [-F <i>FSType</i>] [-V] [-o <i>FSType-specific-options</i>] <i>special</i> [<i>operands</i>]</code>
描述	The <code>labelit</code> utility is used to write or display labels on unmounted disk file systems.
选项	The following options are supported: -F <i>FSType</i> Specify the <i>FSType</i> on which to operate. The <i>FSType</i> should either be specified here or be determinable from <code>/etc/vfstab</code> by matching the <i>special</i> with an entry in the table. If no matching entry is found, the default file system type specified in <code>/etc/default/fs</code> will be used. -V Echo complete command line. This option may be used to verify and validate the command line. Additional information obtained using a <code>/etc/vfstab</code> lookup is included in the output. The command is not executed. -o Specify <i>FSType</i>-specific options. See the manual page for the <code>labelit</code> module specific to the file system type.
操作数	The following operands are supported. If no operands are specified, <code>labelit</code> will display the value of the labels. <i>special</i> The disk partition (for example, <code>/dev/rdsk/c0t3d0s6</code>). The device may not be on a remote machine. <i>operands</i> <i>FSType</i>-specific operands. Consult the manual page of the <i>FSType</i>-specific <code>labelit</code> command for detailed descriptions.
用法	See largefile(5) for the description of the behavior of <code>labelit</code> when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).
退出状态	The following exit values are returned: 0 Write or display of labels was successful. - non-zero An error occurred.
文件	<code>/etc/vfstab</code> List of default parameters for each file system. <code>/etc/default/fs</code> Default local file system type. Default values can be set for the following flags in <code>/etc/default/fs</code>. For example: - LOCAL=ufs - LOCAL The default partition for a command if no <i>FSType</i> is specified.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [labelit_hsf\(1M\)](#), [labelit_udf\(1M\)](#), [labelit_ufs\(1M\)](#), [volcopy\(1M\)](#), [vfstab\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

附注 This utility may not be supported for all *FSTypes*.

引用名	labelit_hsf – provide and print labels for hsf file systems
用法概要	<pre>/usr/sbin/labelit -F hsf [generic_options] [-o specific_options] special</pre>
描述	labelit can be used to provide labels for unmounted CD-ROM images (CD-ROMs may not be labeled, as they are read-only media). <i>generic_options</i> are options supported by the generic labelit command. If no <i>specific_options</i> are specified, labelit prints the current value of all label fields. The <i>special</i> name should be the physical disk section (for example, /dev/dsk/c0d0s6).
选项	-o Use one or more of the following <i>name=value</i> pairs separated by commas (with no intervening spaces) to specify values for specific label fields. According to the ISO 9660 specification, only certain sets of characters may be used to fill in these labels. Thus, ‘d-characters” below refers to the characters ‘A’ through ‘Z’, the digits ‘0’ through ‘9’, and the ‘_’ (underscore) character. “a-characters” below refers to ‘A’ through ‘Z’, ‘0’ through ‘9’, space, and the following characters: !"%&'()*+,-./:;<=>?_. absfile= Abstract file identifier, d-characters, 37 characters maximum. applid= Application identifier, d-characters, 128 characters maximum. bibfile= Bibliographic file identifier, d-characters, 37 characters maximum. copyfile= Copyright file identifier, d-characters, 128 maximum. prepid= Data preparer identifier, d-characters, 128 maximum. pubid= Publisher identifier, d-characters, 128 maximum. sysid= System identifier, a-characters, 32 maximum. volid= Volume identifier, d-characters, 32 maximum. volsetid= Volume set identifier, d-characters, 128 maximum.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [labelit\(1M\)](#), [volcopy\(1M\)](#), [attributes\(5\)](#)

引用名	labelit_udfs – provide and print labels for udf file systems										
用法概要	<code>labelit -F udfs [<i>generic_options</i>] [-o <i>specific_options</i>] <i>special</i> [<i>fsname volume</i>]</code>										
描述	The <code>labelit</code> command writes labels on an unmounted disk that contains a universal disk file (udf) system. These labels can be used to identify volumes.										
选项	The following options are supported: <table><tr><td><i>generic_options</i></td><td>Specify <i>generic_options</i> supported by the generic <code>labelit</code> command. See labelit(1M) for descriptions of supported options.</td></tr><tr><td><code>-o <i>specific_options</i></code></td><td>Specify udfs-file-system-specific options in a comma-separated list with no intervening spaces. The following <i>specific_options</i> are available:</td></tr><tr><td><code>lvinfo1=<i>string</i></code></td><td>Specify information to be inserted in the LVInfo1 field of the Implementation Use Volume Descriptor. Information in LVInfo1 is generally used to identify the person creating the file system. The maximum length of the string specified is 35 bytes.</td></tr><tr><td><code>lvinfo2=<i>string</i></code></td><td>Specify information to be inserted into the LVInfo2 field of the Implementation Use Volume Descriptor. Information in LVInfo2 is generally used to identify the organization responsible for creating the file system. The maximum length of the string specified is 35 bytes.</td></tr><tr><td><code>lvinfo3=<i>string</i></code></td><td>Specify information to be inserted into the LVInfo3 field of the Implementation Use Volume Descriptor. Information in LVInfo3 is generally used to identify the contact information for the medium. The maximum length of the string specified is 35 bytes.</td></tr></table>	<i>generic_options</i>	Specify <i>generic_options</i> supported by the generic <code>labelit</code> command. See labelit(1M) for descriptions of supported options.	<code>-o <i>specific_options</i></code>	Specify udfs-file-system-specific options in a comma-separated list with no intervening spaces. The following <i>specific_options</i> are available:	<code>lvinfo1=<i>string</i></code>	Specify information to be inserted in the LVInfo1 field of the Implementation Use Volume Descriptor. Information in LVInfo1 is generally used to identify the person creating the file system. The maximum length of the string specified is 35 bytes.	<code>lvinfo2=<i>string</i></code>	Specify information to be inserted into the LVInfo2 field of the Implementation Use Volume Descriptor. Information in LVInfo2 is generally used to identify the organization responsible for creating the file system. The maximum length of the string specified is 35 bytes.	<code>lvinfo3=<i>string</i></code>	Specify information to be inserted into the LVInfo3 field of the Implementation Use Volume Descriptor. Information in LVInfo3 is generally used to identify the contact information for the medium. The maximum length of the string specified is 35 bytes.
<i>generic_options</i>	Specify <i>generic_options</i> supported by the generic <code>labelit</code> command. See labelit(1M) for descriptions of supported options.										
<code>-o <i>specific_options</i></code>	Specify udfs-file-system-specific options in a comma-separated list with no intervening spaces. The following <i>specific_options</i> are available:										
<code>lvinfo1=<i>string</i></code>	Specify information to be inserted in the LVInfo1 field of the Implementation Use Volume Descriptor. Information in LVInfo1 is generally used to identify the person creating the file system. The maximum length of the string specified is 35 bytes.										
<code>lvinfo2=<i>string</i></code>	Specify information to be inserted into the LVInfo2 field of the Implementation Use Volume Descriptor. Information in LVInfo2 is generally used to identify the organization responsible for creating the file system. The maximum length of the string specified is 35 bytes.										
<code>lvinfo3=<i>string</i></code>	Specify information to be inserted into the LVInfo3 field of the Implementation Use Volume Descriptor. Information in LVInfo3 is generally used to identify the contact information for the medium. The maximum length of the string specified is 35 bytes.										
操作数	The following operands are supported: <table><tr><td><i>special</i></td><td>Specify <i>special</i> as the physical disk slice, for example, <code>/dev/rdisk/c0t0d0s6</code>. The device can not be on a remote machine.</td></tr><tr><td><i>fsname</i></td><td>Specify <i>fsname</i> as the mount point, (for example, <code>root</code>, <code>u1</code>, and so forth), of the file system.</td></tr><tr><td><i>volume</i></td><td>Specify <i>volume</i> as the physical volume name.</td></tr></table> If none of the options (<i>fsname</i>, <i>volume</i>, <i>specific_options</i>) is specified, <code>labelit</code> prints the current values of <i>fsname</i>, <i>volume</i>, LVInfo1, LVInfo2 and LVInfo3.	<i>special</i>	Specify <i>special</i> as the physical disk slice, for example, <code>/dev/rdisk/c0t0d0s6</code> . The device can not be on a remote machine.	<i>fsname</i>	Specify <i>fsname</i> as the mount point, (for example, <code>root</code> , <code>u1</code> , and so forth), of the file system.	<i>volume</i>	Specify <i>volume</i> as the physical volume name.				
<i>special</i>	Specify <i>special</i> as the physical disk slice, for example, <code>/dev/rdisk/c0t0d0s6</code> . The device can not be on a remote machine.										
<i>fsname</i>	Specify <i>fsname</i> as the mount point, (for example, <code>root</code> , <code>u1</code> , and so forth), of the file system.										
<i>volume</i>	Specify <i>volume</i> as the physical volume name.										

退出状态 The following exit values are returned:

0 Successful completion.

non-zero An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/udfs

另请参见 [labelit\(1M\)](#), [attributes\(5\)](#)

引用名	labelit_ufs – provide and print labels for ufs file systems
用法概要	labelit -F ufs [<i>generic_options</i>] <i>special</i> [<i>fsname</i> <i>volume</i>]
描述	labelit is used to write labels on unmounted disk file systems. Such labels may be used to uniquely identify volumes and are used by volume-oriented programs such as volcopy(1M) .
选项	The following option is supported: <i>generic_options</i> options supported by the generic labelit command. See labelit(1M).
操作数	The following operands are supported: <i>special</i> name should be the physical disk section (for example, /dev/dsk/c0d0s6). The device may not be on a remote machine. <i>fsname</i> represents the mount point (for example, root, u1, and so on) of the file system. <i>volume</i> may be used to represent the physical volume name. If <i>fsname</i> and <i>volume</i> are not specified, labelit prints the current values of these labels. Both <i>fsname</i> and <i>volume</i> are limited to six or fewer characters.
退出状态	The following exit values are returned: 0 Write or display of labels was successful. non-zero An error occurred.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [labelit\(1M\)](#), [volcopy\(1M\)](#), [attributes\(5\)](#), [ufs\(7FS\)](#)

引用名	latencytop – 报告系统和应用程序中与等待时间相关的统计信息
用法概要	<pre>latencytop [-t interval] [-o log_file] [-k log_level] [-f [no]feature,...] [-l log_interval] [-h] [-s pid=PID pgid=PGID]</pre>
描述	LatencyTOP 是一个监测工具，用于报告系统和应用程序中与等待时间相关的统计信息。此工具报告系统中以及系统上正在运行的应用程序中等待时间的发生位置和类型的相关统计信息。可以根据该统计信息来消除所查明的等待时间，从而提高应用程序和系统的性能吞吐量。 此工具定期分析系统活动，并在输出窗口中显示数据。将跟踪两种类型的等待时间：进入休眠状态和从休眠状态唤醒的 LWP，以及调整顺序以获取同步对象的 LWP。此工具使用 Solaris DTrace 框架收集与系统和应用程序 LWP 的这两种非活动情况相对应的统计信息。 输出窗口分成两部分。上部分显示系统范围的统计信息，而下部分显示有关各个进程的统计信息。用户可使用向左方向键和向右方向键在进程列表中导航，并选择感兴趣的列表。然后此工具会在窗口下部分中显示该选定进程的相关统计信息。如果按下 t 或 T 键，此工具将显示该选定进程的特定于 LWP 的视图。可使用 t 或 T 键在进程视图和线程视图之间进行切换。 执行期间，用户可通过按下 r 或 R 键强制刷新分析结果。间隔时间将恢复为缺省值或指定值（如果使用 -t）。要退出应用程序，用户必须按下 q 或 Q 键。
选项	所支持的选项如下： -f, --feature [no]feature1,[no]feature2,... 在 LatencyTOP 中启用/禁用功能。功能只能为下列项之一： [no]filter 过滤较长的中断等待时间，例如，休眠。缺省值为 off。 [no]sched 监视 sched (PID=0)。缺省值为 "off"。 [no]sobj 监视同步对象。缺省值为 on。 [no]low 通过抽样少量的等待时间来降低开销。启用此功能后，将统计评估少量的等待时间，从而降低 CPU 使用率。请将其用于繁重的工作负荷，如繁忙的 Web 服务器。缺省值为 "off"。 -h 显示命令的用法。 -k log_level 指定日志文件中的日志记录的级别。有效值包括： 0 none（缺省值）

- 1 unknown
- 2 all
- l [*log_interval*]
每 *log_interval* 秒向日志文件写入一次数据；*log_interval* 必须大于 60。
- o *log_file*
指定输出将写入的日志文件。缺省日志文件为 `/var/log/latencytop.log`。
- s *pid=PID* | *pgid=PGID*
仅跟踪指定的进程或指定的进程组，并仅显示与此进程或进程组相关的数据。
- t *interval*
指定此工具从系统收集统计信息的时间间隔（以秒为单位）。可能值介于 1 和 60 之间；缺省值为 5 秒。

示例

示例1 运行此工具

以下命令启动此工具，并且使用各个选项的缺省值。

```
% latencytop
```

示例2 设置时间间隔

以下命令将抽样时间间隔设置为两秒。

```
% latencytop -t 2
```

示例3 指定日志文件

以下命令将日志文件设置为 `/tmp/latencytop.log`。

```
% latencytop -o /tmp/latencytop.log
```

示例4 指定日志级别

以下命令将日志级别设置为 `all`。

```
% latencytop -l 2
```

示例5 启用对等待时间的跟踪

以下命令启用对同步对象导致的等待时间的跟踪。

```
% latencytop -f sobj
```

示例6 显示进程组的数据

以下命令显示属于进程组 630 的进程的跟踪日期。

```
% latencytop -s pgid=630
```

退出状态

0
操作成功。

1
出现错误。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
体系结构	x86、SPARC
可用性	diagnostic/latencytop
接口稳定性	Volatile（可变）

另请参见

[dtrace\(1M\)](#)、[kstat\(1M\)](#)、[attributes\(5\)](#)

用法

您必须具有 DTrace 特权才能运行 LatencyTOP。

引用名

ldapaddent – create LDAP entries from corresponding /etc files

用法概要

```
ldapaddent [-cpv] [-a authenticationMethod] [-b baseDN]
           -D bindDN [-w bind_password] [-j passwdFile] [-f filename]
           database

ldapaddent [-cpv] -a sasl/GSSAPI [-b baseDN] [-f filename]
           database

ldapaddent -d [-v] [-a authenticationMethod] [-D bindDN]
           [-w bind_password] [-j passwdFile] database

ldapaddent [-cpv] -h LDAP_server[:serverPort] [-M domainName]
           [-N profileName] [-P certifPath] [-a authenticationMethod]
           [-b baseDN] -D bindDN [-w bind_password] [-f filename]
           [-j passwdFile] database

ldapaddent [-cpv] -h LDAP_server[:serverPort] [-M domainName]
           [-N profileName] [-P certifPath] [-a authenticationMethod]
           [-b baseDN] [-f filename] database

ldapaddent -d [-v] -h LDAP_server[:serverPort] [-M domainName]
           [-N profileName] [-P certifPath] [-a authenticationMethod]
           [-b baseDN] -D bindDN [-w bind_password] [-j passwdFile]
           database
```

描述

ldapaddent creates entries in LDAP containers from their corresponding /etc files. This operation is customized for each of the standard containers that are used in the administration of Solaris systems. The *database* argument specifies the type of the data being processed. Legal values for this type are one of *aliases*, *auto_**, *bootparams*, *ethers*, *group*, *hosts* (including both IPv4 and IPv6 addresses), *ipnodes* (alias for *hosts*), *netgroup*, *netmasks*, *networks*, *passwd*, *shadow*, *protocols*, *publickey*, *rpc*, and *services*. In addition to the preceding, the *database* argument can be one of the RBAC-related files (see [rbac\(5\)](#)):

- /etc/user_attr
- /etc/security/auth_attr
- /etc/security/prof_attr
- /etc/security/exec_attr

By default, ldapaddent reads from the standard input and adds this data to the LDAP container associated with the database specified on the command line. An input file from which data can be read is specified using the *-f* option.

If you specify the *-h* option, ldapaddent establishes a connection to the server indicated by the option in order to obtain a *DUAProfile* specified by the *-N* option. The entries will be stored in the directory described by the configuration obtained.

By default (if the *-h* option is not specified), entries will be stored in the directory based on the client's configuration. To use the utility in the default mode, the Solaris LDAP client must be set up in advance.

The location where entries are to be written can be overridden by using the `-b` option.

If the entry to be added exists in the directory, the command displays an error and exits, unless the `-c` option is used.

Although, there is a shadow database type, there is no corresponding shadow container. Both the shadow and the passwd data is stored in the `people` container itself. Similarly, data from networks and netmasks databases are stored in the `networks` container.

The `user_attr` data is stored by default in the `people` container. The `prof_attr` and `exec_attr` data is stored by default in the `SolarisProfAttr` container.

You must add entries from the passwd database before you attempt to add entries from the shadow database. The addition of a shadow entry that does not have a corresponding passwd entry will fail.

The passwd database must precede the `user_attr` database.

For better performance, the recommended order in which the databases should be loaded is as follows:

- passwd database followed by shadow database
- networks database followed by netmasks database
- bootparams database followed by ethers database

Only the first entry of a given type that is encountered will be added to the LDAP server. The `ldapaddent` command skips any duplicate entries.

选项

The `ldapaddent` command supports the following options:

`-a authenticationMethod`

Specify authentication method. The default value is what has been configured in the profile.

The supported authentication methods are:

```
simple
sasl/CRAM-MD5
sasl/DIGEST-MD5
sasl/GSSAPI
tls:simple
tls:sasl/CRAM-MD5
tls:sasl/DIGEST-MD5
```

Selecting `simple` causes passwords to be sent over the network in clear text. Its use is strongly discouraged. Additionally, if the client is configured with a profile which uses no authentication, that is, either the `credentialLevel` attribute is set to `anonymous` or `authenticationMethod` is set to `none`, the user must use this option to provide an authentication method. If the authentication method is `sasl/GSSAPI`, `bindDN` and `bindPassword` is not required and the `hosts` and `ipnodes` fields of `/etc/nsswitch.conf` must be configured as:

```
hosts: dns files
ipnodes: dns files
```

See `nsswitch.conf(4)`.

-b *baseDN*

Create entries in the *baseDN* directory. *baseDN* is not relative to the client's default search base, but rather, it is the actual location where the entries will be created. If this parameter is not specified, the first search descriptor defined for the service or the default container will be used.

-c

Continue adding entries to the directory even after an error. Entries will not be added if the directory server is not responding or if there is an authentication problem.

-D *bindDN*

Create an entry which has write permission to the *baseDN*. When used with **-d** option, this entry only needs read permission.

-d

Dump the LDAP container to the standard output in the appropriate format for the given database.

-f *filename*

Indicates input file to read in an `/etc/` file format.

-h *LDAP_server[:serverPort]*

Specify an address (or a name) and an optional port of the LDAP server in which the entries will be stored. The current naming service specified in the `nsswitch.conf` file is used. The default value for the port is 389, except when TLS is specified as the authentication method. In this case, the default LDAP server port number is 636.

The format to specify the address and port number for an IPv6 address is:

```
[ipv6_addr]:port
```

To specify the address and port number for an IPv4 address, use the following format:

```
ipv4_addr:port
```

If the host name is specified, use the format:

```
host_name:port
```

-j *passwdFile*

Specify a file containing the password for the bind DN or the password for the SSL client's key database. To protect the password, use this option in scripts and place the password in a secure file. This option is mutually exclusive of the **-w** option.

-M *domainName*

The name of a domain served by the specified server. If not specified, the default domain name will be used.

-N *profileName*

Specify the DUAProfile name. A profile with such a name is supposed to exist on the server specified by -h option. Otherwise, a default DUAProfile will be used. The default value is default.

-P *certifPath*

The certificate path for the location of the certificate database. The value is the path where security database files reside. This is used for TLS support, which is specified in the authenticationMethod and serviceAuthenticationMethod attributes. The default is /var/ldap.

-p

Process the password field when loading password information from a file. By default, the password field is ignored because it is usually not valid, as the actual password appears in a shadow file.

-w *bindPassword*

Password to be used for authenticating the *bindDN*. If this parameter is missing, the command will prompt for a password. NULL passwords are not supported in LDAP.

When you use -w *bindPassword* to specify the password to be used for authentication, the password is visible to other users of the system by means of the ps command, in script files or in shell history.

If you supply “-” (hyphen) as a password, you will be prompted to enter a password.

-v

Verbose.

操作数

The following operands are supported:

database

The name of the database or service name. Supported values are: aliases, auto_*, bootparams, ethers, group, hosts (including IPv6 addresses), netgroup, netmasks, networks, passwd, shadow, protocols, publickey, rpc, and services. Also supported are auth_attr, prof_attr, exec_attr, user_attr, and projects.

示例**示例 1 Adding Password Entries to the Directory Server**

The following example shows how to add password entries to the directory server:

```
example# ldapaddent -D "cn=directory manager" -w secret \
-f /etc/passwd passwd
```

示例 2 Adding Group Entries

The following example shows how to add group entries to the directory server using sasl/CRAM-MD5 as the authentication method:

```
example# ldapaddent -D "cn=directory manager" -w secret \
-a "sasl/CRAM-MD5" -f /etc/group group
```

示例 3 Adding auto_master Entries

The following example shows how to add auto_master entries to the directory server:

```
example# ldapaddent -D "cn=directory manager" -w secret \
-f /etc/auto_master auto_master
```

示例 4 Dumping passwd Entries from the Directory to File

The following example shows how to dump password entries from the directory to a file foo:

```
example# ldapaddent -d passwd > foo
```

示例 5 Adding Password Entries to a Specific Directory Server

The following example shows how to add password entries to a directory server that you specify:

```
example# ldapaddent -h 10.10.10.10:3890 \
-M another.domain.name -N special_duaprofile \
-D "cn=directory manager" -w secret \
-f /etc/passwd passwd
```

退出状态

The following exit values are returned:

- 0
Successful completion.
- >0
An error occurred.

文件

/var/ldap/ldap_client_file
/var/ldap/ldap_client_cred
Files containing the LDAP configuration of the client. These files are not to be modified manually. Their content is not guaranteed to be human readable. Use [ldapclient\(1M\)](#) to update these files.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/nis
Interface Stability	Committed

另请参见

[ldaplist\(1\)](#), [ldapmodify\(1\)](#), [ldapmodrdn\(1\)](#), [ldapsearch\(1\)](#), [idsconfig\(1M\)](#),
[ldapclient\(1M\)](#), [nsswitch.conf\(4\)](#), [attributes\(5\)](#)

[《Oracle Solaris 管理：安全服务》](#)

注意

Currently StartTLS is not supported by `libldap.so.5`, therefore the port number provided refers to the port used during a TLS open, rather than the port used as part of a StartTLS sequence. For example:

```
-h foo:1000 -a tls:simple
```

The preceding refers to a raw TLS open on host `foo` port `1000`, not an open, StartTLS sequence on an unsecured port `1000`. If port `1000` is unsecured the connection will not be made.

引用名	ldap_cachemgr – LDAP daemon to manage client configuration for LDAP based Network Information Service lookups
用法概要	<code>/usr/lib/ldap/ldap_cachemgr [-l <i>log-file</i>] [-g]</code>
描述	The <code>ldap_cachemgr</code> daemon is a process that provides an up-to-date configuration cache for LDAP naming services. It is started during multi-user boot. The <code>ldap_cachemgr</code> utility provides caching for all parameters as specified and used by the LDAP naming service clients. The <code>ldap_cachemgr</code> utility uses the cache files which are originally created by executing the ldapclient(1M) utility, as cold start files. Updates to the cache files take place dynamically if profiles are used to configure the client. See the <code>init</code> option to ldapclient(1M). The <code>ldap_cachemgr</code> utility helps improve the performance of the clients that are using LDAP as the Naming service repository. In order for the LDAP naming services to function properly, the <code>ldap_cachemgr</code> daemon must be running. <code>ldap_cachemgr</code> also improves system security by making the configuration files readable by superuser only. The cache maintained by this daemon is shared by all the processes that access LDAP Naming information. All processes access this cache through a door call. On startup, <code>ldap_cachemgr</code> initializes the cache from the cache files. See ldapclient(1M). Thus, the cache survives machine reboots. The <code>ldap_cachemgr</code> daemon also acts as its own administration tool. If an instance of <code>ldap_cachemgr</code> is already running, commands are passed transparently to the running version.
选项	The following options are supported: <code>-g</code> Print current configuration and statistics to standard output. This is the only option executable without superuser privileges. <code>-l <i>log-file</i></code> Cause <code>ldap_cachemgr</code> to use a log file other than the default <code>/var/ldap/cachemgr.log</code>.
示例	示例 1 Stopping and Restarting the <code>ldap_cachemgr</code> Daemon The following example shows how to stop and to restart the <code>ldap_cachemgr</code> daemon. <pre>example# svcadm disable network/ldap/client example# svcadm enable network/ldap/client</pre> 示例 2 Forcing <code>ldap_cachemgr</code> to Reread Configuration Files The following example shows how to force <code>ldap_cachemgr</code> to reread the <code>/var/ldap/ldap_client_file</code> and <code>/var/ldap/ldap_client_cred</code> files <pre>example# pkill -HUP ldap_cachemgr</pre>

文件 /var/ldap/cachemgr.log
 Default log file.

 /var/ldap/ldap_client_file
 /var/ldap/ldap_client_cred
 Files containing the LDAP configuration of the client. These files are not to be modified manually. Their content is not guaranteed to be human readable. Use [ldapclient\(1M\)](#) to update these files.

警告 The ldap_cachemgr utility is included in the current Solaris release on an uncommitted basis only. It is subject to change or removal in a future minor release.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/nis

另请参见 [ldapadd\(1\)](#), [ldapdelete\(1\)](#), [ldaplist\(1\)](#), [ldapmodify\(1\)](#), [ldapmodrdn\(1\)](#), [ldapsearch\(1\)](#), [pkill\(1\)](#), [svcs\(1\)](#), [idsconfig\(1M\)](#), [ldapaddent\(1M\)](#), [ldapclient\(1M\)](#), [svcadm\(1M\)](#), [signal.h\(3HEAD\)](#), [resolv.conf\(4\)](#), [attributes\(5\)](#), [ldap\(5\)](#), [smf\(5\)](#)

附注 The ldap_cachemgr service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

 svc:/network/ldap/client

 Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名 ldapclient – initialize LDAP client machine or output an LDAP client profile in LDIF format

用法概要

```
/usr/sbin/ldapclient [-v | -q] init [-a profileName=profileName]
    [-a domainName=domain] [-a proxyDN=proxyDN]
    [-a proxyPassword=password]
    [-a authenticationMethod=authenticationMethod]
    [-a enableShadowUpdate=true | false]
    [-a adminDN=adminDN]
    [-a adminPassword=adminPassword]
    [-a certificatePath=path] [-d bindDN] [-w bindPassword]
    [-j passwdFile] [-y passwdFile]
    [-z adminrPasswdFile] LDAP_server[:port_number]

/usr/sbin/ldapclient [-v | -q] manual [-a attrName=attrVal]

/usr/sbin/ldapclient [-v | -q] mod [-a attrName=attrVal]

/usr/sbin/ldapclient [-v | -q] list

/usr/sbin/ldapclient [-v | -q] uninit

/usr/sbin/ldapclient [-v | -q] genprofile -a profileName=profileName
    [-a attrName=attrVal]
```

描述

The `ldapclient` utility can be used to:

- initialize LDAP client machines
- restore the network service environment on LDAP clients
- list the contents of the LDAP client cache in human readable format.

The `init` form of the `ldapclient` utility is used to initialize an LDAP client machine, using a profile stored on an LDAP server specified by `LDAP_server`. The LDAP client will use the attributes in the specified profile to determine the configuration of the LDAP client. Using a configuration profile allows for easy installation of LDAP client and propagation of configuration changes to LDAP clients. The [ldap_cachemgr\(1M\)](#) utility will update the LDAP client configuration when its cache expires by reading the profile. For more information on the configuration profile refer to IETF document A Configuration Schema for LDAP Based Directory User Agents.

The `manual` form of the `ldapclient` utility is used to initialize an LDAP client machine manually. The LDAP client will use the attributes specified on the command line. Any unspecified attributes will be assigned their default values. At least one server must be specified in the `defaultServerList` or the `preferredServerList` attributes. The `domainName` attribute must be specified if the client's `domainName` is not set.

The `mod` form of the `ldapclient` utility is used to modify the configuration of an LDAP client machine that was setup manually. This option modifies only those LDAP client configuration attributes specified on the command line. The `mod` option should only be used on LDAP clients that were initialized using the `manual` option.

Regardless of which method is used for initialization, if a client is to be configured to use a `proxy credentialLevel`, proxy credentials must be provided using `-a proxyDN=proxyDN` and `-a proxyPassword=proxyPassword` options. However, if `-a proxyPassword=proxyPassword` is not specified, `ldapclient` will prompt for it. Note that NULL passwords are not allowed in LDAP. If a `self credentialLevel` is configured, `authenticationMethod` must be `sasl/GSSAPI`.

Similarly, if a client is to be configured to enable shadow information update and use a `proxy credentialLevel`, administrator credentials must be provided using `-a adminDN=adminDN` and `-a adminPassword=adminPassword`. However, the shadow information update does not need the administrator credentials if a `self credentialLevel` is configured.

The naming service-specific configuration properties are stored in the `svc:/network/ldap/client` SMF service. Modifying the SMF properties directly is not advised. Use this tool instead.

Other configuration might be modified during installation. It will be backed up to `/var/ldap/restore`. The files that are typically modified during initialization are:

- `/etc/nsswitch.conf`
- `/etc/defaultdomain` (if it exists)
- `/var/yp/binding/'domainname'` (for a NIS [YP] client)

`ldapclient` does not set up a client to resolve hostnames using DNS. It simply copies `/etc/nsswitch.ldap` to `/etc/nsswitch.conf`. If you prefer to use DNS for host resolution, please refer to the DNS documentation for information on setting up DNS. See [resolv.conf\(4\)](#). If you want to use `sasl/GSSAPI` as the authentication method, you have to use DNS for hosts and ipnodes resolution.

The `list` form of the `ldapclient` utility is used to list the LDAP client configuration. The output will be human readable. LDAP configuration files are not guaranteed to be human readable. Note that for security reason, the values for `adminDN` and `adminPassword` will not be displayed.

The `uninit` form of the `ldapclient` utility is used to uninitialize the network service environment, restoring it to the state it was in prior to the last execution of `ldapclient` using `init` or `manual`. The restoration will succeed only if the machine was initialized with the `init` or `manual` form of `ldapclient`, as it uses the backup files created by these options.

The `genprofile` option is used to write an LDIF formatted configuration profile based on the attributes specified on the command line to standard output. This profile can then be loaded into an LDAP server to be used as the client profile, which can be downloaded by means of the `ldapclient init` command. Loading the LDIF formatted profile to the directory server can be done through [ldapadd\(1\)](#), or through any server specific import tool. Note that the attributes `proxyDN`, `proxyPassword`, `certificatePath`, `domainName`, `enableShadowUpdate`, `adminDN`, and `adminPassword` are not part of the configuration profile and thus are not permitted.

You must have superuser privileges to run the `ldapclient` command, except with the `genprofile` option.

To access the information stored in the directory, clients can either authenticate to the directory, or use an unauthenticated connection. The LDAP client is configured to have a credential level of either anonymous or proxy. In the first case, the client does not authenticate to the directory. In the second case, client authenticates to the directory using a proxy identity for read access, and using a administrator identity for write access if `enableShadowUpdate` is configured. In the third case, client authenticates to the directory using a Kerberos principal that is mapped to an LDAP identity by the LDAP server. Refer to the chapter on implementing security in the [《Oracle Solaris Administration: Naming and Directory Services》](#) or your appropriate directory server documentation for identity mapping details.

If a client is configured to use an identity, you can configure which authentication method the client will use. The LDAP client supports the following authentication methods:

```
none
simple
sasl/CRAM-MD5
sasl/DIGEST-MD5
sasl/GSSAPI
tls:simple
tls:sasl/CRAM-MD5
tls:sasl/DIGEST-MD5
```

Note that some directory servers may not support all of these authentication methods. For `simple`, be aware that the bind password will be sent in the clear to the LDAP server. For those authentication methods using TLS (transport layer security), the entire session is encrypted. You will need to install the appropriate certificate databases to use TLS.

Commands

The following commands are supported:

```
init
    Initialize client from a profile on a server.

manual
    Manually initialize client with the specified attribute values.

mod
    Modify attribute values in the configuration file after a manual initialization of the client.

list
    Write the contents of the LDAP client cache to standard output in human readable form.

uninit
    Uninitialize an LDAP client, assuming that ldapclient was used to initialize the client.
```


genprofile

Generate a configuration profile in LDIF format that can then be stored in the directory for clients to use, with the `init` form of this command.

Attributes

The following attributes are supported:

adminDN

Specify the Bind Distinguished Name for the administrator identity that is used for shadow information update. This option is required if the credential level is `proxy`, and `enableShadowUpdate` is set to `true`. There is no default value.

adminPassword

Specify the administrator password. This option is required if the credential level is `proxy`, and `enableShadowUpdate` is set to `true`. There is no default value.

attributeMap

Specify a mapping from an attribute defined by a service to an attribute in an alternative schema. This can be used to change the default schema used for a given service. The syntax of `attributeMap` is defined in the profile IETF draft. This option can be specified multiple times. The default value for all services is `NULL`. In the example,

```
attributeMap: passwd:uid=employeeNumber
```

the LDAP client would use the LDAP attribute `employeeNumber` rather than `uid` for the `passwd` service. This is a multivalued attribute.

authenticationMethod

Specify the default authentication method used by all services unless overridden by the `serviceAuthenticationMethod` attribute. Multiple values can be specified by using a semicolon-separated list. The default value is `none`. For those services that use `credentialLevel` and `credentialLevel` is `anonymous`, this attribute is ignored. Services such as `pam_ldap` will use this attribute, even if `credentialLevel` is `anonymous`. The supported authentication methods are described above. If the `authenticationMethod` is `sasl/GSSAPI`, the `hosts` and `ipnodes` of `/etc/nsswitch.conf` must be configured with DNS support, for example:

```
hosts: dns files
ipnodes: dns files
```

bindTimeLimit

The maximum time in seconds that a client should spend performing a bind operation. Set this to a positive integer. The default value is 30.

certificatePath

The certificate path for the location of the certificate database. The value is the path where security database files reside. This is used for TLS support, which is specified in the `authenticationMethod` and `serviceAuthenticationMethod` attributes. The default is `/var/ldap`.

credentialLevel

Specify the credential level the client should use to contact the directory. The credential levels supported are either `anonymous` or `proxy`. If a proxy credential level is specified, then the `authenticationMethod` attribute must be specified to determine the authentication mechanism. Also, if the credential level is `proxy` and at least one of the authentication methods require a bind DN, the `proxyDN` and `proxyPassword` attribute values must be set. In addition, if `enableShadowUpdate` is set to `true`, the `adminDN` and `adminPassword` values must be set. If a self credential level is specified, the `authenticationMethod` must be `sasl/GSSAPI`.

defaultSearchBase

Specify the default search base DN. There is no default. The `serviceSearchDescriptor` attribute can be used to override the `defaultSearchBase` for given services.

defaultSearchScope=one | sub

Specify the default search scope for the client's search operations. This default can be overridden for a given service by specifying a `serviceSearchDescriptor`. The default is one level search.

defaultServerList

A space separated list of server names or server addresses, either IPv4 or IPv6. If you specify server names, be sure that the LDAP client can resolve the name without the LDAP name service. You must resolve the LDAP servers' names by using either `files` or `dns`. If the LDAP server name cannot be resolved, your naming service will fail.

The port number is optional. If not specified, the default LDAP server port number 389 is used, except when TLS is specified in the authentication method. In this case, the default LDAP server port number is 636.

The format to specify the port number for an IPv6 address is:

```
[ipv6_addr]:port
```

To specify the port number for an IPv4 address, use the following format:

```
ipv4_addr:port
```

If the host name is specified, use the format:

```
host_name:port
```

If you use TLS, the LDAP server's hostname must match the hostname in the TLS certificate. Typically, the hostname in the TLS certificate is a fully qualified domain name. With TLS, the LDAP server host addresses must resolve to the hostnames in the TLS certificate. You must use `files` or `dns` to resolve the host address.

domainName

Specify the DNS domain name. This becomes the default domain for the machine. The default is the current domain name. This attribute is only used in client initialization.

`enableShadowUpdate=true | false`

Specify whether the client is allowed to update shadow information. If set to `true` and the credential level is `proxy`, `adminDN` and `adminPassword` must be specified.

`followReferrals=true | false`

Specify the referral setting. A setting of `true` implies that referrals will be automatically followed and `false` would result in referrals not being followed. The default is `true`.

`objectclassMap`

Specify a mapping from an `objectclass` defined by a service to an `objectclass` in an alternative schema. This can be used to change the default schema used for a given service. The syntax of `objectclassMap` is defined in the profile IETF draft. This option can be specified multiple times. The default value for all services is `NULL`. In the example,

```
objectclassMap=passwd:posixAccount=unixAccount
```

the LDAP client would use the LDAP `objectclass` of `unixAccount` rather than the `posixAccount` for the `passwd` service. This is a multivalued attribute.

`preferredServerList`

Specify the space separated list of server names or server addresses, either IPv4 or IPv6, to be contacted before servers specified by the `defaultServerList` attribute. If you specify server names, be sure that the LDAP client can resolve the name without the LDAP name service. You must resolve the LDAP servers' names by using either `files` or `dns`. If the LDAP server name cannot be resolved, your naming service will fail.

The port number is optional. If not specified, the default LDAP server port number 389 is used, except when TLS is specified in the authentication method. In this case, the default LDAP server port number is 636.

The format to specify the port number for an IPv6 address is:

```
[ipv6_addr]:port
```

To specify the port number for an IPv4 address, use the following format:

```
ipv4_addr:port
```

If the host name is specified, use the format:

```
host_name:port
```

If you use TLS, the LDAP server's hostname must match the hostname in the TLS certificate. Typically, the hostname in the TLS certificate is a fully qualified domain name. With TLS, the LDAP server host addresses must resolve to the hostnames in the TLS certificate. You must use `files` or `dns` to resolve the host address.

`profileName`

Specify the profile name. For `ldapclient init`, this attribute is the name of an existing profile which may be downloaded periodically depending on the value of the `profileTTL` attribute. For `ldapclient genprofile`, this is the name of the profile to be generated. The default value is `default`.

profileTTL

Specify the TTL value in seconds for the client information. This is only relevant if the machine was initialized with a client profile. If you do not want [ldap_cachemgr\(1M\)](#) to attempt to refresh the LDAP client configuration from the LDAP server, set `profileTTL` to 0 (zero). Valid values are either zero 0 (for no expiration) or a positive integer in seconds. The default value is 12 hours.

proxyDN

Specify the Bind Distinguished Name for the proxy identity. This option is required if the credential level is proxy, and at least one of the authentication methods requires a bind DN. There is no default value.

proxyPassword

Specify client proxy password. This option is required if the credential level is proxy, and at least one of the authentication methods requires a bind DN. There is no default.

searchTimeLimit

Specify maximum number of seconds allowed for an LDAP search operation. The default is 30 seconds. The server may have its own search time limit.

serviceAuthenticationMethod

Specify authentication methods to be used by a service in the form *servicename:authenticationmethod*, for example:

```
pam_ldap:tls:simple
```

For multiple authentication methods, use a semicolon-separated list. The default value is no service authentication methods, in which case, each service would default to the `authenticationMethod` value. The supported authentications are described above.

Three services support this feature: `passwd -cmd`, `keyserv`, and `pam_ldap`. The `passwd -cmd` service is used to define the authentication method to be used by [passwd\(1\)](#) to change the user's password and other attributes. The `keyserv` service is used to identify the authentication method to be used by the [chkey\(1\)](#) and [newkey\(1M\)](#) utilities. The `pam_ldap` service defines the authentication method to be used for authenticating users when [pam_ldap\(5\)](#) is configured. If this attribute is not set for any of these services, the `authenticationMethod` attribute is used to define the authentication method. This is a multivalued attribute.

serviceCredentialLevel

Specify credential level to be used by a service. Multiple values can be specified in a space-separated list. The default value for all services is `NULL`. The supported credential levels are: `anonymous` or `proxy`. At present, no service uses this attribute. This is a multivalued attribute.

serviceSearchDescriptor

Override the default base DN for LDAP searches for a given service. The format of the descriptors also allow overriding the default search scope and search filter for each service.

The syntax of `serviceSearchDescriptor` is defined in the profile IETF draft. The default value for all services is NULL. This is a multivalued attribute. In the example,

```
serviceSearchDescriptor=passwd:ou=people,dc=a1,dc=acme,dc=com?one
```

the LDAP client would do a one level search in `ou=people,dc=a1,dc=acme,dc=com` rather than `ou=people,defaultSearchBase` for the `passwd` service.

选项

The following options are supported:

`-a attrName=attrValue`

Specify `attrName` and its value. See SYNOPSIS for a complete list of possible attribute names and values.

`-D bindDN`

Specifies an entry that has read permission for the requested database.

`-j passwdFile`

Specify a file containing the password for the bind DN or the password for the SSL client's key database. To protect the password, use this option in scripts and place the password in a secure file. This option is mutually exclusive of the `-w` option.

`-q`

Quiet mode. No output is generated.

`-v`

Verbose output.

`-w bindPassword`

Password to be used for authenticating the bind DN. If this parameter is missing, the command will prompt for a password. NULL passwords are not supported in LDAP.

When you use `-w bindPassword` to specify the password to be used for authentication, the password is visible to other users of the system by means of the `ps` command, in script files, or in shell history.

If you supply “-” (hyphen) as a password, the command will prompt for a password.

`-y passwdFile`

Specify a file containing the password for the proxy DN. To protect the password, use this option in scripts and place the password in a secure file. This option is mutually exclusive of the `-a proxyPassword` option.

`-z adminrPasswdFile`

Specify a file containing the password for the `adminDN`. To protect the password, use this option in scripts and place the password in a secure file. This option is mutually exclusive of the `-a adminPassword` option.

操作数

The following operand is supported:

LDAP_server

An address or a name for the LDAP server from which the profile will be loaded. The current naming service specified in the `nsswitch.conf` file is used. Once the profile is loaded, the `preferredServerList` and `defaultServerList` specified in the profile are used.

示例**示例 1** Setting Up a Client By Using the Default Profile Stored on a Specified LDAP Server

The following example shows how to set up a client using the default profile stored on the specified LDAP server. This command will only be successful if either the credential level in the profile is set to anonymous or the authentication method is set to none.

```
example# ldapclient init 172.16.100.1
```

示例 2 Setting Up a Client By Using the simple Profile Stored on a Specified LDAP Server

The following example shows how to set up a client using the simple profile stored on the specified LDAP server. The domainname is set to `xyz.mycompany.com` and the `proxyPassword` is secret.

```
example# ldapclient init -a profileName=simple \
-a domainName=xyz.mycompany.com \
-a proxyDN=cn=proxyagent,ou=profile,dc=xyz,dc=mycompany,dc=com \
-a proxyPassword=secret '["fe80::a00:20ff:fea3:388"]':386
```

示例 3 Setting Up a Client Using Only One Server

The following example shows how to set up a client using only one server. The authentication method is set to none, and the search base is `dc=mycompany,dc=com`.

```
example# ldapclient manual -a authenticationMethod=none \
-a defaultSearchBase=dc=mycompany,dc=com \
-a defaultServerList=172.16.100.1
```

示例 4 Setting Up a Client Using Only One Server That Does Not Follow Referrals

The following example shows how to set up a client using only one server. The credential level is set to proxy. The authentication method of is `sasl/CRAM-MD5`, with the option not to follow referrals. The domain name is `xyz.mycompany.com`, and the LDAP server is running on port number 386 at IP address `172.16.100.1`.

```
example# ldapclient manual \
-a credentialLevel=proxy \
-a authenticationMethod=sasl/CRAM-MD5 \
-a proxyPassword=secret \
-a proxyDN=cn=proxyagent,ou=profile,dc=xyz,dc=mycompany,dc=com \
-a defaultSearchBase=dc=xyz,dc=mycompany,dc=com \
-a domainName=xyz.mycompany.com \
-a followReferrals=false \
```

示例 4 Setting Up a Client Using Only One Server That Does Not Follow Referrals (续)

```
-a defaultServerList=172.16.100.1:386
```

示例 5 Using genprofile to Set Only the defaultSearchBase and the Server Addresses

The following example shows how to use the genprofile command to set the defaultSearchBase and the server addresses.

```
example# ldapclient genprofile -a profileName=myprofile \
-a defaultSearchBase=dc=eng,dc=sun,dc=com \
-a "defaultServerList=172.16.100.1 172.16.234.15:386" \
> myprofile.ldif
```

示例 6 Creating a Profile on IPv6 servers

The following example creates a profile on IPv6 servers

```
example# ldapclient genprofile -a profileName=eng \
-a credentialLevel=proxy \
-a authenticationMethod=sasl/DIGEST-MD5 \
-a defaultSearchBase=dc=eng,dc=acme,dc=com \
-a "serviceSearchDescriptor=passwd:ou=people,dc=a1,dc=acme,dc=com?one" \
-a preferredServerList= '['fe80::a00:20ff:fea3:388']' \
-a "defaultServerList='['fec0::111:a00:20ff:fea3:edcf']' \
    '['fec0::111:a00:20ff:feb5:e41']'" > eng.ldif
```

示例 7 Creating a Profile That Overrides Every Default Value

The following example shows a profile that overrides every default value.

```
example# ldapclient genprofile -a profileName=eng \
-a credentialLevel=proxy -a authenticationMethod=sasl/DIGEST-MD5 \
-a bindTimeLimit=20 \
-a defaultSearchBase=dc=eng,dc=acme,dc=com \
-a "serviceSearchDescriptor=passwd:ou=people,dc=a1,dc=acme,dc=com?one" \
-a serviceAuthenticationMethod=pam_ldap:tls:simple \
-a defaultSearchScope=sub \
-a attributeMap=passwd:uid=employeeNumber \
-a objectclassMap=passwd:posixAccount=unixAccount \
-a followReferrals=false -a profileTTL=6000 \
-a preferredServerList=172.16.100.30 -a searchTimeLimit=30 \
-a "defaultServerList=172.16.200.1 172.16.100.1 192.168.5.6" > eng.ldif
```

退出状态

The following exit values are returned:

- 0 The command successfully executed.
- 1 An error occurred. An error message is output.
- 2 proxyDN and proxyPassword attributes are required, but they are not provided.

- 文件

/var/ldap/ldap_client_cred

/var/ldap/ldap_client_file

Contain the LDAP configuration of the client. These files are not to be modified manually. Their content is not guaranteed to be human readable. Use `ldapclient` to update them.

/etc/defaultdomain

System default domain name, matching the domain name of the data in the LDAP servers. See `defaultdomain(4)`.

/etc/nsswitch.conf

Configuration file for the name-service switch. See `nsswitch.conf(4)`.

/etc/nsswitch.ldap

Sample configuration file for the name-service switch configured with LDAP and files.
- 属性

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/nis
Interface Stability	Committed

另请参见

`chkey(1)`, `ldapadd(1)`, `ldapdelete(1)`, `ldaplist(1)`, `ldapmodify(1)`, `ldapmodrdn(1)`, `ldapsearch(1)`, `idsconfig(1M)`, `ldapaddent(1M)`, `ldap_cachemgr(1M)`, `defaultdomain(4)`, `nsswitch.conf(4)`, `resolv.conf(4)`, `attributes(5)`

注意

Currently StartTLS is not supported by `libldap.so.5`, therefore the port number provided refers to the port used during a TLS open, rather than the port used as part of a StartTLS sequence. To avoid timeout delays, mixed use of TLS and non-TLS authentication mechanisms is not recommended.

For example:

```
-h foo:1000 -a authenticationMethod=tls:simple
```

...or:

```
defaultServerList= foo:1000
authenticationMethod= tls:simple
```

The preceding refers to a raw TLS open on host `foo` port 1000, not an open, StartTLS sequence on an unsecured port 1000. If port 1000 is unsecured the connection will not be made.

As a second example, the following will incur a significant timeout delay while attempting the connection to `foo:636` with an unsecured bind.

```
defaultServerList= foo:636 foo:389
authenticationMethod= simple
```


引用名 ldmad – Logical Domains Agents daemon

用法概要 /usr/lib/ldoms/ldmad

描述 The ldmad daemon is part of the framework that enables Logical Domain agents to run on a Logical Domain. A Logical Domain agent is a component that interacts with the control domain to provide features or information.

ldmad is responsible for running agents on a Logical Domain. ldmad must be enabled to ensure the proper functionality of all features provided by the domain manager on the control domain. The daemon is started at boot time and has no configuration options.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/ldoms
Interface Stability	Unstable

另请参见 [svcs\(1\)](#), [svcadm\(1M\)](#), [syslog\(3C\)](#), [syslog.conf\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

《Oracle VM Server for SPARC 2.1 管理指南》

错误 ldmad uses [syslog\(3C\)](#) to report status and error messages. Error messages are logged with the LOG_ERR and LOG_NOTICE priorities. Informational messages are logged with the LOG_INFO priority. The default entries in the /etc/syslog.conf file specify that all ldmad error messages are written to the /var/adm/messages log.

附注 The ldmad service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/ldoms/agents:default

Administrative actions on this service, such as enabling, disabling, or requesting restart can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名 link, unlink – link and unlink files and directories

用法概要

```
/usr/sbin/link existing-file new-file
/usr/xpg4/bin/link existing-file new-file
/usr/sbin/unlink file
```

描述 The link and unlink commands link and unlink files and directories. Only super-users can use these commands on directories.

Use link to create a new file that points to an existing file. The *existing-file* and *new-file* operands specify the existing file and newly-created files. See OPERANDS.

Note that the ZFS file system does not support links between directories.

link and unlink directly invoke the [link\(2\)](#) and [unlink\(2\)](#) system calls, performing exactly what they are told to do and abandoning all error checking. This differs from the [ln\(1\)](#) command. See [ln\(1\)](#).

While linked files and directories can be removed using unlink, it is safer to use [rm\(1\)](#) and [rmdir\(1\)](#) instead. See [rm\(1\)](#) and [rmdir\(1\)](#).

/usr/xpg4/bin/link If the existing file being hard linked is itself a symbolic link, then the newly created file (*new-file*) will be a hard link to the file referenced by the symbolic link, not to the symbolic link object itself (*existing-file*).

操作数 The following operands are supported:

existing-file Specifies the name of the existing file to be linked.

file Specifies the name of the file to be unlinked.

new-file Specifies the name of newly created (linked) file.

环境变量 See [environ\(5\)](#) for descriptions of the following environment variables that affect the execution of link: LANG, LC_ALL, LC_CTYPE, LC_MESSAGES, and NLSPATH.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

/usr/xpg4/bin/link

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/xopen/xcu4
Interface Stability	Committed
Standard	See standards(5) .

另请参见

[ln\(1\)](#), [rm\(1\)](#), [link\(2\)](#), [unlink\(2\)](#), [attributes\(5\)](#), [environ\(5\)](#), [standards\(5\)](#)

引用名 llc2_loop – loopback diagnostics to test the driver, adapter and network.

用法概要

```
/usr/lib/llc2/llc2_loop2 [-v] ppa
/usr/lib/llc2/llc2_loop3 ppa sap frames
/usr/lib/llc2/llc2_loop3 ppa type frames
/usr/lib/llc2/llc2_loop4 [-v] ppa
```

描述

Loop 2 The loop2 test sends a NULL XID frame to the broadcast (all 1's) destination MAC address. The source SAP (Service Access Point) value used is 0x04 (SNA's SAP). Therefore, if SNA is running on the system, the loop2 test will fail. The destination SAP value is the NULL SAP (0x00). This test finds out who is listening and can receive frames sent out from a node. The verbose (-v) option displays the MAC address of responding nodes. All possible responders may not be displayed, since the loop2 test only waits for responses for 2 seconds, but during this time 50-200 nodes may be displayed. The most likely error is:

Unexpected DLPI primitive *x*, expected *y*.

where *x* = 5 and *y* = 6. From /usr/include/sys/dlpi.h, the expected return value from one of the DLPI primitives is 6 (DL_OK_ACK), but instead a 5 (DL_ERROR_ACK) was received. This can occur for two reasons:

- The loop2 command was issued to a non-existent PPA (Physical Point of Attachment).
- The SAP (0x04) is already in use (for example, the SNA subsystem is up).

Loop 3 The loop3 test sends 1,495 byte Unnumbered Information (UI) frames to the NULL (all 0's) destination MAC address. This should be used along with data capture either on the local node or another node on the same LAN to verify the transmission of data. The *ppa* argument specifies the adapter on which to run the test. The *ppa* is the relative physical position of the adapter and may be ascertained by viewing the adapter configuration (see [llc2_config\(1\)](#)). For Token Ring or Ethernet, specify an even *sap* value from 2 through 254, or, for Ethernet only, any *type* value from 1519 (0x05ef) through 65535 (0xffff). It is advised to pick a value that is easily recognized when the data capture output is viewed. *frames* is the decimal number of 1,495 bytes packets to transmit. The test will only display a message if a failure occurs.

Loop 4 The loop4 test sends a TEST frame (no information field) to the broadcast (all 1's) destination MAC address. The source SAP value used is 0x04 (SNA's SAP). Therefore, if SNA is running on the system, the loop4 test will fail. The destination SAP value is the NULL SAP (0x00). This test finds out who is listening and can receive frames sent out from a node. The verbose (-v) option displays the MAC address of responding nodes. All possible responders may not be displayed since the loop4 test only waits for responses for 2 seconds, but during this time 50-200 nodes may be displayed. The loop4 test displays information similar to the following example if other nodes are listening and respond (verbose mode):

-Attaching
-Binding

```
-Sending TEST
-Responders
  1-0000c0c12449
  2-08000e142990
  3-08000e142a51
  4-0000c0450044
  5-0000c0199e46
-Unbinding
-Detaching
5 nodes responding
```

The errors displayed are the same as for loop2.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/llc2

另请参见

[llc2_config\(1\)](#), [llc2\(4\)](#), [attributes\(5\)](#), [llc2\(7D\)](#)

附注

For information about how to start the service, see [llc2\(7D\)](#)

引用名

lldpadm – Link Layer Discovery Protocol administration tool

用法概要

```
lldpadm set-agentprop -p prop[+|-]=value[,...] lldp_agent
lldpadm reset-agentprop -p prop[,...] lldp_agent
lldpadm show-agentprop [[-c] -o field,...] -p prop[,...] [lldp_agent]

lldpadm set-tlvprop -p prop[+|-]=value[,...] tlv_name
lldpadm reset-tlvprop -p prop[,...] tlv_name
lldpadm show-tlvprop [[-c] -o field,...] -p prop[,...] [tlv_name]

lldpadm set-agenttlvprop -p prop[+|-]=value[,...] -a lldp_agent
                        tlv_name
lldpadm reset-agenttlvprop -p prop[,...] -a lldp_agent tlv_name
lldpadm show-agenttlvprop [[-c] -o field,...] -p prop[,...]
                        [-a lldp_agent] [tlv_name]

lldpadm show-agent [-c] [-s] [-v] -o field,...] [-l|-r] [lldp_agent]
```

描述

The `lldpadm` command is used to enable or disable a Link Layer Discovery Protocol (LLDP) agent on a physical datalink. `lldpadm` is also used to configure the behavior of an LLDP agent. The LLDP agent implements the LLDP protocol for a given physical datalink. LLDP is a one-way link layer protocol that allows an IEEE 802 LAN station to advertise the capabilities and current status of the system to other stations attached to the same LAN. The LLDP agent can also receive information about the capabilities and current status of the system associated with a remote station. LLDP agent can either be enabled for transmission only, for reception only, or for both.

Information to be exchanged is packed as a sequence of type, length, and value (TLVs), wherein the type field identifies the type of information, the length field indicates the length of the information field in octets, and the value field contains the information itself.

操作数

Each `lldpadm` subcommand operates on one of the following objects:

lldp_agent

An LLDP agent implements the LLDP protocol for a given physical datalink that is connected to IEEE 802 LAN. The only supported physical links are the ones of media type Ethernet. Thus, LLDP can be enabled on all the links displayed in `dladm show-phys` output that are of media type Ethernet. The name of the *lldp_agent* is the name of the datalink itself.

tlv_name

Name of the TLV whose value can be modified. The supported modifiable TLVs are:

- `syscapab` and `mgmtaddr`. These form Global TLVs that are common to all the LLDP agents on the system.
- `pfc` and `appln`. Per-LLDP agent TLVs.

子命令

lldpdm supports the following subcommands.

lldpdm set-agentprop|set-ap -p prop[+|-]=value[...] lldp_agent

Sets the value of one or more LLDP agent properties to the value specified. If the property takes multiple values then the value should be specified with a comma as the delimiter. The value is always made persistent and thus will be reapplied on system reboot or [lldpd\(1M\)](#) daemon restart. The list of properties supported and each property's possible values can be retrieved using show-agentprop subcommand.

-p prop[+|-]=value[...], --prop prop[+|-]=value[...]

A comma-separated list of properties to be set to the specified values. It also provides the following qualifiers to perform add and delete operations in addition to assignment.

- +** Adds the given value to the current list of value(s).
- Removes the given value from the current list of value(s).
- =** Makes a new assignment and removes all the current value(s).

See EXAMPLES for more information on how to use the qualifiers.

lldpdm reset-agentprop|reset-ap -p prop[...] lldp_agent

Resets one or more properties to their default values. The default values for properties can be retrieved using show-agentprop subcommand.

-p prop[...], --prop prop[...]

A comma-separated list of properties to reset.

lldpdm show-agentprop|show-ap [[-c] -o field,...] -p prop[...] [lldp_agent]

Show the current value of one or more properties, either for all of the LLDP agents or for the specified LLDP agent. Several properties of interest can be retrieved at a time by providing comma-separated property names to -p option. If the -p option is not specified, all available properties are displayed.

-o field[...], --output field[...]

A case-insensitive, comma-separated list of output fields to display. The field name must be one of the fields listed below, or the special value **all** to display all fields. For each LLDP agent, the following fields can be displayed:

AGENT

The name of the LLDP agent.

PROPERTY

The name of the property.

PERM

The read/write permissions of the property. The value shown will be **r** (read only), **w** (write only) or **rw** (read/write).

VALUE

The current value of the property. If the value is not set, it is shown as -. If it is unknown, the value is shown as ?.

DEFAULT

The default value of the property. If the property has no default value, - is shown.

POSSIBLE

A comma-separated list of the values the property can have. If the values span a numeric range, *min - max* might be shown as shorthand. If the possible values are unknown or unbounded, - is shown.

-c, --parsable

Display using a stable machine-parsable format. The -o option is required with this option. See “Parsable Output Format”, below.

-p prop[,...], --prop prop[,...]

A comma-separated list of properties to display.

For the supported list of agent properties, see “Agent Properties” section below.

lldpdm set-tlvprop|set-tp -p prop[+|-]=value[,...] tlv_name

Sets the value of one or more TLV properties to the value specified. If the property takes multiple values, the value should be specified with a comma as the delimiter. The value is always persisted and will be reapplied on system reboot or **lldpd(1M)** daemon restart. The list of properties supported and each property's possible values can be retrieved using **show-tlvprop** subcommand.

注 – The TLVs modified using this subcommand apply to all the LLDP agents running on the system.

-p prop[+|-]=value[,...], --prop prop[+|-]=value[,...]

See the description of this option under the **set-agentprop** subcommand, above.

lldpdm reset-tlvprop|reset-tp -p prop[,...] tlv_name

Resets one or more properties to their default values. The default values for properties can be retrieved using **show-tlvprop** subcommand.

-p prop[,...], --prop prop[,...]

A comma-separated list of properties to reset.

lldpdm show-tlvprop|show-tp [[-c] -o field,...] -p prop[,...] [lldp_agent]

Show the current value of one or more properties, either for all the TLVs or for a specified TLV. Several properties of interest can be retrieved at a time by providing comma-separated property names to -p option. If the -p option is not specified, all available properties are displayed.

`-o field[,...], --output field[,...]`

A case-insensitive, comma-separated list of output fields to display. The field name must be one of the fields listed below, or the special value `all` to display all fields. For each TLV, the following fields can be displayed:

TLVNAME

The name of the TLV.

PROPERTY

The name of the property.

PERM

The read/write permissions of the property. The value shown will be `r` (read only), `w` (write only) or `rw` (read/write).

VALUE

The current value of the property. If the value is not set, it is shown as `--`. If it is unknown, the value is shown as `?`.

DEFAULT

The default value of the property. If the property has no default value, `--` is shown.

POSSIBLE

A comma-separated list of the values the property can have. If the values span a numeric range, *min* - *max* might be shown as shorthand. If the possible values are unknown or unbounded, `--` is shown.

`-c, --parsable`

Display using a stable machine-parsable format. The `-o` option is required with this option. See “Parsable Output Format”, below.

`-p prop[,...], --prop prop[,...]`

A comma-separated list of properties to display.

`lldpdm set-agent tlvprop | set-atp -p prop[+|-]=value[,...] -a lldp_agent tlv_name`

Sets the value of one or more TLV properties to the value specified. The `-a` option is mandatory and identifies the agent for which the TLV property needs to be set. In this way, the TLV property modification is reflected only on the specified agent. If the property takes multiple values then the value should be specified with a comma as the delimiter. The value is always made persistent and will be reapplied on system reboot or [lldpd\(1M\)](#) daemon restart. The list of properties supported and each property's possible values can be retrieved using `show-agent tlvprop` subcommand.

`-p prop[+|-]=value[,...], --prop prop[+|-]=value[,...]`

See the description of this option under the `set-agentprop` subcommand, above.

`-a lldp_agent, --agent=lldp_agent`

The name of the LLDP agent for which TLV properties need to be displayed.

```
lldpdm reset-agent tlvprop | reset-atp -p prop[,...]=value[,...] -a lldp_agent tlv_name
```

Resets one or more properties to their default values. The `-a` option is mandatory and identifies the name of the agent for which the TLV property needs to be reset. The default values for properties can be retrieved using `show-agent tlvprop` subcommand.

```
-p prop[,...], --prop prop[,...]
```

A comma-separated list of properties to display.

```
-a lldp_agent, --agent=lldp_agent
```

The name of the LLDP agent for which TLV properties need to be displayed.

```
lldpdm show-agent tlvprop | show-atp [[-c] -o field,...] -p prop[,...] [-a lldp_agent]
[tlv_name]
```

Show the current value of one or more properties, either for all of the TLVs or for a specified TLV. If a single LLDP agent is not specified (using `-a`), TLV properties for all LLDP agents are displayed. Several properties of interest can be retrieved at a time by providing comma-separated property names to the `-p` option. If the `-p` option is not specified, all available properties are displayed.

```
-o field[,...], --output field[,...]
```

A case-insensitive, comma-separated list of output fields to display. The field name must be one of the fields listed below, or the special value `all` to display all fields. For each LLDP agent, the following fields can be displayed:

AGENT

The name of the LLDP agent.

TLVNAME

The name of the TLV.

PROPERTY

The name of the property.

PERM

The read/write permissions of the property. The value shown will be `r` (read only), `w` (write only) or `rw` (read/write).

VALUE

The current value of the property. If the value is not set, it is shown as `--`. If it is unknown, the value is shown as `?`.

DEFAULT

The default value of the property. If the property has no default value, `--` is shown.

POSSIBLE

A comma-separated list of the values the property can have. If the values span a numeric range, `min - max` might be shown as shorthand. If the possible values are unknown or unbounded, `--` is shown.

-c, --parsable

Display using a stable machine-parsable format. The -o option is required with this option. See “Parsable Output Format”, below.

-p *prop*[,...], --prop *prop*[,...]

A comma-separated list of properties to display.

-a *lldp_agent*, --agent=*lldp_agent*

The name of the LLDP agent for which TLV properties need to be displayed.

For the supported list of TLV properties that apply on a per-LLDP agent basis, see the “Per-LLDP agent TLV properties” section, below.

`lldpdm show-agent [-c] [-s] [-v] -o field,... [-l | -r] [lldp_agent]`

Show the information advertised by the specified LLDP agent or information advertised by the adjacent neighbors to the specified LLDP agent. If no LLDP agent is specified, then the local or remote information will be displayed for all the LLDP agents. The information is displayed as a multi-line output, with each line containing information about a single TLV in the following format:

<Name of the TLV> : *<TLV Information expressed as a string>*

The following lines can be displayed:

```

Agent:
Chassis ID Subtype:
Chassis ID:
Port ID Subtype:
Port ID:
Port Description:
Time to Live:
System Name:
System Description:
Supported Capabilities:
Enabled Capabilities:
Management Address:
Maximum Frame Size:
Port VLAN ID:
VLAN Name/ID:
VNIC PortID/VLAN ID:
Aggregation Information:
PFC Willing:
PFC Cap:
PFC MBC:
PFC Enable:
Application(s)(ID/Sel/Pri):
Next Packet Transmission: [displayed only when -l is used]
Information Valid Until: [displayed only when -r is used]
```

-c, --parsable

Display using a stable machine-parsable format. The -o option is required with this option. See “Parsable Output Format”, below.

-l, --local

Displays information advertised by the local LLDP agent. This option is mutually exclusive of the -r option.

-r, --remote

Displays information advertised by the adjacent neighbors. This option is mutually exclusive of the -l option.

-s

Show the statistics for the specified LLDP agents or for all the LLDP agents on the system. Every LLDP agent maintains statistical counters that are used to count significant events in the transmit and receive state machines. These counters are defined to be 32-bit unsigned integers.

-o *field*[,...], --output *field*[,...]

A case-insensitive, comma-separated list of output fields to display. The field name must be one of the fields listed below, or the special value `all` to display all fields. For each TLV, the following fields can be displayed:

AGENT

The name of the LLDP agent.

IFRAMES

Count of all LLDP frames received by AGENT.

TIERR

Count of all LLDPDUs received with one or more detectable errors.

IDISCARD

Count of all LLDPDUs received and then discarded for any of the following reasons:

- incorrectly formed LLDPDUs with respect to the first three mandatory TLVs
- insufficient space to store the incoming LLDPDU

OFRAMES

Count of all LLDP frames transmitted by AGENT.

OLENERR

Count of times the outgoing LLDPDU exceeded the length restrictions of 1500 bytes.

TLVDISCARD

Count of all TLVs received and then discarded because the TLVs did not adhere to the TLV usage rules as specified by the standard.

TLVUNRECOG

Count of all TLVs received that are not recognized by the LLDP agent.

AGEOUT

Count of the times that a neighbor's information has been deleted because of aging.

-v, --verbose

Displays detailed information.

Parsable Output Format

The `lldpadm show` subcommands have a `-c` option that displays output in a machine-parsable format. The output format is one or more lines of colon (:) delimited fields. The fields displayed are specific to the subcommand used and are listed under the entry for the `-o` option for a given subcommand. Output includes only those fields requested by means of the `--o` option, in the order requested. Note that the `-o all` option, which displays all the fields for a given subcommand, cannot be used with the parsable output option.

When you request multiple fields, any literal colon characters are escaped by a backslash (\) before being output. Similarly, literal backslash characters are also escaped with a backslash. This escape format is parsable by using shell [read\(1\)](#) functions with the environment variable set as `IFS=`. Note that escaping is not done when you request only a single field.

Agent Properties

The following LLDP agent properties are supported:

mode

Configures the operation mode of the LLDP agent. Possible values are:

txonly

Enables LLDP for transmission only.

rxonly

Enables LLDP for receiving only.

both

Enables LLDP for both transmission and receiving.

disable

Disables LLDP on the LLDP agent.

Every LLDP packet (LLDPDU) transmitted by an LLDP agent contains multiple TLVs. The following four TLVs are mandatory and therefore included in all the LLDPDUs transmitted by an agent configured in the `txonly` or `both` mode:

CHASSIS ID

The value transmitted in the Chassis ID TLV is:

- subtype = 7 (locally assigned)
- Chassis ID = `hostid(1)`

PORT ID

The value transmitted in the Port ID TLV is:

- subtype = 3 (MAC address)
- Port ID = primary MAC address of the agent

TTL

The duration for which this packet is valid. The default value is 30 seconds.

End of PDU TLV

End of PDU indicator.

Optional TLVs that can be advertised are configured using the following properties:

basic-tlv

Configures the Basic Management TLVs that should be advertised by the LLDP agent. The possible values are:

portdesc

Alphanumeric string that identifies the datalink. Value set to the linkname.

sysname

Alphanumeric string that identifies the system. Value set to the output of 'uname -n'.

sysdesc

Alphanumeric string that describes the system. Value set to the output of 'uname -a'.

syscapab

Indicates the systems supported and enabled capabilities.

mgmtaddr

Indicates the IP address of the system that can be used by network management.

dot1-tlv

Configures the IEEE 802.1 Organizationally Specific TLVs that should be advertised by the LLDP agent. The possible values are:

vlanname

Indicates the names and IDs of all the VLANs configured on the datalink.

pvid

Indicates the default VLAN ID associated with the given datalink. It corresponds to the `default_tag` datalink property that is managed by means of the [dladm\(1M\)](#) utility.

linkaggr

Indicates whether underlying datalink is in an aggregation or is capable of being part of an aggregation.

pfc

Indicates whether underlying datalink supports PFC (Priority Flow Control) and the priorities for which the PFC pause frame is enabled. Also indicates whether the local endpoint is willing to negotiate the PFC configuration.

appln

Indicates the priority that will be used by an application.

dot3-tlv

Configures the IEEE 802.3 Organizationally Specific TLVs that should be advertised by the LLDP agent. The possible values are:

max-framesize

Indicates the maximum supported frame size for the underlying datalink.

virt-tlv

Configures the Solaris's Virtualization Specific TLVs that should be advertised by the LLDP agent. The possible values are:

vnic

Indicates the MAC address of the Virtual NIC created on top of the underlying physical link. Also indicates any VLAN id associated with the VNIC. See [dladm\(1M\)](#) for more information on VNIC.

Global TLVs and Their Properties

The following Global TLV properties are supported:

syscapab (TLV name)

This property can one of the following values:

supported

Indicates the supported capabilities on the system. The default supported capabilities are: bridge, router, and station.

enabled

Indicates the enabled capabilities on the system. The enabled capabilities must be a subset of the supported capabilities.

mgmtaddr (TLV name)

This property can have the following value:

ipaddr

The IP address(es), either IPv4 or IPv6, associated with the local LLDP agent that will be used to reach higher layer entities to assist discovery by network management.

Per-Agent TLVs and Their Properties

The following Agent TLV properties are supported:

appln (TLV name)

This property can have the following value:

apt

Configures the Application Priority Table for an Application TLV. One can add or remove entries from this table using the + and - qualifiers. Each entry in the table indicates the application and the priority that will be used for that application. Its value is of the form:

id/selector/priority

The meaning of the *id* is determined by the *selector* field. The *selector* field can be any one of the following:

- 1 — *id* indicates an Ethertype (an L2 protocol), therefore *id*'s value should be greater than 1536
- 2 — *id* indicates a port number over TCP or SCTP
- 3 — *id* indicates a port number over UDP or DCCP
- 4 — *id* indicates a port number over TCP, SCTP, UDP, or DCCP

The *priority* indicates the priority value (0-7) that will be used for given application.

Authorizations

The following subcommands require solaris.network.lldp authorization:

- set-agentprop
- reset-agentprop
- set-tlvprop
- reset-tlvprop
- set-agenttlvprop
- reset-agenttlvprop

The various show-* subcommands do not need any authorization.

示例

示例 1 Enabling LLDP Protocol on an LLDP Agent

The following command enables the LLDP protocol on an LLDP agent for both transmission and reception of LLDPDUs.

```
# lldpadm set-agentprop -p mode=both net0
```

示例 2 Disabling LLDP Protocol on an LLDP Agent

The following command disables the LLDP protocol on an LLDP agent.

```
# lldpadm set-agentprop -p mode=disable net0
```

示例 3 Configuring TLVs

The following command configures transmission of the Port Description and System Name TLV.

```
# lldpadm set-agentprop -p basic-tlv=portdesc,sysname net0
```

The following command configures transmission of a VLAN Name TLV and a Link Aggregation TLV.

```
# lldpadm set-agentprop -p dot1-tlv=vlanname net0
# lldpadm set-agentprop -p dot1-tlv+=linkaggr net0
```

The following command configures transmission of all dot3-tlvs.

```
# lldpadm set-agentprop -p dot3-tlv=all net0
```

All the above lldpadm invocations can be combined into the following, single invocation.

示例 3 Configuring TLVs (续)

```
# lldpadmin set-agentprop -p basic-tlv=portdesc,sysname,\
dot1-tlv=vlanname,linkaggr,dot3-tlv=all net0
```

示例 4 Disabling Transmission

The following command disables the transmission of all dot1-tlvs out of an LLDP agent.

```
# lldpadmin set-agentprop -p dot1-tlv=none net0
```

The following command is equivalent to the preceding.

```
# lldpadmin reset-agentprop -p dot1-tlv net0
```

示例 5 Configuring Enabled Capabilities

The following command configures the enabled capabilities on a system.

```
# lldpadmin set-tlvprop -p enabled=router syscapab
```

With this configuration, when an LLDP agent is enabled for advertising a System Capabilities TLV, the adjacent neighbors would learn of the local system's capabilities.

示例 6 Configuring a Management Address for Subsequent Advertising

The following command configures the management address that will be advertised by means of the Management Address TLV.

```
# lldpadmin set-tlvprop -p ipaddr=192.168.1.2 mgmtaddr
```

Note that this address would be identified as an address associated with the local LLDP agent that will be used to reach higher layer entities to assist discovery by network management.

示例 7 Configuring an Application TLV

The following sequence of commands configures the application TLV to advertise the priority that will be used by FCoE.

```
# lldpadmin set-agenttlvprop -p apt=8906/1/4 -a net0 appln
# lldpadmin show-agenttlvprop -a net0 appln
```

AGENT	TLVNAME	PROPERTY	PERM	VALUE	DEFAULT	POSSIBLE
net0	appln	apt	rw	8906/1/4	--	--

示例 8 Show Local Information Advertised by LLDP Agent

The following commands show, respectively, brief and detailed local information advertised by an LLDP agent.

Brief information:

```
# lldpadmin show-agent -l net0
```

AGENT	CHASSISID	PORTID

示例 8 Show Local Information Advertised by LLDP Agent (续)

net0 004bb87f 00:14:4f:01:77:5d

Detailed information:

```
# lldpdm show-agent -lv net0
    Agent: net0
    Chassis ID Subtype: Local(7)
    Chassis ID: 004bb87f
    Port ID Subtype: MacAddress(3)
    Port ID: 00:14:4f:01:77:5d
    Port Description: net0
    Time to Live: 81 (seconds)
    System Name: hosta.example.com
    System Description: SunOS 5.11 dcb-clone-x-01-19-11 i86pc
    Supported Capabilities: bridge, router, station
    Enabled Capabilities: router
    Management Address: 192.168.1.2
    Maximum Frame Size: 3000
    Port VLAN ID: --
    VLAN Name/ID: vlan25/25
    VNIC PortID/VLAN ID : --
    Aggregation Information: Capable, Not Aggregated
    PFC Willing: --
    PFC Cap: --
    PFC MBC: --
    PFC Enable: --
    Application(s)(ID/Sel/Pri): --
    Next Packet Transmission: 18 (seconds)
```

示例 9 Show Remote Information about Adjacent Devices for an LLDP Agent

The following commands show, respectively, brief and detailed remote information about adjacent devices for a given LLDP agent.

Brief information:

```
# lldpdm show-agent -r net0
AGENT      SYSNAME      CHASSISID      PORTID
net0       hostb       0083b390      00:14:4f:01:59:ab
```

Detailed information:

```
# lldpdm show-agent -rv net0
    Agent: net0
    Chassis ID Subtype: Local(7)
    Chassis ID: 0083b390
    Port ID Subtype: MacAddress(3)
    Port ID: 00:14:4f:01:59:ab
```

示例 9 Show Remote Information about Adjacent Devices for an LLDP Agent (续)

```
Port Description: net0
Time to Live: 121 (seconds)
System Name: hostb.example.com
System Description: SunOS 5.11 dcb-clone-x-01-19-11 i86pc
Supported Capabilities: bridge, router, station
Enabled Capabilities: router
Management Address: 192.168.1.3
Maximum Frame Size: 3000
Port VLAN ID: --
VLAN Name/ID: vlan25/25
VNIC PortID/VLAN ID : 02:08:20:72:71:31
Aggregation Information: Capable, Not Aggregated
PFC Willing: --
PFC Cap: --
PFC MBC: --
PFC Enable: --
Application(s) (ID/Sel/Pri): --
Information Valid Until: 117 (seconds)
```

示例 10 Show LLDP Agent Statistics
The following command displays LLDP agent statistics.

```
# lldpadm show-agent -s net0
AGENT IFRAMES IERR IDISCARD OFRAMES OLENERR TLVDISCARD TLVUNRECOG AGEOUT
net0      44      0      0      57      0      0      0      0
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

/sbin

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/lldp
Interface Stability	Committed

另请参见

[hostid\(1\)](#), [read\(1\)](#), [uname\(1\)](#), [dladm\(1M\)](#), [lldpd\(1M\)](#), [svcadm\(1M\)](#), [attributes\(5\)](#)

IEEE Std 802.1AB-2009, IEEE Standard for Local and Metropolitan Area Networks: Station and Media Access Control Connectivity Discovery

IEEE Draft 802.1Qbb, Virtual Bridged Local Area Networks - Amendment : Priority-based Flow Control

IEEE Draft 802.1Qaz, Virtual Bridged Local Area Networks - Amendment XX: Enhanced Transmission Selection for Bandwidth Sharing between Traffic Classes

附注 The [lldpd\(1M\)](#) daemon that implements the LLDP protocol must be first enabled before using the `lldpadm` command. The `lldpd` daemon is controlled through the service management facility (SMF) service instance:

```
svc:/network/lldp:default
```

Use [svcadm\(1M\)](#) to enable this service.

引用名

lldpd – Link Layer Discovery Protocol daemon

用法概要

```
/usr/lib/lldpd
svc:/network/lldp:default
```

描述

lldpd is a system daemon that implements the Link Layer Discovery Protocol as specified in IEEE Std. 802.1AB. The daemon also manages LLDP agents on physical datalinks. Please see [lldpadm\(1M\)](#) for more information on administering LLDP on LLDP agents.

The lldpd daemon is controlled through the service management facility (SMF) service instance:

```
svc:/network/lldp:default
```

This means, [svcadm\(1M\)](#) must be used to start, stop, restart, and refresh this daemon. This daemon is disabled by default. To enable LLDP on any of the physical datalinks, this daemon should be running.

The lldpd daemon is a project private interface and has no user-accessible options.

示例

示例 1 Enabling the Daemon

The following command enables the lldpd daemon.

```
# svcadm enable network/lldp
# svcs network/lldp
STATE          STIME      FMRI
online         Jan_26     svc:/network/lldp:default
```

示例 2 Disabling the Daemon

The following command disables the lldpd daemon.

```
# svcadm disable network/lldp
# svcs network/lldp
STATE          STIME      FMRI
disabled       5:11:35    svc:/network/lldp:default
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/lldp
Interface Stability	Project Private

另请参见

[lldpadm\(1M\)](#), [svcadm\(1M\)](#), [attributes\(5\)](#)

IEEE Std 802.1AB-2009, IEEE Standard for Local and Metropolitan Area Networks: Station and Media Access Control Connectivity Discovery

引用名 lms – 允许应用程序访问 Intel 主动管理技术

用法概要 lms

描述 本地管理服务 lms 允许应用程序使用本地 HECI 接口访问 Intel 主动管理技术 (Intel Active Management Technology, Intel AMT) ME (Management Engine, 管理引擎)。LMS 依赖于 HECI 驱动程序。

要使用 lms，请确保已安装 Intel AMT 管理接口驱动程序，并已将其连接至 ME。

lms 将作为一个守护进程运行。来自该服务的消息被发送至 syslog。LMS 消息被标记为以 "LMS" 为来源。

有关如何使用 SMF 启用 LMS 服务的示例，请参见“示例”部分。

示例 示例 1 使用 SMF 启用 LMS 服务

以下示例使用 SMF 启用 LMS 服务：

```
svcadm enable network/lms
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
体系结构	x86
可用性	system/management/intel-amt
接口稳定性	Volatile (可变)

另请参见 [attributes\(5\)](#)、[smf\(5\)](#)

引用名 localectr – customize and build new locales

用法概要

```
/usr/bin/localectr

/usr/bin/localectr -l locale1,locale2... [-d destination_path]
               [-c compiler_path] [-r 0 | 1 | 2] [-i pkginfo_template]
               [-p pkgname_prefix]

/usr/bin/localectr -h

/usr/bin/localectr -q

/usr/bin/localectr -V
```

描述

The `localectr` utility allows new locales to be customized and built. The output of `localectr` is an installable package containing the compiled shared object binary which contains the locale data information as well as a number of other supporting files that are required to have a fully working locale on the system.

Once generated, the package can be added to the system by using the [pkgadd\(1M\)](#) command and removed with [pkgrm\(1M\)](#).

Depending on your default system login, you might have to reset your user environment after you add a locale. If `dtlogin(1X)` (for the CDE user environment) is the default system login, you need to restart `dtlogin`. No action is required if `gdm(1)` (for the Gnome user environment) is the default login.

There are two interfaces to `localectr`, command line (CLI) and graphical user interface (GUI). To customize the locale data, you must use the GUI. To create locales with standard locale data according to Unicode's Common Locale Data Repository (CLDR), the CLI is sufficient.

With the CLI it is also possible to generate several locales in a single step, with a separate package being generated for each locale. With the GUI, a single locale is processed at a time.

To launch the GUI, run the `localectr` command with no options. To run from the CLI, use the appropriate options as described below.

`localectr` uses the [localedef\(1\)](#) utility to build the locale data binary shared object. Therefore, access to a C compiler is required in order to run `localectr` successfully, as this is also required by `localedef`.

`localectr` is mainly concerned with locale data. However, in order to create a fully working locale on the system with `localectr`, many features, such as fonts, translations and input methods, are also required. Depending on what locales are already installed on the system, the relevant features might or might not be present on the system. `localectr` bundles locale data for the latest set of locales available in Unicode's CLDR. A user can also create a locale not available in CLDR by supplying her own data in the `localectr` GUI.

选项

The following options are supported:

- c *compiler_path*
Specify the path to the C compiler that is used to compile the locale data into a shared object binary. Not required if the compiler is already in the user's PATH.
- d *destination_path*
Specify the path to the directory where the created package is to be stored.
- h
Displays the usage message.
- i *template_file*
The full path to user defined `pkginfo(4)` template file.
- l *locale1,locale2...*
Specify a comma-separated list of locale(s) to generate. Locale names are in the form: *locale.codeset@variant*, where *codeset* and *variant* are optional. The default and only allowed *codeset* is UTF-8. The default *variant* is `localectr`.
- p
The package name prefix
- q
Queries `localectr` for a complete list of locales for which locale data is defined in `localectr`. When `localectr` is run from the CLI, the locale(s) specified with the -l option must be on this list in order for an installable locale package to be generated.
- r
Specify the range of Unicode characters for which locale data rules in the LC_CTYPE and LC_COLLATE categories are to be generated. There are three valid options:
 - 0
Locale data rules are restricted to the exemplar or to commonly used characters of the locale in question.
 - 1
Locale data rules are restricted to the Unicode plane 0 characters, whose codepoints fall in the range `u0000-uFFFF`.
 - 2
Locale data rules are generated for all codepoints defined in the latest version of Unicode that is supported by the system on which `localectr` is being run.
- V
Shows the version of this software.

示例

示例 1 Launching the GUI

The following example launches the `localectr` GUI.

```
example% localectr
```


示例 2 Generating Locale for Afrikaans (South Africa) with Default Locale Data

The following example generates a package in the specified destination directory, which can be used to install the Afrikaans (af_ZA.UTF-8) locale on the system. The package name is composed of a prefix followed by the hyphen separated ISO-639 language code, the ISO-3166 country code, the locale encoding and an optional user-defined tag. The resulting package can then be added to the system using [pkgadd\(1M\)](#).

```
example% localectr -l af_ZA -d /tmp
```

示例 3 Generating Several South Africa Locales with Full Unicode Range of Characters

The following example generates an installable package for each of the specified locales.

```
example% localectr -l af_ZA,en_ZA,xh_ZA,zu_ZA -d /tmp -r 2
```

示例 4 Generating the Irish Locale with a User-Specified Tag

The following example will generate a locale whose full name is ga_IE.UTF-8@mycompanyname.

```
example% localectr -l ga_IE@mycompanyname -d /tmp
```

退出状态

The following exit codes are returned:

0
Successful completion

>0
An error occurred.

文件

/usr/bin/localectr
Wrapper script that launches locale creator.

/usr/lib/localectr
Jar files, scripts, and locale data repository needed to run the application.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWlocalecreator
Interface Stability	See below.

Interface stability is Committed for command-line options and is Uncommitted for other interfaces.

另请参见

[locale\(1\)](#), [localedef\(1\)](#), [pkgadd\(1M\)](#), [pkgrm\(1M\)](#), [pkginfo\(4\)](#), [attributes\(5\)](#)

[dtlogin\(1X\)](#), [gdm\(1\)](#)(these are not SunOS man pages)

International Language Environments Guide

引用名	locator – location indicator control
用法概要	<code>/usr/sbin/locator [-f -n]</code>
描述	The locator command sets or queries the state of the system locator if such a device exists. Without options, the locator command reports the current state of the system. The privileges required to use this command are hardware dependent. Typically, only the super user can get or set a locator.
选项	The following options are supported: -f Turns the locator off. -n Turns the locator on.
示例	示例 1 Using the locator Command on a Platform Which Has a System Locator LED When issued on a platform which has a system locator LED, the following command turns the locator on: <pre># locator -n # locator The 'system' locator is on</pre> 示例 2 Using the locator Command on a Platform Which Does Not Have a System Locator LED When issued on a platform which does not have a system locator LED, the following command attempts to turn the locator on. The command returns an error message. <pre># locator -n 'system' locator not found</pre>
退出状态	The following exit values are returned: 0 Successful completion. 1 Invalid command line input. 2 The requested operation failed.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/picl

另请参见 [attributes\(5\)](#)

引用名 lockd – network lock daemon

用法概要 `/usr/lib/nfs/lockd [-g graceperiod] [-l listen_min_backlog]
[-t timeout] [nthreads]`

描述 The lockd utility is part of the NFS lock manager, which supports record locking operations on NFS files. See [fcntl\(2\)](#) and [lockf\(3C\)](#). The lock manager provides the following two functions:

- It forwards [fcntl\(2\)](#) locking requests for NFS mounted file systems to the lock manager on the NFS server.
- It generates local file locking operations in response to requests forwarded from lock managers running on NFS client machines.

State information kept by the lock manager about these locking requests can be lost if the lockd is killed or the operating system is rebooted. Some of this information can be recovered as follows. When the server lock manager restarts, it waits for a grace period for all client-site lock managers to submit reclaim requests. Client-site lock managers, on the other hand, are notified by the status monitor daemon, [statd\(1M\)](#), of the restart and promptly resubmit previously granted lock requests. If the lock daemon fails to secure a previously granted lock at the server site, then it sends SIGLOST to a process.

Administrators can make changes to the startup parameters for lockd by logging in as root and using the [sharectl\(1M\)](#) command.

SMF Management The lockd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/nfs/nlockmgr
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

If it is disabled, it will be enabled by [mount_nfs\(1M\)](#), [share_nfs\(1M\)](#), and [automountd\(1M\)](#) unless its application/auto_enable property is set to false.

The [sharectl\(1M\)](#) command is used to manipulate the startup SMF parameters for lockd. Currently supported parameters are as follows:

`lockd_listen_backlog=num`

Set connection queue length for lockd over a connection-oriented transport. The default and minimum value is 32. Equivalent to -l option.

`lockd_servers=num`

Maximum number of concurrent lockd requests. The default is 1024. Equivalent to the *nthreads* operand.

`lockd_retransmit_timeout=num`

Retransmit timeout, in seconds, before lockd retries. The default is 5. Equivalent to `-t` option.

`grace_period=num`

Grace period, in seconds, that all clients (both NLM and NFSv4) have to reclaim locks after a server reboot. This parameter also controls the NFSv4 lease interval. The default is 90. Equivalent to `-g` option.

See EXAMPLES, below.

选项

The following options are supported:

`-g graceperiod`

Deprecated in favor of `grace_period`. Specify the number of seconds that all clients (both NLM and NFSv4) have to reclaim locks after the server reboots. It also controls the NFSv4 lease interval. This option is equivalent to the `grace_period` property described above.

`-l listen_min_backlog`

Specify the listener backlog (`listen_min_backlog`). `listen_min_backlog` is the number connect requests that are queued and waiting to be processed before new connect requests start to get dropped. Equivalent of the `lockd_listen_backlog` property described above.

`-t timeout`

Specify the number of seconds to wait before retransmitting a lock request to the remote server. The default value is 5 seconds. Equivalent of the `lockd_retransmit_timeout` property described above.

操作数

`nthreads`

Specify the maximum number of concurrent threads that the server can handle. This concurrency is achieved by up to `nthreads` threads created as needed in the kernel. `nthreads` should be based on the load expected on this server. If `nthreads` is not specified, the maximum number of concurrent threads will default to 1024. Equivalent of the `lockd_servers` property described above.

示例

示例 1 Setting a lockd Property

The following command sets `lockd_listen_backlog` to a new value:

```
# sharectl set -p lockd_listen_backlog=40 nfs
```

The `lockd_listen_backlog` and other lockd properties are described under NOTES, below.

示例 2 Getting a lockd Property Value

The following command retrieves the value of the `lockd_listen_backlog` property.

```
% sharectl get -p lockd_listen_backlog nfs
lockd_listen_backlog=40
```

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/nfs

另请参见 [svcs\(1\)](#), [automountd\(1M\)](#), [clear_locks\(1M\)](#), [mount_nfs\(1M\)](#), [share\(1M\)](#), [share_nfs\(1M\)](#), [sharectl\(1M\)](#), [statd\(1M\)](#), [svcadm\(1M\)](#), [fcntl\(2\)](#), [lockf\(3C\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The lockd daemon might not exist in a future release of Solaris.

引用名	lockfs – change or report file system locks
用法概要	<code>/usr/sbin/lockfs [-adefhnuw] [-c <i>string</i>] [<i>file-system</i>]</code> ...
描述	lockfs is used to change and report the status of file system locks. lockfs reports the lock status and unlocks the file systems that were improperly left locked. Using lockfs to lock a file system is discouraged because this requires extensive knowledge of SunOS internals to be used effectively and correctly. When invoked with no arguments, lockfs lists the UFS file systems that are locked. If <i>file-system</i> is not specified, and -a is specified, lockfs is run on all mounted, UFS type file systems.
选项	The options are mutually exclusive: wndheu<code>f</code>. If you do specify more than one of these options on a lockfs command line, the utility does not protest and invokes only the last option specified. In particular, you cannot specify a flush (-f) and a lock (for example, -w) on the same command line. However, all locking operations implicitly perform a flush, so the -f is superfluous when specifying a lock. You must be super-user to use any of the following options, with the exception of -a, -f and -v. The following options are supported. -a Apply command to all mounted, UFS type file systems. <i>file-system</i> is ignored when -a is specified. -c <i>string</i> Accept a string that is passed as the comment field. The -c only takes affect when the lock is being set using the -d, -h, -n, -u, or -w options. -d Delete-lock (dlock) the specified <i>file-system</i>. dlock suspends access that could remove directory entries. -e Error-lock (elock) the specified <i>file-system</i>. elock blocks all local access to the locked file system and returns EWOULDBLOCK on all remote access. File systems are elocked by UFS on detection of internal inconsistency. They may only be unlocked after successful repair by fsck, which is usually done automatically (see mount_ufs(1M)). elocked file systems can be unmounted. -f Force a synchronous flush of all data that is dirty at the time fsflush is run to its backing store for the named file system (or for all file systems.)

It is a more reliable method than using [sync\(1M\)](#) because it does not return until all possible data has been pushed. In the case of UFS filesystems with logging enabled, the log is also rolled before returning. Additional data can be modified by the time `fsflush` exits, so using one of the locking options is more likely to be of general use.

- h
Hard-lock (`hlock`) the specified *file-system*. `hlock` returns an error on every access to the locked file system, and cannot be unlocked. `hlocked` file systems can be unmounted.
- n
Name-lock (`nlock`) the specified *file-system*. `nlock` suspends accesses that could change or remove existing directories entries.
- u
Unlock (`unlock`) the specified *file-system*. `unlock` awakens suspended accesses.
- v
Enable verbose output.
- w
Write-lock (`wlock`) the specified *file-system*. `wlock` suspends writes that would modify the file system. Access times are not kept while a file system is write-locked.

操作数 The following operands are supported.

file-system
A list of path names separated by whitespace. Note that *file-system* can be a directory rather than the specific name of a file system, such as `/` or `/usr`. For example, if you specify `/export/home` as an argument to a `lockfs` command and `/export/home` is mounted on the root (`/`) file system, the `lockfs` command will take effect on the root file system.

用法 See [largefile\(5\)](#) for the description of the behavior of `lockfs` when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).

示例 示例1 Using `lockfs -a`

In the following examples, *filesystem* is the pathname of the mounted-on directory (mount point). Locktype is one of “write,” “name,” “delete,” “hard,” or “unlock”. When enclosed in parenthesis, the lock is being set. Comment is a string set by the process that last issued a lock command.

The following example shows the `lockfs` output when only the `-a` option is specified.

example# **`/usr/sbin/lockfs -a`**

Filesystem	Locktype	Comment
<code>/</code>	unlock	

示例1 Using lockfs -a (续)

/var	unlock

example#

示例2 Using lockfs -w

The following example shows the lockfs output when the -w option is used to write lock the /var file system and the comment string is set using the -c option. The -a option is then specified on a separate command line.

```
example# /usr/sbin/lockfs -w -c "lockfs: write lock example" /var
example# /usr/sbin/lockfs -a
```

Filesystem	Locktype	Comment
/	unlock	
/var	write	lockfs: write lock example

example#

示例3 Using lockfs -u

The following example shows the lockfs output when the -u option is used to unlock the /var file system and the comment string is set using the -c option.

```
example# /usr/sbin/lockfs -uc "lockfs: unlock example" /var
example# /usr/sbin/lockfs /var
```

Filesystem	Locktype	Comment
/var	unlock	lockfs: unlock example

example#

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	system/core-os

另请参见

[kill\(1\)](#), [mount_ufs\(1M\)](#), [sync\(1M\)](#), [attributes\(5\)](#), [largefile\(5\)](#), [ufs\(7FS\)](#),
《Oracle Solaris 管理：常见任务》

诊断

file system: Not owner

You must be root to use this command.

file system :Deadlock condition detected/avoided

A file is enabled for accounting or swapping, on *file system*.

file system: Device busy

Another process is setting the lock on *file system*.

引用名	lockstat – report kernel lock and profiling statistics
用法概要	<pre>lockstat [-ACEHI] [-e event_list] [-i rate] [-b -t -h -s depth] [-n nrecords] [-l lock [, size]] [-d duration] [-f function [, size]] [-T] [-ckgwWRpP] [-D count] [-o filename] [-x opt [=val]] command [args]</pre>
描述	The lockstat utility gathers and displays kernel locking and profiling statistics. lockstat allows you to specify which events to watch (for example, spin on adaptive mutex, block on read access to rwlock due to waiting writers, and so forth) how much data to gather for each event, and how to display the data. By default, lockstat monitors all lock contention events, gathers frequency and timing data about those events, and displays the data in decreasing frequency order, so that the most common events appear first. lockstat gathers data until the specified command completes. For example, to gather statistics for a fixed-time interval, use <code>sleep(1)</code> as the command, as follows: <pre>example# lockstat sleep 5</pre> When the <code>-I</code> option is specified, lockstat establishes a per-processor high-level periodic interrupt source to gather profiling data. The interrupt handler simply generates a lockstat event whose caller is the interrupted PC (program counter). The profiling event is just like any other lockstat event, so all of the normal lockstat options are applicable. lockstat relies on DTrace to modify the running kernel's text to intercept events of interest. This imposes a small but measurable overhead on all system activity, so access to lockstat is restricted to super-user by default. The system administrator can permit other users to use lockstat by granting them additional DTrace privileges. Refer to the Solaris Dynamic Tracing Guide for more information about DTrace security features.
选项	The following options are supported:
Event Selection	If no event selection options are specified, the default is <code>-C</code>. <code>-A</code> Watch all lock events. <code>-A</code> is equivalent to <code>-CH</code>. <code>-C</code> Watch contention events. <code>-E</code> Watch error events. <code>-e event_list</code> Only watch the specified events. <i>event list</i> is a comma-separated list of events or ranges of events such as 1,4-7,35. Run lockstat with no arguments to get a brief description of all events.

	-H Watch hold events. -I Watch profiling interrupt events. -i <i>rate</i> Interrupt rate (per second) for -I. The default is 97 Hz, so that profiling doesn't run in lockstep with the clock interrupt (which runs at 100 Hz).
Data Gathering	-x <i>arg[=val]</i> Enable or modify a DTrace runtime option or D compiler option. The list of options is found in dtrace(1M). Boolean options are enabled by specifying their name. Options with values are set by separating the option name and value with an equals sign (=).
Data Gathering (Mutually Exclusive)	-b Basic statistics: lock, caller, number of events. -h Histogram: Timing plus time-distribution histograms. -s <i>depth</i> Stack trace: Histogram plus stack traces up to <i>depth</i> frames deep. -t Timing: Basic plus timing for all events [default].
Data Filtering	-d <i>duration</i> Only watch events longer than <i>duration</i>. -f <i>func[,size]</i> Only watch events generated by <i>func</i>, which can be specified as a symbolic name or hex address. <i>size</i> defaults to the ELF symbol size if available, or 1 if not. -l <i>lock[,size]</i> Only watch <i>lock</i>, which can be specified as a symbolic name or hex address. <i>size</i> defaults to the ELF symbol size or 1 if the symbol size is not available. -n <i>nrecords</i> Maximum number of data records. -T Trace (rather than sample) events [off by default].
Data Reporting	-c Coalesce lock data for lock arrays (for example, <code>pse_mutex[]</code>). -D <i>count</i> Only display the top <i>count</i> events of each type.

- g
Show total events generated by function. For example, if `foo()` calls `bar()` in a loop, the work done by `bar()` counts as work generated by `foo()` (along with any work done by `foo()` itself). The `-g` option works by counting the total number of stack frames in which each function appears. This implies two things: (1) the data reported by `-g` can be misleading if the stack traces are not deep enough, and (2) functions that are called recursively might show greater than 100% activity. In light of issue (1), the default data gathering mode when using `-g` is `-s 50`.
- k
Coalesce PCs within functions.
- o *filename*
Direct output to *filename*.
- P
Sort data by (*count * time*) product.
- p
Parsable output format.
- R
Display rates (events per second) rather than counts.
- W
Whichever: distinguish events only by caller, not by lock.
- w
Wherever: distinguish events only by lock, not by caller.

Display Formats

The following headers appear over various columns of data.

Count or ops/s

Number of times this event occurred, or the rate (times per second) if `-R` was specified.

indv

Percentage of all events represented by this individual event.

genr

Percentage of all events generated by this function.

cuml

Cumulative percentage; a running total of the individuals.

rcnt

Average reference count. This will always be 1 for exclusive locks (mutexes, spin locks, rwlocks held as writer) but can be greater than 1 for shared locks (rwlocks held as reader).

nsec

Average duration of the events in nanoseconds, as appropriate for the event. For the profiling event, duration means interrupt latency.

- Lock
- Address of the lock; displayed symbolically if possible.
- CPU+PIL
- CPU plus processor interrupt level (PIL). For example, if CPU 4 is interrupted while at PIL 6, this will be reported as `cpu[4]+6`.
- Caller
- Address of the caller; displayed symbolically if possible.

示例

```
示例1 Measuring Kernel Lock Contention

example# lockstat sleep 5
Adaptive mutex spin: 2210 events in 5.055 seconds (437 events/sec)

Count indiv cuml rcnt      nsec Lock                      Caller
-----
 269 12%  12% 1.00    2160 service_queue      background+0xdc
 249 11%  23% 1.00     86 service_queue      qenable_locked+0x64
 228 10%  34% 1.00    131 service_queue      background+0x15c
  68 3%   37% 1.00     79 0x30000024070      timeout+0x1c
  59 3%   40% 1.00    384 0x300066fa8e0      background+0xb0
  43 2%   41% 1.00     30 rqcred_lock          svc_getreq+0x3c
  42 2%   43% 1.00    341 0x30006834eb8      background+0xb0
  41 2%   45% 1.00    135 0x30000021058      timeout+0x1c
  40 2%   47% 1.00     39 rqcred_lock          svc_getreq+0x260
  37 2%   49% 1.00   2372 0x300068e83d0      hmemstart+0x1c4
  36 2%   50% 1.00     77 0x30000021058      timeout_common+0x4
  36 2%   52% 1.00    354 0x300066fa120      background+0xb0
  32 1%   53% 1.00     97 0x30000024070      timeout_common+0x4
  31 1%   55% 1.00   2923 0x300069883d0      hmemstart+0x1c4
  29 1%   56% 1.00    366 0x300066fb290      background+0xb0
  28 1%   57% 1.00    117 0x3000001e040      timeout+0x1c
  25 1%   59% 1.00     93 0x3000001e040      timeout_common+0x4
  22 1%   60% 1.00     25 0x30005161110      sync_stream_buf+0xdc
  21 1%   60% 1.00    291 0x30006834eb8      putq+0xa4
  19 1%   61% 1.00     43 0x3000515dcb0      mdf_alloc+0xc
  18 1%   62% 1.00    456 0x30006834eb8      qenable+0x8
  18 1%   63% 1.00     61 service_queue      queuerun+0x168
  17 1%   64% 1.00    268 0x30005418ee8      vmem_free+0x3c
[...]

R/W reader blocked by writer: 76 events in 5.055 seconds (15 events/sec)

Count indiv cuml rcnt      nsec Lock                      Caller
-----
  23 30%  30% 1.00 22590137 0x300098ba358      ufs_dirlook+0xd0
  17 22%  53% 1.00 5820995 0x3000ad815e8      find_bp+0x10
  13 17%  70% 1.00 2639918 0x300098ba360      ufs_iget+0x198
   4 5%   75% 1.00 3193015 0x300098ba360      ufs_getattr+0x54
```

示例1 Measuring Kernel Lock Contention (续)

```
3 4% 79% 1.00 7953418 0x3000ad817c0 find_bp+0x10
3 4% 83% 1.00 935211 0x3000ad815e8 find_read_lof+0x14
2 3% 86% 1.00 16357310 0x300073a4720 find_bp+0x10
2 3% 88% 1.00 2072433 0x300073a4720 find_read_lof+0x14
2 3% 91% 1.00 1606153 0x300073a4370 find_bp+0x10
1 1% 92% 1.00 2656909 0x300107e7400 ufs_iget+0x198
[...]
```

示例2 Measuring Hold Times

example# **lockstat -H -D 10 sleep 1**
Adaptive mutex spin: 513 events

Count	indv	cuml	rcnt	nsec	Lock	Caller
480	5%	5%	1.00	1136	0x300007718e8	putnext+0x40
286	3%	9%	1.00	666	0x3000077b430	getf+0xd8
271	3%	12%	1.00	537	0x3000077b430	msgio32+0x2fc
270	3%	15%	1.00	3670	0x300007718e8	strgetmsg+0x3d4
270	3%	18%	1.00	1016	0x300007c38b0	getq_noenab+0x200
264	3%	20%	1.00	1649	0x300007718e8	strgetmsg+0xa70
216	2%	23%	1.00	6251	tcp_mi_lock	tcp_snmp_get+0xfc
206	2%	25%	1.00	602	thread_free_lock	clock+0x250
138	2%	27%	1.00	485	0x300007c3998	putnext+0xb8
138	2%	28%	1.00	3706	0x300007718e8	strrput+0x5b8

[...]

示例3 Measuring Hold Times for Stack Traces Containing a Specific Function

example# **lockstat -H -f tcp_rput_data -s 50 -D 10 sleep 1**
Adaptive mutex spin: 11 events in 1.023 seconds (11 events/sec)

Count	indv	cuml	rcnt	nsec	Lock	Caller
9	82%	82%	1.00	2540	0x30000031380	tcp_rput_data+0x2b90

nsec	-----	Time Distribution	-----	count	Stack
256		@@@@@@@@@@@@@@@@		5	tcp_rput_data+0x2b90
512		@@@@@		2	putnext+0x78
1024		@@@		1	ip_rput+0xec4
2048				0	_c_putnext+0x148
4096				0	hmerread+0x31c
8192				0	hmeintr+0x36c
16384		@@@		1	sbus_intr_wrapper+0x30

[...]

示例 3 Measuring Hold Times for Stack Traces Containing a Specific Function (续)

```
Count  indiv  cuml  rcnt      nsec Lock      Caller
      1    9%  91%  1.00    1036 0x30000055380    freemsg+0x44

      nsec ----- Time Distribution ----- count      Stack
    1024 |@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@ 1    freemsg+0x44
                                           tcp_rput_data+0x2fd0
                                           putnext+0x78
                                           ip_rput+0xec4
                                           _c_putnext+0x148
                                           hmeread+0x31c
                                           hmeintr+0x36c

sbus_intr_wrapper+0x30
-----
[...]
```

示例 4 Basic Kernel Profiling

For basic profiling, we don't care whether the profiling interrupt sampled `foo()+0x4c` or `foo()+0x78`; we care only that it sampled somewhere in `foo()`, so we use `-k`. The CPU and PIL aren't relevant to basic profiling because we are measuring the system as a whole, not a particular CPU or interrupt level, so we use `-W`.

```
example# lockstat -kIW -D 20 ./polltest
Profiling interrupt: 82 events in 0.424 seconds (194
events/sec)
```

Count	indv	cuml	rcnt	nsec	Hottest	CPU+PIL	Caller
8	10%	10%	1.00	698	cpu[1]		utl0
6	7%	17%	1.00	299	cpu[0]		read
5	6%	23%	1.00	124	cpu[1]		getf
4	5%	28%	1.00	327	cpu[0]		fifo_read
4	5%	33%	1.00	112	cpu[1]		poll
4	5%	38%	1.00	212	cpu[1]		uiomove
4	5%	43%	1.00	361	cpu[1]		mutex_tryenter
3	4%	46%	1.00	682	cpu[0]		write
3	4%	50%	1.00	89	cpu[0]		pcache_poll
3	4%	54%	1.00	118	cpu[1]		set_active_fd
3	4%	57%	1.00	105	cpu[0]		syscall_trap32
3	4%	61%	1.00	640	cpu[1]		(usermode)
2	2%	63%	1.00	127	cpu[1]		fifo_poll
2	2%	66%	1.00	300	cpu[1]		fifo_write
2	2%	68%	1.00	669	cpu[0]		releasef
2	2%	71%	1.00	112	cpu[1]		bt_getlowbit
2	2%	73%	1.00	247	cpu[1]		splx

示例 4 Basic Kernel Profiling (续)

2	2%	76%	1.00	503	cpu[0]	mutex_enter
2	2%	78%	1.00	467	cpu[0]+10	disp_lock_enter
2	2%	80%	1.00	139	cpu[1]	default_copyin

[...]						

示例 5 Generated-load Profiling

In the example above, 5% of the samples were in `poll()`. This tells us how much time was spent inside `poll()` itself, but tells us nothing about how much work was *generated* by `poll()`; that is, how much time we spent in functions called by `poll()`. To determine that, we use the `-g` option. The example below shows that although `polltest` spends only 5% of its time in `poll()` itself, `poll()`-induced work accounts for 34% of the load.

Note that the functions that generate the profiling interrupt (`lockstat_intr()`, `cyclic_fire()`, and so forth) appear in every stack trace, and therefore are considered to have generated 100% of the load. This illustrates an important point: the generated load percentages do *not* add up to 100% because they are not independent. If 72% of all stack traces contain both `foo()` and `bar()`, then both `foo()` and `bar()` are 72% load generators.

example# **lockstat -kgIW -D 20 ./polltest**

Profiling interrupt: 80 events in 0.412 seconds (194 events/sec)

Count	genr	cuml	rcnt	nsec	Hottest	CPU+PIL	Caller

80	100%	----	1.00	310	cpu[1]		lockstat_intr
80	100%	----	1.00	310	cpu[1]		cyclic_fire
80	100%	----	1.00	310	cpu[1]		cbe_level14
80	100%	----	1.00	310	cpu[1]		current_thread
27	34%	----	1.00	176	cpu[1]		poll
20	25%	----	1.00	221	cpu[0]		write
19	24%	----	1.00	249	cpu[1]		read
17	21%	----	1.00	232	cpu[0]		write32
17	21%	----	1.00	207	cpu[1]		pcache_poll
14	18%	----	1.00	319	cpu[0]		fifo_write
13	16%	----	1.00	214	cpu[1]		read32
10	12%	----	1.00	208	cpu[1]		fifo_read
10	12%	----	1.00	787	cpu[1]		utl0
9	11%	----	1.00	178	cpu[0]		pcacheset_resolve
9	11%	----	1.00	262	cpu[0]		uimove
7	9%	----	1.00	506	cpu[1]		(usermode)
5	6%	----	1.00	195	cpu[1]		fifo_poll
5	6%	----	1.00	136	cpu[1]		syscall_trap32
4	5%	----	1.00	139	cpu[0]		releasef
3	4%	----	1.00	277	cpu[1]		polllock

[...]							

示例 6 Gathering Lock Contention and Profiling Data for a Specific Module

In this example we use the -f option not to specify a single function, but rather to specify the entire text space of the sbus module. We gather both lock contention and profiling statistics so that contention can be correlated with overall load on the module.

```
example# modinfo | grep sbus
24 102a8b6f b8b4 59 1 sbus (SBus (sysio) nexus driver)

example# lockstat -kICE -f 0x102a8b6f,0xb8b4 sleep 10
Adaptive mutex spin: 39 events in 10.042 seconds (4 events/sec)
```

Count	indv	cuml	rcnt	nsec	Lock	Caller
15	38%	38%	1.00	206	0x30005160528	sync_stream_buf
7	18%	56%	1.00	14	0x30005160d18	sync_stream_buf
6	15%	72%	1.00	27	0x300060c3118	sync_stream_buf
5	13%	85%	1.00	24	0x300060c3510	sync_stream_buf
2	5%	90%	1.00	29	0x300060c2d20	sync_stream_buf
2	5%	95%	1.00	24	0x30005161cf8	sync_stream_buf
1	3%	97%	1.00	21	0x30005161110	sync_stream_buf
1	3%	100%	1.00	23	0x30005160130	sync_stream_buf
[...]						

Adaptive mutex block: 9 events in 10.042 seconds (1 events/sec)

Count	indv	cuml	rcnt	nsec	Lock	Caller
4	44%	44%	1.00	156539	0x30005160528	sync_stream_buf
2	22%	67%	1.00	763516	0x30005160d18	sync_stream_buf
1	11%	78%	1.00	462130	0x300060c3510	sync_stream_buf
1	11%	89%	1.00	288749	0x30005161110	sync_stream_buf
1	11%	100%	1.00	1015374	0x30005160130	sync_stream_buf
[...]						

Profiling interrupt: 229 events in 10.042 seconds (23 events/sec)

Count	indv	cuml	rcnt	nsec	Hottest CPU+PIL	Caller
89	39%	39%	1.00	426	cpu[0]+6	sync_stream_buf
64	28%	67%	1.00	398	cpu[0]+6	sbus_intr_wrapper
23	10%	77%	1.00	324	cpu[0]+6	iommu_dvma_kaddr_load
21	9%	86%	1.00	512	cpu[0]+6	iommu_tlb_flush
14	6%	92%	1.00	342	cpu[0]+6	iommu_dvma_unload
13	6%	98%	1.00	306	cpu[1]	iommu_dvma_sync
5	2%	100%	1.00	389	cpu[1]	iommu_dma_bindhdl
[...]						

示例7 Determining the Average PIL (processor interrupt level) for a CPU

```
example# lockstat -Iw -l cpu[3] ./testprog
```

Profiling interrupt: 14791 events in 152.463 seconds (97 events/sec)

Count	indv	cuml	rcnt	nsec	CPU+PIL	Hottest Caller
13641	92%	92%	1.00	253	cpu[3]	(usermode)
579	4%	96%	1.00	325	cpu[3]+6	ip_ocsum+0xe8
375	3%	99%	1.00	411	cpu[3]+10	splx
154	1%	100%	1.00	527	cpu[3]+4	fas_intr_svc+0x80
41	0%	100%	1.00	293	cpu[3]+13	send_mondo+0x18
1	0%	100%	1.00	266	cpu[3]+12	zsa_rxint+0x400

[...]

示例8 Determining which Subsystem is Causing the System to be Busy

```
example# lockstat -s 10 -I sleep 20
```

Profiling interrupt: 4863 events in 47.375 seconds (103 events/sec)

Count	indv	cuml	rcnt	nsec	CPU+PIL	Caller
1929	40%	40%	0.00	3215	cpu[0]	usec_delay+0x78
nsec ----- Time Distribution ----- count Stack						
4096		@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@@	1872			ata_wait+0x90
8192			27			acersb_get_intr_status+0x34
16384			29			ata_set_feature+0x124
32768			1			ata_disk_start+0x15c
						ata_hba_start+0xbc
						ghd_waitq_process_and \
						_mutex_hold+0x70
						ghd_waitq_process_and \
						_mutex_exit+0x4
						ghd_transport+0x12c
						ata_disk_tran_start+0x108

[...]

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/dtrace

另请参见

[dtrace\(1M\)](#), [plockstat\(1M\)](#), [attributes\(5\)](#), [lockstat\(7D\)](#), [mutex\(9F\)](#), [rwlock\(9F\)](#)

Solaris Dynamic Tracing Guide

附注

The profiling support provided by `lockstat -I` replaces the old (and undocumented) `/usr/bin/kgmon` and `/dev/profile`.

Tail-call elimination can affect call sites. For example, if `foo()+0x50` calls `bar()` and the last thing `bar()` does is call `mutex_exit()`, the compiler can arrange for `bar()` to branch to `mutex_exit()` with a return address of `foo()+0x58`. Thus, the `mutex_exit()` in `bar()` will appear as though it occurred at `foo()+0x58`.

The PC in the stack frame in which an interrupt occurs can be bogus because, between function calls, the compiler is free to use the return address register for local storage.

When using the `-I` and `-s` options together, the interrupted PC will usually not appear anywhere in the stack since the interrupt handler is entered asynchronously, not by a function call from that PC.

The `lockstat` technology is provided on an as-is basis. The format and content of `lockstat` output reflect the current Solaris kernel implementation and are therefore subject to change in future releases.

引用名	lofiadm – 通过 lofi 管理作为块设备提供的文件
用法概要	<pre> /usr/sbin/lofiadm -a file [device] /usr/sbin/lofiadm -c crypto_algorithm -a file [device] /usr/sbin/lofiadm -c crypto_algorithm -k raw_key_file -a file [device] /usr/sbin/lofiadm -c crypto_algorithm -T token_key -a file [device] /usr/sbin/lofiadm -c crypto_algorithm -T token_key -k wrapped_key_file -a file [device] /usr/sbin/lofiadm -c crypto_algorithm -e -a file [device] /usr/sbin/lofiadm -C algorithm [-s segment_size] file /usr/sbin/lofiadm -d file device /usr/sbin/lofiadm -U file /usr/sbin/lofiadm [file device] </pre>
描述	lofiadm 管理回送文件驱动程序 lofi。lofi 可用于将文件与块设备关联，然后可通过块设备访问此文件。当文件包含某些文件系统的映像（例如 CD-ROM 映像）时，这会很有用，因为块设备随后可与常规系统实用程序一起用于挂载、检查或修复文件系统。请参见 fsck(1M) 和 mount(1M)。 可以使用 lofiadm 将文件添加为回送设备、删除这样的关联或显示有关当前关联的信息。 潜在 lofi 设备的数量受 zone.max-lofi rctl 的限制，后者可通过 zonecfg(1M) 在全局区域(global zone)中进行设置。有关 zone.max-lofi 的说明，请参见 resource_controls(5)。 不能在命令行上同时使用加密和压缩选项。此外，加密文件无法再进行压缩，而压缩文件也无法再进行加密。 在全局区域中，全局区域设备以及系统中其他非全局区域所拥有的所有设备都可以使用 lofiadm。
选项	所支持的选项如下： -a file [device] 将 file 添加为块设备。 如果没有指定 device，则会选择一个可用设备。 如果指定了 device，lofiadm 会尝试将其分配给 file。device 必须可用，否则 lofiadm 将失败。指定设备的功能可用于在脚本中重新建立一组特定的关联。 -C {gzip gzip-N lzma} 使用指定压缩算法压缩文件。

`gzip` 压缩算法使用与开源 `gzip` 命令相同的压缩。要指定 `gzip` 级别，可使用值 `gzip-N`，其中 `N` 为 6（快速）或 9（最佳压缩比）。当前，不含数字的 `gzip` 相当于 `gzip-6`（这也是 `gzip` 命令的缺省值）。

`lzma` 表示 LZMA (Lempel-Ziv-Markov) 压缩算法。

请注意，不能对压缩文件进行写入，也不能以读/写方式挂载压缩文件。

-d *file* | *device*

根据 *file* 或 *device* 名称删除关联（如果关联的块设备未处于忙碌状态），并解除块设备的分配。

-s *segment_size*

用于划分所压缩文件的段大小。*segment_size* 可以是 512 的整数倍。

-U *file*

解压缩某个压缩文件。

加密文件时可以使用下列选项：

-c *crypto_algorithm*

选择加密算法。启用加密时必须指定算法，因为磁盘映像中并未存储算法。

如果 `-e`、`-k` 或 `-T` 均未指定，则 `lofiadm` 会提示输入最短八个字符长的口令短语。口令短语用于使用 PKCS#5 PBKD2 派生出一个对称加密密钥。

-k *raw_key_file* | *wrapped_key_file*

原始或包装的对称加密密钥的路径。如果还使用 `-T` 选项指定了 PKCS#11 对象，则密钥将由该对象包装。如果未指定 `-T`，则以原始状态使用密钥。

-T *token_key*

PKCS#11 标记中的密钥，用于加密或用于解开密钥文件的包装。

如果同时指定了 `-k`，则 `-T` 表示解开包装的密钥，该密钥必须为 RSA 私钥。

-e

生成临时对称加密密钥。

操作数

支持下列操作数：

crypto_algorithm

`aes-128-cbc`、`aes-192-cbc`、`aes-256-cbc`、`des3-cbc`、`blowfish-cbc` 之一。

device

显示与块设备 *device* 关联的文件名。

如果不指定参数，则会显示当前关联的列表。文件名必须为有效的绝对路径名。

添加文件时，文件会打开，供 `root` 读取或写入。任何限制均适用（例如通过 NFS 的受限 `root` 访问）。文件将保持打开状态，直至删除关联。在使用块设备之前无法实际访问文件，因此如果仅以只读方式打开块设备，将永远无法对文件进行写入。

请注意，如果无法在当前上下文中解析路径（例如，如果它是非全局区域中的一个 NFS 路径），则文件名可能显示为问号 (?)。

file

显示与 *file* 关联的块设备。

raw_key_file

要用作原始对称加密密钥的适当长度（以位为单位）的文件的途径。

token_key

以下格式的 PKCS#11 标记对象：

token_name:manufacturer_id:serial_number:key_label

除密钥标签外的其他各项均为可选项且可以为空。例如，要只使用其密钥标签 MylofiKey 指定标记对象，请使用：

-T ::MylofiKey

wrapped_key_file

包含由 **-T** 指定的 RSA 私钥包装的对称加密密钥的文件的途径。

示例

示例 1 挂载现有 CD-ROM 映像

在创建 CD 之前应确保 Solaris 了解此映像。lofi 允许您挂载映像并查看其是否有效。

此示例将挂载一个从 Internet 下载的 Red Hat 6.0 CD 的现有 CD-ROM 映像 (sparc.iso)。它是使用来自 Internet 的 mkisofs 实用程序创建的。

使用 lofiadm 为其关联一个块设备：

```
# lofiadm -a /home/mike_s/RH6.0/sparc.iso
/dev/lofi/1
```

lofiadm 将选择设备并将设备名称打印到标准输出中。您可以通过发出以下命令再次运行 lofiadm：

```
# lofiadm
Block Device      File                                Options
/dev/lofi/1       /home/mike_s/RH6.0/sparc.iso      -
```

或者，您可以通过发出以下命令为其提供一个名称并查询另一个名称：

```
# lofiadm /dev/lofi/1
/home/mike_s/RH6.0/sparc.iso
```

使用 mount 命令挂载映像：

```
# mount -F hsfs -o ro /dev/lofi/1 /mnt
```

检查以确保 Solaris 了解该映像：

```
# df -k /mnt
Filesystem          kbytes    used   avail capacity  Mounted on
```

示例1 挂载现有 CD-ROM 映像 (续)

```
/dev/lofi/1          512418 512418          0 100%  /mnt
# ls /mnt
./          RedHat/      doc/          ls-lR        rr_moved/
../         TRANS.TBL   dosutils/    ls-lR.gz    sbin@
.buildlog   bin@        etc@         misc/       tmp/
COPYING     boot/       images/      mnt/        usr@
README     boot.cat*   kernels/     modules/
RPM-PGP-KEY dev@        lib@         proc/
```

Solaris 可挂载 CD-ROM 映像并可识别文件名。映像创建正确，现在您可以放心地创建 CD-ROM。

最后一步，取消挂载并分离映像：

```
# umount /mnt
# lofiadm -d /dev/lofi/1
# lofiadm
Block Device          File          Options
```

示例2 在文件上创建 UFS 文件系统

在文件上创建 UFS 文件系统可能很有用，尤其是在测试套件需要临时文件系统时。单纯为了测试套件而被迫对磁盘进行重新分区会令人感到痛苦（或烦恼），但您不必如此。您可以通过 lofi 对文件进行 newfs。

创建文件：

```
# mkfile 35m /export/home/test
```

将其关联到块设备。您还会获得 newfs 所需的字符设备，因此 newfs 将如下所示：

```
# lofiadm -a /export/home/test
/dev/lofi/1
# newfs /dev/rlofi/1
newfs: construct a new file system /dev/rlofi/1: (y/n)? y
/dev/rlofi/1: 71638 sectors in 119 cylinders of 1 tracks, 602 sectors
          35.0MB in 8 cyl groups (16 c/g, 4.70MB/g, 2240 i/g)
super-block backups (for fsck -F ufs -o b=#) at:
    32, 9664, 19296, 28928, 38560, 48192, 57824, 67456,
```

请注意，ufs 可能无法使用整个文件。挂载并使用文件系统：

```
# mount /dev/lofi/1 /mnt
# df -k /mnt
Filesystem          kbytes    used    avail capacity  Mounted on
/dev/lofi/1          33455      9    30101      1%    /mnt
# ls /mnt
./          ../          lost+found/
# umount /mnt
```


示例2 在文件上创建 UFS 文件系统 (续)

```
# lofiadm -d /dev/lofi/1
```

示例3 在 Unix 文件上创建 PC (FAT) 文件系统

下列命令将在 Unix 文件上创建 FAT 文件系统。此文件与由 lofiadm 创建的块设备关联。

```
# mkfile 10M /export/test/testfs
# lofiadm -a /export/test/testfs
/dev/lofi/1
Note use of rlofi, not lofi, in following command.
# mkfs -F pcfs -o nofdisk,size=20480 /dev/rlofi/1
Construct a new FAT file system on /dev/rlofi/1: (y/n)? y
# mount -F pcfs /dev/lofi/1 /mnt
# cd /mnt
# df -k .


| Filesystem  | kbytes | used | avail | capacity | Mounted on |
|-------------|--------|------|-------|----------|------------|
| /dev/lofi/1 | 10142  | 0    | 10142 | 0%       | /mnt       |


```

示例4 压缩现有 CD-ROM 映像

以下示例将揭示如何压缩一个现有 CD-ROM 映像 (solaris.iso)，验证映像是否已压缩，然后将其解压缩。

```
# lofiadm -C gzip /export/home/solaris.iso
```

使用 lofiadm 为其关联一个块设备：

```
# lofiadm -a /export/home/solaris.iso
/dev/lofi/1
```

检查映射的映像是否已压缩：

```
# lofiadm


| Block Device | File                     | Options          |
|--------------|--------------------------|------------------|
| /dev/lofi/1  | /export/home/solaris.iso | Compressed(gzip) |
| /dev/lofi/2  | /export/home/regular.iso | -                |


```

取消映射压缩的映像并将其解压缩：

```
# lofiadm -d /dev/lofi/1
# lofiadm -U /export/home/solaris.iso
```

示例5 在文件上创建加密的 UFS 文件系统

本例与上述在文件上创建 UFS 文件系统的示例类似。

创建文件：

```
# mkfile 35m /export/home/test
```

示例 5 在文件上创建加密的 UFS 文件系统 (续)

将文件关联到块设备并指定文件映像已加密。作为此命令的结果，您将获得字符设备，该设备随后由 newfs 使用：

```
# lofiadm -c aes-256-cbc -a /export/home/secrets
Enter passphrase: My-M0th3r;l0v3s_m3+4lw4ys!      (not echoed)
Re-enter passphrase: My-M0th3r;l0v3s_m3+4lw4ys!   (not echoed)
/dev/lofi/1

# newfs /dev/rlofi/1
newfs: construct a new file system /dev/rlofi/1: (y/n)? y
/dev/rlofi/1: 71638 sectors in 119 cylinders of 1 tracks, 602 sectors
35.0MB in 8 cyl groups (16 c/g, 4.70MB/g, 2240 i/g)
super-block backups (for fsck -F ufs -o b=#) at:
32, 9664, 19296, 28928, 38560, 48192, 57824, 67456,
```

映射的文件系统显示已启用加密：

```
# lofiadm
Block Device      File              Options
/dev/lofi/1       /export/home/secrets Encrypted
```

挂载并使用文件系统：

```
# mount /dev/lofi/1 /mnt
# cp moms_secret_*_recipe /mnt
# ls /mnt
./          moms_secret_cookie_recipe  moms_secret_soup_recipe
../         moms_secret_fudge_recipe   moms_secret_stuffing_recipe
lost+found/ moms_secret_meatloaf_recipe moms_secret_waffle_recipe
# umount /mnt
# lofiadm -d /dev/lofi/1
```

使用错误密钥或错误加密算法映射文件系统的后续尝试将失败：

```
# lofiadm -c blowfish-cbc -a /export/home/secrets
Enter passphrase: mommy      (not echoed)
Re-enter passphrase: mommy   (not echoed)
lofiadm: could not map file /root/lofi: Invalid argument
# lofiadm
Block Device      File              Options
#
```

在不使用加密的情况下映射文件系统的尝试将成功，但是挂载和使用文件系统的尝试将失败：

```
# lofiadm -a /export/home/secrets
/dev/lofi/1
# lofiadm
Block Device      File              Options
```

示例 5 在文件上创建加密的 UFS 文件系统 (续)

```
/dev/lofi/1      /export/home/secrets      -
# mount /dev/lofi/1 /mnt
mount: /dev/lofi/1 is not this fstype
#
```

环境变量 有关下列影响 lofiadm 的执行的環境變量的說明，請參見 [environ\(5\)](#)：LC_CTYPE、LC_MESSAGES 和 NLSPATH。

退出状态 将返回下列退出值：

```
0
    成功完成。
>0
    出現錯誤。
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见 [fsck\(1M\)](#)、[mount\(1M\)](#)、[mount_ufs\(1M\)](#)、[newfs\(1M\)](#)、[zonecfg\(1M\)](#)、[attributes\(5\)](#)、[resource](#)

附注 就像您不会直接访问挂载了文件系统的磁盘设备一样，您也不能访问与块设备关联的文件，除非通过 lofi 文件驱动程序。除此之外，最好再确保文件设有适当权限以禁止这类访问。

lofiadm 的功能以及可以使用它们的用户是由 /dev/lofictl 的权限控制的。读取访问权限允许执行查询操作，例如列出所有关联。执行任何状态更改操作都需要写入访问权限，例如添加关联。出厂时 /dev/lofictl 归 root 所有，位于组 sys 中且模式为 0644，因此所有用户均可执行查询操作，但只有 root 能够进行任何更改。管理员可赋予用户写入访问权限，允许他们添加或删除关联，但这很有可能成为安全漏洞，因此可能只应赋予受信任的组。

挂载文件系统映像时，请注意使用适当挂载选项。尤其是，nosuid 挂载选项可能适合于来源未知的 UFS 映像。此外，某些选项可能没有用或者不适合，例如 logging 或 forcedirectio 对于 UFS。出于兼容性目的，原始设备也会随块设备一起导出。例如，newfs(1M) 就需要一个原始设备。

lofiadm（不含参数）的输出在将来的发行版中可能会有所变化。

引用名 logadm – manage endlessly growing log files

用法概要 logadm

logadm [-options] logname...

描述 logadm is a general log rotation tool that is suitable for running from [cron\(1M\)](#).

Without arguments, logadm reads the /etc/logadm.conf file, and for every entry found in that file checks the corresponding log file to see if it should be rotated. Typically this check is done each morning by an entry in the root's crontab.

If the *logname* argument is specified, logadm renames the corresponding log file by adding a suffix so that the most recent log file ends with .0 (that is, *logfile.0*), the next most recent ends with .1 (that is, *logfile.1*), and so forth. By default, ten versions of old log files are kept (that is, *logfile.0* through *logfile.9*). At the point when what would be the eleventh file is logged, logadm automatically deletes the oldest version to keep the count of files at ten.

logadm takes a number of *options*. You can specify these options on the command line or in the /etc/logadm.conf file. The logadm command searches /etc/logadm.conf for lines of the form *logname options*

logname

Identifies an entry in /etc/logadm.conf. This can be a name or the pathname of the log file. If you specify a log file, rather than a name, for this field, it must be a fully qualified pathname.

options

Identifies command line options exactly as they would be entered on the command line. This allows commonly used log rotation policies to be stored in the /etc/logadm.conf file. See EXAMPLES.

If *options* are specified both in /etc/logadm.conf and on the command line, those in the /etc/logadm.conf file are applied first. Therefore, the command line options override those in /etc/logadm.conf.

Log file names specified in /etc/logadm.conf may contain filename substitution characters such as * and ?, that are supported by [csh\(1\)](#).

Two options control when a log file is rotated. They are: -s size -p period.

When using more than one of these options at a time, there is an implied *and* between them. This means that all conditions must be met before the log is rotated.

If neither of these two options are specified, the default conditions for rotating a log file are: -s 1b -p 1w, which means the log file is only rotated if the size is non-zero and if at least 1 week has passed since the last time it was rotated.

By specifying `-p never` as a rotation condition, any other rotation conditions are ignored and `logadm` moves on to the expiration of old log files. By specifying `-p now` as a rotation condition, a log rotation is forced.

Unless specified by the `-o`, `-g`, or `-m` options, `logadm` replaces the log file (after renaming it) by creating an empty file whose owner, group ID, and permissions match the original file.

Three options control when old log files are expired: `-A age` `-C count` `-S size`. These options expire the oldest log files until a particular condition or conditions are met. For example, the combination `-C 5` and the `-S 10m` options expires old log files until there are no more than 5 of the *and* their combined disk usage is no more than 10 megabytes. If none of these options are specified, the default expiration is `-C 10` which keeps ten old log files. If no files are to be expired, use `-C 0` to prevent expiration by default.

`logadm` stores timestamps in the file `/var/logadm/timestamps`. For users of previous versions of `logadm`, the utility automatically moves timestamps from `/etc/logadm.conf`, their previous repository, to `/var/logadm/timestamps`.

选项

The following options are supported:

`-a post_command`

Execute the *post_command* after renaming the log file. *post_command* is passed to `sh -c`.

Specify *post_command* as a valid shell command. Use quotes to protect spaces or shell metacharacters in *post_command*.

This option can be used to restart a daemon that is writing to the file. When rotating multiple logs with one `logadm` command, *post_command* is executed only once after all the logs are rotated, not once per rotated log.

`-A age`

Delete any versions that have not been modified for the amount of time specified by *age*.

Specify *age* as a number followed by an `h` (hours), `d` (days), `w`(weeks), `m` (months), or `y` (years).

`-b pre_command`

Execute *pre_command* before renaming the log file. *pre_command* is passed to `sh -c`.

Specify *pre_command* as a valid shell command. Use quotes to protect spaces or shell metacharacters in the *pre_command*.

This option can be used to stop a daemon that is writing to the file. When rotating multiple logs with one `logadm` command, *pre_command* is executed only once before all the logs are rotated, not once per rotated log.

`-c`

Rotate the log file by copying it and truncating the original logfile to zero length, rather than renaming the file.

-C *count*

Delete the oldest versions until there are not more than *count* files left.

If no expire options (-A, -C, or -S) are specified, -C 10 is the default. To prevent the default expire rule from being added automatically, specify -C 0.

-e *mail_addr*

Send error messages by email to *mail_addr*.

As logadm is typically run from [cron\(1M\)](#), error messages are captured by cron and mailed to the owner of the crontab.

This option is useful if you want the mail regarding error messages to go to another address instead. If no errors are encountered, no mail message is generated.

-E *cmd*

Execute *cmd* to expire the file, rather than deleting the old log file to expire it.

cmd is passed it to `sh -c`. The file is considered expired after *cmd* completes. If the old log file is not removed or renamed by the *cmd*, logadm considers it for expiration the next time that it runs on the specified log file. If present, the keyword `$file` is expanded in the specified *cmd* to the name of the file being expired.

This option is useful for tasks such as mailing old log files to administrators, or copying old log files to long term storage.

-f *conf_file*

Use *conf_file* instead of `/etc/logadm.conf`.

This option allows non-root users to keep their own logadm configuration files.

-F *timestamp_file*

Use *timestamp_file* instead of `/var/logadm/timestamps` to store logadm timestamps.

-g *group*

Create a new empty file with the ID specified by *group*, instead of preserving the group ID of the log file.

Specify *group* by name or by numeric group ID, as accepted by [chgrp\(1\)](#).

This option requires the ability to change file group ownership using the [chgrp\(1\)](#) command.

-h

Print a help message that describes logadm's options.

-l

Use local time rather than the Coordinated Universal Time (UTC) when naming rotated log files (see the discussion of percent sequences in the templates supplied with the -t option).

-m *mode*

Create a new empty file with the mode specified by *mode*, instead of preserving the mode of the log file.

Specify *mode* in any form that is accepted by the [chmod\(1\)](#) command.

-M *cmd*

Use *cmd* to rename the log file. If the keyword *\$file* is specified, it is expanded to the name of the log file. Similarly, the keyword *\$nfile* is expanded to the new name of the log file. The *\$nfile* keyword is only available with commands provided with the **-M** option. After the command completes, the log file is replaced by the rotate file. The default *cmd* is `"/bin/mv $file$nfile"`.

-n

Print the actions that the `logadm` command will perform without actually performing them.

This option is useful for checking arguments before making any changes to the system.

It is important to remember, however, that since log rotating actions are only printed with this option, `logadm` might not find files that need expiring, but if run without the **-n** `logadm` might create a file that needs expiring by performing the log rotating actions. Therefore, if you see no files being expired with the **-n** option, files still might be expired without it.

-N

Prevent an error message if the specified logfile does not exist. Normally, `logadm` produces an error message if the log file is not found. With **-N**, if the log file doesn't exist `logadm` moves on to the expire rules (if any) and then to the next log file (if any), without creating the empty replacement log file.

-o *owner*

Create the new empty file with *owner*, instead of preserving the owner of the log file.

Specify *owner* in any form that is accepted by the [chown\(1\)](#) command.

-p *period*

Rotate a log file after the specified time period (*period*).

Specify *period* as a number followed by **d** for days, **h** for hours, **w** for weeks, **m** for months (30 days) or **y** for years. There are also two special values for period: **now** and **never**. `"-p now"` forces log rotation. `"-p never"` forces no log rotation.

-P *timestamp*

Used by `logadm` to record the last time the log was rotated in `/var/logadm/timestamps`.

This option uses *timestamp* to determine if the log rotation period has passed. The format of *timestamp* matches the format generated by [ctime\(3C\)](#), with quotes around it to protect embedded spaces. *timestamp* is always recorded in the Coordinated Universal Time (UTC) timezone.

-r
Remove any entries corresponding to the specified *logname* from the `/etc/logadm.conf`.

-R *cmd*
Run the *cmd* when an old log file is created by a log rotation. If the keyword `$file` is embedded in the specified command, it is expanded to the name of the old log file just created by log rotation.

This option is useful for processing log file contents after rotating the log. *cmd* is executed by passing it to `sh -c`. When rotating multiple logs with one `logadm` command, the command supplied with `-R` is executed once every time a log is rotated. This is useful for post-processing a log file (that is, sorting it, removing uninteresting lines, etc.). The `-a` option is a better choice for restarting daemons after log rotation.

-s *size*
Rotate the log file only if its size is greater than or equal to *size*.

Specify *size* as a number followed by the letter `b` for bytes, `k` for kilobytes, `m` for megabytes, or `g` for gigabytes.

-S *size*
Delete the oldest versions until the total disk space used by the old log files is less than the specified size.

Specify *size* as a number followed by the letter `b` for bytes, `k` for kilobytes, `m` for megabytes, or `g` for gigabytes.

-t *template*
Specify the template to use when renaming log files.

template can be a simple name, such as `/var/adm/oldfile`, or it can contain special keywords which are expanded by `logadm` and are in the form `$word`. Allowed sequences are:

\$basename
The log file name, without the directory name

\$dirname
The directory of the file to be rotated

\$domain
Expands to the output of `domainname`

\$file
The full path name of the file to be rotated

\$isa
Expands to the output of `uname -p`

\$machine
Expands to the output of `uname -m`

\$n

The version number, 0 is most recent, 1 is next most recent, and so forth

\$N

The same as *\$n*, but starts at 1 instead of zero

\$nodename

Expands to the output of `uname -n`

\$platform

Expands to the output of `uname -i`

\$release

Expands to the output of `uname -r`

\$secs

The number of seconds since 00:00:00 UTC, January 1, 1970

\$zonename

Expands to the output of `zonename(1)`.

To actually have the dollar sign character in the file name, use `$$`. Any percent sequences allowed by `strftime(3C)` are also allowed, for example, `%d` expands to the day of the month. To actually have a percent sign character in the file name, use `%%`. Both dollar-sign keywords and percent sequences can appear anywhere in the template. If the template results in a pathname with non-existent directories, they are created as necessary when rotating the log file.

If no `-t` option is specified, the default template is `$file.$n`. Actual *rotation* of log files, where each version is shifted up until it expires is done using the *\$n* keyword. If the template does not contain the *\$n* keyword, the log file is simply renamed to the new name and then the expire rules, if any, are applied.

`-T pattern`

Normally `logadm` looks for a list of old log files by turning the template (specified with the `-t` option) into a pattern and finding existing files whose names match that pattern. The `-T` option causes the given pattern to be used instead.

This option is useful if another program fiddles with the old log file names, like a `cron` job to compress them over time. The pattern is in the form of a pathname with special characters such as `*` and `?` as supported by `csh(1)` filename substitution.

`-v`

Print information about the actions being executed in verbose mode.

`-V`

Validate the configuration file.

This option validates that an entry for the specified *logname* exists in the `/etc/logadm.conf` file and is syntactically correct. If *logname* is not specified, all entries in

the configuration file are validated. If a *logname* argument is specified, the command validates the syntax of that entry. If the entry is found, it is printed and the exit value of the command is true. Otherwise the exit value is false.

-w *entryname*

Write an entry into the config file (that is, `/etc/logadm.conf`) that corresponds to the current command line arguments. If an entry already existed for the specified *entryname*, it is removed first. This is the preferred method for updating `/etc/logadm.conf`, because it prevents syntax errors. The *entryname* is an argument to an invocation of `logadm`. *entryname* might be chosen as something easy to remember or it can be the pathname of the log file. If a pathname, rather than a name is used, it must be a fully qualified pathname.

If no log file name is provided on a `logadm` command line, the entry name is assumed to be the same as the log file name. For example, the following two lines achieve the same thing, keeping two copies of rotated log files:

```
% logadm -C2 -w mylog /my/really/long/log/file/name
% logadm -C2 -w /my/really/long/log/file/name
```

-z *count*

Compress old log files after all other commands have been executed. *count* of the most recent log files are left uncompressed, therefore making the *count* most recent files easier to peruse. Use *count* of zero to compress all old logs.

The compression is done with `gzip(1)` and the resulting log file has the suffix of `.gz`.

操作数

The following operands are supported:

logname

Identifies the name of the entry in `/etc/logadm.conf`. If the log file name is specified in the *logname* field, it is assumed that *logname* is the same as the actual log file name.

示例

示例 1 Rotating a File and Keeping Previous Versions

The following example rotates the `/var/adm/exacct/proc` file, keeping ten previous versions in `/var/adm/exacct/proc.0` through `/var/adm/exacct/proc.9`.

Tell `logadm` to copy the file and truncate it.

```
% logadm -c /var/adm/exacct/proc
```

示例 2 Rotating syslog

The following example rotates `syslog` and keeps eight log files. Old log files are put in the directory `/var/oldlogs` instead of `/var/log`:

```
% logadm -C8 -t'/var/oldlogs/syslog.$n' /var/log/syslog
```

示例 3 Rotating /var/adm/sulog and Expiring Based on Age

The following entry in the `/etc/logadm.conf` file rotates the `/var/adm/sulog` file and expires any copies older than 30 days.

```
/var/adm/sulog -A 30d
```

示例 4 Rotating Files and Expiring Based on Disk Usage

The following entry in the `/etc/logadm.conf` file rotates the `/var/adm/sulog` file and expires old log files when more than 100 megabytes are used by the sum of all the rotated log files.

```
/var/adm/sulog -S 100m
```

示例 5 Creating an Entry that Stores the Logfile Name

This example creates an entry storing the log file name and the fact that we want to keep 20 copies in `/etc/logadm.conf`, but the `-p` never means the entry is ignored by the normal `logadm` run from root's crontab every morning.

```
% logadm -w locallog /usr/local/logfile -C20 -p never
```

Use the following entry on the command line to override the `-p` never option:

```
% logadm -p now locallog
```

示例 6 Rotating the apache Error and Access Logs

The following example rotates the apache error and access logs monthly to filenames based on current year and month. It keeps the 24 most recent copies and tells apache to restart after renaming the logs.

This command is run once, and since the `-w` option is specified, an entry is made in `/etc/logadm.conf` so the apache logs are rotated from now on.

```
% logadm -w apache -p 1m -C 24\  
-t '/var/apache/old-logs/${basename.%Y-%m}'\  
-a '/usr/apache/bin/apachectl graceful'\  
'/var/apache/logs/*{access,error}_log'
```

This example also illustrates that the entry name supplied with the `-w` option doesn't have to match the log file name. In this example, the entry name is `apache` and once the line has been run, the entry in `/etc/logadm.conf` can be forced to run by executing the following command:

```
% logadm -p now apache
```

Because the expression matching the apache log file names was enclosed in quotes, the expression is stored in `/etc/logadm.conf`, rather than the list of files that it expands to. This means that each time `logadm` runs from `cron` it expands that expression and checks all the log files in the resulting list to see if they need rotating.

示例 6 Rotating the apache Error and Access Logs (续)

The following command is an example without the quotes around the log name expression. The shell expands the last argument into a list of log files that exist at the time the command is entered, and writes an entry to `/etc/logadm.conf` that rotates the files.

文件

`logadm -w apache /var/apache/logs/*_log`
`/etc/logadm.conf`
configuration file for `logadm` command
`/var/logadm/timestamps`
repository for logging timestamps

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[chgrp\(1\)](#), [chmod\(1\)](#), [chown\(1\)](#), [csh\(1\)](#), [gzip\(1\)](#), [cron\(1M\)](#), [ctime\(3C\)](#), [strftime\(3C\)](#), [logadm.conf\(4\)](#), [attributes\(5\)](#)

附注

When `logadm` applies expire conditions (supplied by the `-A`, `-C`, and `-S` options), it deletes files, the oldest first, until the conditions are satisfied. If the template used for naming the old logs contained `$n` or `$N`, `logadm` picks the highest value of `$n` or `$N` found in the old log file names first. If the template used is something else, `logadm` uses the modification time to determine which files to expire first. This may not be the expected behavior if an old log file has been modified since it was rotated.

Depending on log file sizes and number of log files, log file rotations can be very time-consuming.

By default, `logadm` works in UTC. Therefore, all entries written to the `/etc/logadm.conf` file (see [logadm.conf\(4\)](#)) will have a UTC timestamp. Users can use the `-l` option to set `logadm` to local time.

The `-f` and `-F` options can specify the same file, in which case `logadm` reverts to the same behavior as in prior releases. That is, timestamps are written to the configuration file.

引用名	logins – list user and system login information
用法概要	<pre>/usr/bin/logins [-admpstux] [-g group...] [-l login_name...]</pre>
描述	This command displays information on user and system logins known to the system. Contents of the output is controlled by the command options and can include the following: user or system login, user id number, passwd account field value (user name or other information), primary group name, primary group id, multiple group names, multiple group ids, home directory, login shell, and four password aging parameters. The default information is the following: login id, user id, primary group name, primary group id and the account field value. Output is sorted by user id, system logins, followed by user logins.
选项	Options may be used together. If so, any login that matches any criteria are displayed. The following options are supported: -a Add two password expiration fields to the display. The fields show how many days a password can remain unused before it automatically becomes inactive, and the date that the password expires. -d Selects logins with duplicate uids. -g <i>group</i> Selects all users belonging to <i>group</i>, sorted by login. Multiple groups can be specified as a comma-separated list. When the -l and -g options are combined, a user is only listed once, even if the user belongs to more than one of the selected groups. -l <i>login_name...</i> Selects the requested login. Multiple logins can be specified as a comma-separated list. Depending on the nameservice lookup types set in <code>/etc/nsswitch.conf</code>, the information can come from the <code>/etc/passwd</code> and <code>/etc/shadow</code> files and other nameservices. When the -l and -g options are combined, a user is only listed once, even if the user belongs to more than one of the selected groups. -m Displays multiple group membership information. -o Formats output into one line of colon-separated fields. -p Selects logins with no passwords. -s Selects all system logins. -t Sorts output by login instead of by uid.

- u
Selects all user logins.
- x
Prints an extended set of information about each selected user. The extended information includes home directory, login shell and password aging information, each displayed on a separate line. The password information currently consists of password status:
 - NP
Account has no password
 - LK
Account is locked for UNIX authentication
 - NL
Account is a no login account
 - UP
This account has not yet been activated by the administrator and cannot be used.
 - PS
Account probably has a valid password
 - UN
Account password status is unknown. That is, it is not a recognizable hashed password or any of the above entries. See [crypt\(3C\)](#) for valid password hashes.

If the login is passworded, status is followed by the date the password was last changed, the number of days required between changes, and the number of days allowed before a change is required. The password aging information shows the time interval that the user receives a password expiration warning message (when logging on) before the password expires.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [passwd\(1\)](#), [crypt\(3C\)](#), [attributes\(5\)](#)

引用名	lpget – get printing configuration
用法概要	lpget [-k <i>key</i>] [<i>destination...</i> <i>list</i>]
描述	The lpget utility reads printing configuration information from the configuration databases in <code>\$HOME/.printers</code> and <code>printers.org_dir</code> printer. This information, called a <i>configuration report</i>, is displayed to the standard output. lpget displays a configuration report for all keys for the specified destination or destinations by default. Use the -k option to display a configuration report for specific keys. Use the <i>list</i> operand to display a configuration report for all configured destinations.
选项	The following option is supported: -k <i>key</i> Displays a configuration report for <i>key</i>.
操作数	The following operands are supported: <i>destination</i> Displays a configuration report for <i>destination</i>. Destination can be either a printer or a class of printers. Specify <i>destination</i> using atomic or POSIX-style (<i>server:destination</i>) names. See standards(5) for information concerning POSIX. <i>list</i> Displays a configuration report for all configured destinations.
示例	示例 1 Displaying a Configuration Report for the <code>bsdaddr</code> Key The following example displays a configuration report for the <code>bsdaddr</code> key for printer <code>catalpa</code>. <pre>example% lpget -k bsdaddr catalpa</pre> 示例 2 A Configuration Report for all Keys for all Configured Destinations The following example displays a configuration report for all keys for all configured destinations. <pre>example% lpget list</pre>
退出状态	The following exit values are returned: 0 Successful completion. non-zero An error occurred.
文件	<code>\$HOME/.printers</code> User-configurable printer database.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	print/lp/print-client-commands
Interface Stability	Committed

另请参见

[lpset\(1M\)](#), [attributes\(5\)](#), [standards\(5\)](#), [ldap\(5\)](#)

《[Oracle Solaris Administration: Naming and Directory Services](#)》

附注

If the LDAP database is used as the name service. If the [ldapclient\(1M\)](#) server is a replica LDAP server, LDAP printer database updates may not appear immediately, as the replica server may not have been updated by the master server and can be out of sync. For example, a printer that you deleted by using [lpset\(1M\)](#) may still appear in the printer list you display with `lpget` until the replica is updated from the master. Replica servers vary as to how often they are updated from the master. Refer to the 《[Oracle Solaris Administration: Naming and Directory Services](#)》 for more information on LDAP replication.

引用名	lpset – set printing configuration in supported databases
用法概要	<pre>lpset [-n system fnsldap] [-x] [[-D binddn] [-w passwd] [-h ldaphost]] [-a key=value] [-d key] destination</pre>
描述	The lpset utility sets printing configuration information in the system configuration databases. See nsswitch.conf(4). Only a superuser or a member of Group 14 may execute lpset.
选项	The following options are supported: -n ldap Create or update the configuration information for the <i>destination</i> entry in LDAP printer contexts. ldap specifies that the information is written to an LDAP server. See NOTES. -x Remove all configuration for the <i>destination</i> entry from the database specified by the -n option. -a key=value Configure the specified <i>key=value</i> pair for the <i>destination</i>. -d key Delete the configuration option specified by <i>key</i> for the <i>destination</i> entry. -D binddn Use the distinguished name (DN) binddn to bind to the LDAP directory server. -w passwd Use passwd as the password for authentication to the LDAP directory server. -h ldaphost Specify an alternate host on which the LDAP server is running. This option is only used when ldap is specified as the naming service. If this option is not specified, the default is the current host system.
操作数	The following operand is supported: <i>destination</i> Specifies the entry in LDAP, in which to create or modify information. <i>destination</i> names a printer of class of printers. Specify <i>destination</i> using atomic names. POSIX-style destination names are not acceptable. See standards(5) for information regarding POSIX.
示例	示例 1 Setting a key=value Pair The following example sets the user-equivalence=<i>true</i> <i>key=value</i> pair for destination <i>tabloid</i> in the NIS context: <pre>example% lpset -n nis -a user-equivalence=true tabloid</pre> 示例 2 Setting a key=value Pair in LDAP <pre>example% lpset -n ldap -h ldap1.xyz.com -D "cn=Directory Manager" \ -w passwd -a key1=value1 printer1</pre>

退出状态

The following exit values are returned:

- 0 Successful completion.
- non-zero An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	print/lp/print-client-commands
Interface Stability	Committed

另请参见

[ldapclient\(1M\)](#), [lpget\(1M\)](#), [nsswitch.conf\(4\)](#), [attributes\(5\)](#), [ldap\(5\)](#), [standards\(5\)](#)
《Oracle Solaris Administration: Naming and Directory Services》

附注

If the ldap database is used, the printer administrator should be mindful of the following when updating printer information.

1. Because the domain information for the printer being updated is extracted from the [ldapclient\(1M\)](#) configuration, the LDAP server being updated must host the same domain that is used by the current [ldapclient\(1M\)](#) server.
2. If the LDAP server being updated is a replica LDAP server, the updates will be referred to the master LDAP server and completed there. The updates might be out of sync and not appear immediately, as the replica server may not have been updated by the master server. For example, a printer that you deleted by using lpset may still appear in the printer list you display with lpget until the replica is updated from the master. Replica servers vary as to how often they are updated from the master.
3. Although users can use the LDAP command line utilities [ldapadd\(1\)](#) and [ldapmodify\(1\)](#) to update printer entries in the directory, the preferred method is to use lpset. Otherwise, if the ldapadd and ldapmodify utilities are used, the administrator must ensure that the printer-name attribute value is unique within the ou=printers container on the LDAP server. If the value is not unique, the result of modifications done using lpset might be unpredictable.

引用名	luxadm – administer Sun Fire 880 storage subsystem and FC_AL devices
用法概要	<pre>luxadm [options]... subcommand [options]... enclosure [,dev] pathname...</pre>
描述	The luxadm program is an administrative command that manages the SENA, Sun Fire 880 internal storage subsystem, and individual Fiber Channel Arbitrated Loop (FC_AL) devices. luxadm performs a variety of control and query tasks depending on the command line arguments and options used. The command line must contain a subcommand. The command line may also contain options, usually at least one enclosure name or pathname, and other parameters depending on the subcommand. You need specify only as many characters as are required to uniquely identify a subcommand. Specify the device that a subcommand interacts with by entering a pathname. For the SENA subsystem, a disk device or enclosure services controller may instead be specified by entering the World Wide Name (WWN) for the device or a port to the device. The device may also be specified by entering the name of the SENA enclosure, and an optional identifier for the particular device in the enclosure. The individual FC_AL devices may be specified by entering the WWN for the device or a port to the device.
Pathname	Specify the device or controller by either a complete physical pathname or a complete logical pathname. For SENA, a typical physical pathname for a device is: <pre>/devices/sbus@1f,0/SUNW,socal@1,0/sf@0,0/ssd@w2200002037000f96, 0:a,raw</pre> For all SENA IBs (Interface Boards) and Sun Fire 880 SES device controllers on the system, a logical link to the physical paths is kept in the directory /dev/es. An example of a logical link is /dev/es/ses0. The WWN may be used in place of the pathname to select an FC_AL device, SENA subsystem IB, or Sun Fire 880 internal storage subsystem. The WWN is a unique 16 hexadecimal digit value that specifies either the port used to access the device or the device itself. A typical WWN value is: <pre>2200002037000f96</pre> See NOTES for more information on the WWN formats. For a disk in a Sun Fire 880 internal storage subsystem, a typical physical pathname is: <pre>/devices/pci@8,600000/SUNW,qlc@2/fp@0,0/ssd@w2100002037a6303c,0:a</pre> and a typical logical pathname is: <pre>/dev/rdisk/c2t8d0s2</pre>

For individual FC_AL devices, a typical physical pathname is:

```
/devices/sbus@3.0/SUNW,socal@d,10000/sf@0,0/ssd@w2200002037049fc3,0:a,raw
```

and a typical logical pathname is:

```
/dev/rdisk/c1t0d0s2
```

Enclosure

For SENA, a device may be identified by its enclosure name and slotname:

```
box_name[, fslot_number]
```

```
box_name[, rslot_number]
```

box_name is the name of the SENA enclosure, as specified by the *enclosure_name* subcommand. When used without the optional *slot_number* parameter, the *box_name* identifies the SENA subsystem IB.

f or r specifies the front or rear slots in the SENA enclosure.

slot_number specifies the slot number of the device in the SENA enclosure, 0-6 or 0-10.

For a Sun Fire 880 internal storage subsystem, a device may also be identified by its enclosure name and slot name. However, there is only one set of disks:

```
box_name[, sslot_number]
```

box_name is the name of the Sun Fire 880 enclosure, as specified by the *enclosure_name* subcommand. When used without the optional *slot_number* parameter, *box_name* identifies the Sun Fire 880 internal storage subsystem enclosure services device. Use *s* to specify the disk slot number in the Sun Fire 880 internal storage subsystem, 0 - 11.

See [disks\(1M\)](#) and [devlinks\(1M\)](#) for additional information on logical names for disks and subsystems.

选项

The following options are supported by all subcommands:

-e Expert mode. This option is not recommended for the novice user.

-v Verbose mode.

Options that are specific to particular subcommands are described with the subcommand in the USAGE section.

操作数

The following operands are supported:

enclosure

The *box_name* of the SENA or Sun Fire 880 internal storage subsystem.

fibre_channel_HBA_port

The path to the host controller port. A typical path is:

```
/devices/pci@8,600000/pci@1/SUNW,qlc@4/fp@0,0:devctl
```

pathname

The logical or physical path of a SENA IB, Sun Fire 880 internal storage subsystem, or disk device. *pathname* can also be the WWN of a SENA IB, SENA disk, or individual FC_AL device.

用法

Subcommands

```
display enclosure[,dev] ... | pathname ...
```

```
display -p pathname ...
```

```
display -r enclosure[,dev] ... | pathname ...
```

```
display -v enclosure[,dev] ... | pathname ...
```

Displays enclosure or device specific data.

Subsystem data consists of enclosure environmental sense information and status for all subsystem devices, including disks.

Disk data consists of inquiry, capacity, and configuration information.

- p Displays performance information for the device or subsystem specified by *pathname*. This option only applies to subsystems that accumulate performance information.
- r Displays error information for the FC_AL device specified by the *pathname*, or, if the path is a SENA, for all devices on the loop. The -r option only applies to SENA subsystems and individual FC_AL devices.
- v Displays in verbose mode, including mode sense data.

```
download [ -s ] [ -f filename_path ] enclosure. . .
```

Download the prom image pointed to the SENA subsystem Interface Board unit or the Sun Fire 880 internal storage subsystem specified by the enclosure or *pathname*.

When the SENA's download is complete, the SENA will be reset and the downloaded code executed. If no filename is specified, the default prom image will be used. The default prom image for the SENA is in the directory `usr/lib/locale/C/LC_MESSAGES` and is named `ibfirmware`

When the Sun Fire 880 internal storage subsystem's download is complete, the subsystem resets and the downloaded code begins execution. The default firmware image for the Sun Fire 880 internal storage subsystem is in:

```
/usr/platform/SUNW,Sun-Fire-880/lib/images/int_fcbpl_fw.
```

- s Save. The -s option is used to save the downloaded firmware in the FEPRM. If -s is not specified, the downloaded firmware will not be saved across power cycles.

The -s option does not apply to the Sun Fire 880 internal storage subsystem as it always stores downloaded firmware in the flash memory.

When using the `-s` option, the `download` subcommand modifies the FEPROM on the subsystem and should be used with *caution*.

`enclosure_name new_name enclosure | pathname`

Change the enclosure name of the enclosure or enclosures specified by the enclosure or pathname. The new name (*new_name*) must be 16 or less characters. Only alphabetic or numeric characters are acceptable. This subcommand applies only to the SENA and the Sun Fire 880 internal storage subsystem.

`failover primary | secondary pathname`

Select which Sun Storage T3 storage array partner group controller accesses a given logical volume. If `primary` is specified, the logical volume is accessed through the primary controller. If `secondary` is specified, the logical volume is accessed through the secondary controller specified by *pathname*.

`fcal_s_download [ -f fcode-file ]`

Download the fcode contained in the file *fcode-file* into *all* the FC100/S Sbus Cards. This command is interactive and expects user confirmation before downloading the fcode.

Use `fcal_s_download` *only* in single-user mode. Using `fcal_s_download` to update a host adapter while there is I/O activity through that adapter *will* cause the adapter to reset. Newly updated FCode will not be executed or visible until a system reboot.

`-f fcode-file` When invoked without the `-f` option, the current version of the fcode in each FC100/S Sbus card is printed.

`fcode_download -p`

`fcode_download -d dir-name`

Locate the installed FC/S, FC100/S, FC100/P, or FC100/2P host bus adapter cards and download the FCode files in *dir-name* to the appropriate cards. The command determines the correct card for each type of file, and is interactive. User confirmation is required before downloading the FCode to each device.

Use `fcode_download` to load FCode only in single-user mode. Using `fcode_download` to update a host adapter while there is I/O activity through that adapter causes the adapter to reset. Newly updated FCode will not be executed or visible until a system reboot.

`-d dir-name` Download the FCode files contained in the directory *dir-name* to the appropriate adapter cards.

`-p` Prints the current version of FCode loaded on each card. No download is performed.

`inquiry enclosure[,dev] ... | pathname ...`

Display the inquiry information for the selected device specified by the enclosure or pathname.

`insert_device [ enclosure,dev ... ]`

Assist the user in the hot insertion of a new device or a chain of new devices. Refer to NOTES for limitations on hotplug operations. This subcommand applies only to the SENA, Sun Fire 880 internal storage subsystem, and individual FC_AL drives. For the SENA, if more than one enclosure has been specified, concurrent hot insertions on multiple busses can be performed. With no arguments to the subcommand, entire enclosures or individual FC_AL drives can be inserted. For the SENA or the Sun Fire 880 internal storage subsystem, this subcommand guides the user interactively through the hot insertion steps of a new device or chain of devices. If a list of disks was entered it will ask the user to verify the list of devices to be inserted is correct, at which point the user can continue or quit. It then interactively asks the user to insert the disk(s) or enclosure(s) and then creates and displays the logical pathnames for the devices.

`led enclosure,dev ... | pathname ...`

Display the current state of the LED associated with the disk specified by the enclosure or pathname. This subcommand only applies to subsystems that support this functionality.

`led_blink enclosure,dev ... | pathname ...`

Requests the subsystem to start blinking the LED associated with the disk specified by the enclosure or pathname. This subcommand only applies to subsystems that support this functionality.

`led_off enclosure,dev ... | pathname ...`

Requests the subsystem to disable (turn off) the LED associated with the disk specified by the enclosure or pathname. On a SENA subsystem, this may or may not cause the LED to turn off or stop blinking depending on the state of the SENA subsystem. Refer to the SENA Array Installation and Service Manual (p/n 802-7573). This subcommand only applies to subsystems that support this functionality.

`led_on pathname ...`

Requests the subsystem to enable (turn on) the LED associated with the disk specified by the pathname. This subcommand only applies to subsystems that support this functionality.

`power_off [ -F ] enclosure[,dev] ... | pathname ...`

When a SENA is addressed, this subcommand causes the SENA subsystem to go into the power-save mode. The SENA drives are not available when in the power-save mode. When a drive in a SENA is addressed the drive is set to the drive off/unmated state. In the drive off/unmated state, the drive is spun down (stopped) and in bypass mode. This command does not apply to the Sun Fire 880 internal storage subsystem.

-F The force option only applies to the SENA. Instructs luxadm to attempt to power off one or more devices even if those devices are being used by this host (and are, therefore, busy).

Warning: Powering off a device which has data that is currently being used will cause unpredictable results. Users should attempt to power off the device normally (without -F) first, only resorting to this option when sure of the consequences of

overriding normal checks.

`power_on enclosure[,dev] . .`

Causes the SENA subsystem to go out of the power-save mode, when this subcommand is addressed to a SENA.. When this subcommand is addressed to a drive the drive is set to its normal start-up state. This command does not apply to the Sun Fire 880 internal storage subsystem.

`probe [-p]`

Finds and displays information about all attached SENA subsystems, Sun Fire 880 internal storage subsystems, and individual FC_AL devices, including the logical pathname, the WWNs, and enclosure names. This subcommand warns the user if it finds different SENAs with the same enclosure names.

`-p` Includes the physical pathname in the display.

`qlgc_s_download [-f fcode-file]`

Download the FCode contained in the file *fcode-file* into all the FC100/P, FC100/2P PCI host adapter cards. This command is interactive and expects user confirmation before downloading the FCode to each device. Only use `qlgc_s_download` in single-user mode. Using `qlgc_s_download` to update a host adapter while there is I/O activity through that adapter will cause the adapter to reset. Newly updated FCode will not be executed or visible until a system reboot.

`-f fcode-file` When invoked without the `-f` option, the current version of the FCode in each FC100/P, FC100/2P PCI card is printed.

`release pathname`

Release a reservation held on the specified disk. The pathname should be the physical or logical pathname for the disk.

This subcommand is included for historical and diagnostic purposes only.

`remove_device [-F] enclosure[,dev] ... | pathname ...`

Assists the user in hot removing a device or a chain of devices. This subcommand can also be used to remove entire enclosures. This subcommand applies to the SENA, Sun Fire 880 internal storage subsystem, and individual FC_AL drives. Refer to NOTES for limitations on hotplug operations. For the SENA, Sun Fire 880 internal storage subsystem, and individual FC_AL devices, this subcommand guides the user through the hot removal of a device or devices. During execution it will ask the user to verify the list of devices to be removed is correct, at which point the user can continue or quit. It then prepares the disk(s) or enclosure(s) for removal and interactively asks the user to remove the disk(s) or enclosure(s).

For Multi-Hosted disk, the steps taken are:

- Issue the `luxadm remove_device` command on the first host. When prompted to continue, wait.

- Issue the `luxadm remove_device` command on the secondary hosts. When prompted to continue, wait.
 - Continue with the `remove_device` command on the first host. Remove the device when prompted to do so.
 - Complete the `luxadm remove_device` command on the additional hosts.
- F Instructs `luxadm` to attempt to hot plug one or more devices even if those devices are being used by this host (and are, therefore, *busy* or *reserved*), to *force* the hotplugging operation.

Warning: Removal of a device which has data that is currently being used will cause unpredictable results. Users should attempt to hotplug normally (without -F) first, only resorting to this option when sure of the consequences of overriding normal hotplugging checks.

<code>reserve pathname</code>	Reserve the specified disk for exclusive use by the issuing host. The pathname used should be the physical or logical pathname for the disk.
	This subcommand is included for historical and diagnostic purposes only.
<code>set_boot_dev [-y] pathname</code>	Set the boot-device variable in the system PROM to the physical device name specified by <i>pathname</i> , which can be a block special device or the pathname of the directory on which the boot file system is mounted. The command normally runs interactively requesting confirmation for setting the default boot-device in the PROM. The -y option can be used to run it non-interactively, in which case no confirmation is requested or required.
<code>start pathname</code>	Spin up the specified disk(s) in a SENA.
<code>stop pathname...</code>	Spin down the specified disks in a SENA.

SENA, Sun Fire 880
Internal Storage
Subsystem, and
Individual FC_AL Drive
Expert Mode
Subcommands

The following subcommands are for expert use only, and are applicable only to the SENA, Sun Fire 880 internal storage subsystem, and fiber channel loops. They should only be used by users that are knowledgeable about the SENA subsystem and fiber channel loops.

If you specify a disk to an expert subcommand that operates on a bus, the subcommand operates on the bus to which the specified disk is attached.

-e bypass [-ab] *enclosure,dev*
-e bypass -f *enclosure*

Request the enclosure services controller to set the LRC (Loop Redundancy Circuit) to the bypassed state for the port and device specified.

	This subcommand supports the following options: - a Bypass port a of the device specified. - b Bypass port b of the device specified.
-e dump_map <i>fibre_channel_HBA_port</i>	Display WWN data for a target device or host bus adapter on the specified fibre channel port. If there are no target devices on the specified port, an error is returned.
-e enable [-ab] <i>enclosure,dev</i> -e enable -f <i>enclosure</i>	Request the enclosure services controller to set the LRC (Loop Redundancy Circuit) to the enabled state for the port and device specified.
	This subcommand supports the following options: - a Enable port a of the device specified. - b Enable port b of the device specified.
-e forcelp <i>enclosure[,dev] ... pathname ...</i>	Force the link to reinitialize, using the Loop Initialization Primitive (LIP) sequence. The enclosure or pathname can specify any device on the loop. Use the pathname to specify a specific path for multiple loop configurations.
	This is an expert only command and should be used with caution. It will reset all ports on the loop.
-e rdls <i>enclosure[,dev] ... pathname ...</i>	Read and display the link error status information for all available devices on the loop that contains the device specified by the enclosure or pathname.
Other Expert Mode Subcommands	See NOTES for limitations of these subcommands. They should only be used by users that are knowledgeable about the systems they are managing.

These commands do not apply to the Sun Fire 880 internal storage subsystem.

<code>-e bus_getstate <i>pathname</i></code>	Get and display the state of the specified bus.
<code>-e bus_quiesce <i>pathname</i></code>	Quiesce the specified bus.
<code>-e bus_reset <i>pathname</i></code>	Reset the specified bus only.
<code>-e bus_resetall <i>pathname</i></code>	Reset the specified bus and all devices.
<code>-e bus_unquiesce <i>pathname</i></code>	Unquiesce the specified bus. the specified device.
<code>-e dev_getstate <i>pathname</i></code>	Get and display the state of the specified device.
<code>-e dev_reset <i>pathname</i></code>	Reset the specified device.
<code>-e offline <i>pathname</i></code>	Take the specified device offline.
<code>-e online <i>pathname</i></code>	Put the specified device online.

示例

示例 1 Displaying the SENAs and Individual FC_AL Devices on a System

The following example finds and displays all of the SENAs and individual FC_AL devices on a system:

```
example% luxadm probe
```

示例 2 Displaying a SENA or Sun Fire 880 Internal Storage Subsystem

The following example displays a SENA or Sun Fire 880 internal storage subsystem:

```
example% luxadm display /dev/es/ses0
```

示例 3 Displaying Two Subsystems

The following example displays two subsystems using the enclosure names:

```
example% luxadm display BOB system1
```

示例 4 Displaying Information about the First Disk

The following example displays information about the first disk in the front of the enclosure named BOB. Use `f` to specify the front disks. Use `r` to specify the rear disks.

```
example% luxadm display BOB,f0
```

示例 5 Displaying Information on a Sun Fire 880 Internal Storage Subsystem

The Sun Fire 880 internal storage subsystem has only one set of disks. In this case, use `s` to specify the slot:

```
example% luxadm display BOB,s0
```

示例 6 Displaying Information about a SENA disk, an Enclosure, or an Individual FC_AL Drive

The following example displays information about a SENA disk, an enclosure, or an individual FC_AL drive with the port WWN of 2200002037001246:

```
example% luxadm display 2200002037001246
```

示例 7 Using Unique Characters to Issue a Subcommand

The following example uses only as many characters as are required to uniquely identify a subcommand:

```
example% luxadm disp BOB
```

示例 8 Displaying Error Information

The following example displays error information about the loop that the enclosure BOB is on:

```
example% luxadm display -r BOB
```

示例 9 Downloading New Firmware into the Interface Board

The following example downloads new firmware into the Interface Board in the enclosure named BOB (using the default path for the file to download):

```
example% luxadm download -s BOB
```

示例 10 Displaying Information from the SCSI Inquiry Command

The following example displays information from the SCSI inquiry command from all individual disks on the system, using only as many characters as necessary to uniquely identify the inquiry subcommand:

```
example% luxadm inq /dev/rdisk/c?t?d?s2
```

示例 11 Hotplugging

The following example hotplugs a new drive into the first slot in the front of the enclosure named BOB:

```
example% luxadm insert_device BOB,f0
```

The following example hotplugs a new drive into the first slot in the Sun Fire 880 internal storage subsystem named SF880-1:

```
example% luxadm insert_device SF880-1,s0
```

示例 12 Running an Expert Subcommand

The following example runs an expert subcommand. The subcommand forces a loop initialization on the loop that the enclosure BOB is on:

```
example% luxadm -e forcelp BOB
```

示例 13 Using the Expert Mode Hot Plugging Subcommands

An example of using the expert mode hot plugging subcommands to hot remove a disk follows. See NOTES for hot plugging limitations.

The first step reserves the SCSI device so that it can't be accessed by way of its second SCSI bus:

```
example# luxadm reserve /dev/rdisk/c1t8d0s2
```

示例 14 Taking the Disk to be Removed Offline

The next two steps take the disk to be removed offline then quiesce the bus:

```
example# luxadm -e offline /dev/rdisk/c1t8d0s2
example# luxadm -e bus_quiesce /dev/rdisk/c1t8d0s2
```

示例 15 Unquiescing the Bus

The user then removes the disk and continues by unquiescing the bus, putting the disk back online, then unreserving it:

```
example# luxadm -e bus_unquiesce /dev/rdisk/c1t8d0s2
example# luxadm -e online /dev/rdisk/c1t8d0s2
example# luxadm release /dev/rdisk/c1t8d0s2
```

环境变量

See [environ\(5\)](#) for a description of the LANG environment variable that affects the execution of luxadm.

退出状态

The following exit values are returned:

- 0 Successful completion.
- 1 An error occurred.

文件

- usr/lib/firmware/fc_s/fc_s_fcode
- usr/lib/locale/C/LC_MESSAGES/ibfirmware

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

usr/sbin

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/storage/fc-utilities

另请参见

[devlinks\(1M\)](#), [disks\(1M\)](#), [attributes\(5\)](#), [environ\(5\)](#), [ses\(7D\)](#)

《SAN Configuration and Multipathing Guide》

Oracle Solaris Cluster 3.3 with Fibre Channel JBOD Storage Device Manual

附注

Refer to *Tutorial for SCSI use of IEEE Company_ID*, R. Snively, for additional information regarding the IEEE extended WWN. Currently, only some device drivers support hot plugging. If hot plugging is attempted on a disk or bus where it is not supported, an error message of the form:

```
luxadm: can't acquire "PATHNAME": No such file or directory
```

...will be displayed.

You must be careful not to quiesce a bus that contains the root or the /usr filesystems or any swap data. If you do quiesce such a bus a deadlock can result, requiring a system reboot.

参 考 文 档

系统管理命令—第2单元

引用名	mail.local – store mail in a mailbox
用法概要	<code>/usr/lib/mail.local [-f <i>sender</i>] [-d] <i>recipient</i></code>
描述	<code>mail.local</code> reads the standard input up to an end-of-file and appends it to each user's mail file (mailbox). This program is intended to be used by sendmail(1M) as a mail delivery agent for local mail. It is not a user interface agent. Messages are appended to the user's mail file in the <code>/var/mail</code> directory. The user must be a valid user name. Each delivered mail message in the mailbox is preceded by a "Unix From line" with the following format: <pre>From <i>sender_address</i> <i>time_stamp</i></pre> The <i>sender_address</i> is extracted from the SMTP envelope address (the envelope address is specified with the <code>-f</code> option). A trailing blank line is also added to the end of each message. The mail files are locked with a <code>.lock</code> file while mail is appended. The mail files are created with mode <code>660</code>, owner is set to <i>recipient</i>, and group is set to <code>mail</code>. If the "biff" service is returned by getservbyname(3SOCKET), the biff server is notified of delivered mail. This program also computes the Content-Length: header which will be used by the mailbox reader to mark the message boundary.
选项	The following options are supported: <code>-f <i>sender</i></code> Specifies the "envelope from address" of the message. This flag is technically optional, but should be used. <code>-d</code> Specifies the recipient of the message. This flag is also optional and is supported here for backward compatibility. That is, <code>mail.local <i>recipient</i></code> is the same as <code>mail.local -d <i>recipient</i></code>. <code>-l</code> Turn on LMTP mode. <code>-r <i>from</i></code> Specify the sender's name (for backward compatibility). <code>-7</code> Do not advertise 8BITMIME support in LMTP mode. <code>-b</code> Return a permanent error instead of a temporary error if a mailbox exceeds quota.
操作数	The following operand is supported: <i>recipient</i> The recipient of the mail message.

环境变量 TZ Used to set the appropriate time zone on the timestamp.

退出状态 The following exit values are returned:

0 Successful operation.

>0 An error occurred.

文件 /tmp/local.XXXXXX temporary files
/tmp/lochd.XXXXXX temporary files
/var/mail/user_name user's mail file

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/smtp/sendmail

另请参见 [mail\(1\)](#), [comsat\(1M\)](#), [sendmail\(1M\)](#), [getservbyname\(3SOCKET\)](#), [attributes\(5\)](#)

引用名

makedbm – make a dbm file, or get a text file from a dbm file

用法概要

```
makedbm [-b] [-l] [-s] [-E] [-i yp_input_file]
        [-o yp_output_name] [-d yp_domain_name]
        [-m yp_master_name] [-S delimiter]
        [-D number_of_delimiters] infile outfile

makedbm [-u dbmfilename]
```

描述

The makedbm utility takes the *infile* and converts it to a pair of files in ndbm format (see [ndbm\(3C\)](#)), namely *outfile.pag* and *outfile.dir*. Each line of the input file is converted to a single dbm record. All characters up to the first TAB or SPACE form the key, and the rest of the line is the data. If a line ends with ‘\’ (backslash), the data for that record is continued on to the next line. makedbm does not treat ‘#’ (pound-sign) as a special character.

Because makedbm is mainly used in generating dbm files for the NIS name service, it generates a special entry with the key *yp_last_modified*, which is the date of *infile* (or the current time, if *infile* is ‘-’). The entries that have keys with the prefix *yp_* are interpreted by NIS server utilities.

选项

The following options are supported:

- b Insert the YP_INTERDOMAIN into the output. This key causes [ypserv\(1M\)](#) to use DNS for host name and address lookups for hosts not found in the maps.
- d *yp_domain_name* Create a special entry with the key *yp_domain_name*.
- D *number_of_delimiters* Specify *number_of_delimiters* to skip before forming the key.
- E Delimiters are escaped.
- i *yp_input_file* Create a special entry with the key *yp_input_file*.
- l Lower case. Convert the keys of the given map to lower case, so that, for example, host name matches succeed independent of upper or lower case distinctions.
- m *yp_master_name* Create a special entry with the key *yp_master_name*. If no master host name is specified, *yp_master_name* is set to the local host name.
- o *yp_output_name* Create a special entry with the key *yp_output_name*.
- s Secure map. Accept connections from secure NIS networks only.
- S *delimiter* Specify the *delimiter* to use instead of the default delimiter for forming the key.
- u *dbmfilename* Undo a dbm file. Prints out the file in text format, one entry per line, with a single space separating keys from values.

操作数 The following operands are supported:

infile Input file for makedbm. If *infile* is '-' (dash), the standard input is read.

outfile One of two output files in ndbm format: *outfile.pag* and *outfile.dir*.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [ypserv\(1M\)](#), [ndbm\(3C\)](#), [attributes\(5\)](#)

引用名 makemap – create database maps for sendmail

用法概要 makemap [-N] [-d] [-f] [-o] [-r] [-s] [-v] [-C *file*]
[-c *cache*size] [-D *comment*char] [-e] [-l] [-t *delim*]
[-u] *mantype mapname*

描述 makemap creates the database maps used by the keyed map lookups in [sendmail\(1M\)](#). makemap reads from the standard input and outputs to the specified *mapname*.

In all cases, makemap reads lines from the standard input consisting of two words separated by whitespace. The first is the database key, the second is the value. The value may contain *%n* strings to indicated parameter substitution. Literal percents should be doubled (%%). Blank lines and lines beginning with # are ignored.

makemap handles three different database formats. Database format is selected using the *maptype* parameter. See OPERANDS.

选项 The following options are supported:

- c *cache*size Use the specified hash and B-Tree cache size (*cache*size).
- C *file* Use the specified sendmail configuration file (*file*) for looking up the TrustedUser option.
- d Allow duplicate keys in the map. This is only allowed on B-Tree format maps. If two identical keys are read, both be inserted into the map.
- D *comment*char Use the specified character to indicate a comment (which is ignored) instead of the default of '#'.
- e Allow empty value (right hand side).
- f Normally, all upper case letters in the key are folded to lower case. This flag disables that behavior. This is intended to mesh with the -f flag in the K line in `sendmail.cf`. The value is never case folded.
- l List supported map types.
- N Include the null byte that terminates strings in the map. This must match the -N flag in the K line in `sendmail.cf`
- o Append to an old file. This allows you to augment an existing file.
- r Allow replacement of existing keys. Normally makemap complains if you repeat a key, and does not do the insert.
- s Ignore safety checks on maps being created. This includes checking for hard or symbolic links in world writable directories.
- t *delim* Use the specified delimiter (*delim*) instead of whitespace.

- u Dump (unmap) the content of the database to standard output. Note that, if the -t option is also provided, the specified delimiter is used when the content is dumped instead of whitespace.
- v Verbosely print what it is doing.

操作数

The following operands are supported:

mapname File name of the database map being created.

maptype Specifies the database format. The following *maptype* parameters are available:

- dbm Specifies DBM format maps.
- btree Specifies B-Tree format maps.
- hash Specifies hash format maps.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/smtp/sendmail

另请参见

[editmap\(1M\)](#), [sendmail\(1M\)](#), [attributes\(5\)](#)

引用名	makeuuid – generate Universal Unique Identifiers
用法概要	makeuuid [-e <i>ether</i>] [-n <i>count</i>] [-R <i>root</i>]
描述	The makeuuid command generates UUIDs (Universal Unique Identifiers) conforming to the OSF DCE specification for UUIDs. The specification states: “A UUID is an identifier that is unique across both space and time, with respect to the space of all UUIDs. A UUID can be used for multiple purposes, from tagging objects with an extremely short lifetime, to reliably identifying very persistent objects across a network. “The generation of UUIDs does not require a registration authority for each single identifier. Instead, it requires a unique value over space for each UUID generator. This spatially unique value is [normally] specified as an IEEE 802 address, which is usually already applied to network-connected systems.” The makeuuid command generates one or more UUIDs on the standard output.
选项	The makeuuid command supports the following options: -e <i>ether</i> Supplies an alternate address to be used in the generation of the UUIDs. Normally, the system's Ethernet address is acquired and used during the generation of a UUID. However, this requires root privileges to open and read the network devices. If this is not possible, you must supply an alternate Ethernet address. -n <i>count</i> Generate multiple UUIDs. This option generates the specified number of UUIDs, one per line. Using this form is more efficient than, and functionally equivalent to, calling the makeuuid command multiple times. This can be used, for example, when a large number of UUIDs need to be generated for a given application. -R <i>root</i> Use <i>root</i> as the root filesystem path when updating the shared state file (see FILES). The shared state file must be writable by the user running makeuuid, otherwise no UUIDs will be generated and the command will return in failure. 注 – The root file system of any non-global zones must not be referenced with the -R option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See zones(5).
用法	Normally, you run the makeuuid command with root privileges, as the Ethernet address and state files can be easily accessed and updated. If this is not possible, you must use the -R and -e options to specify an alternate root and Ethernet address to use when calculating the UUIDs.
示例	示例 1 Generating Multiple UUIDs The following command generates 3000 UUIDs: <pre>example# makeuuid -n 3000</pre>

示例 2 Invoking Without Root Privileges

If you cannot obtain root privileges, you must specify an alternate Ethernet address and state file location:

```
example% makeuuid -e 11:22:33:44:55:66 -R /export/root/example2
```

See the caveat on the use of the -R option in the description of that option, above.

退出状态

The following exit values are returned:

- 0 Successful completion.
- 1 Out of memory.
- 1 Invalid Ethernet address given or access denied.

文件

`/var/sadm/system/uuid_state` UUID state file. Use of time values is one way that UUID generators, such as `makeuuid`, guarantee uniqueness. A state file is a mechanism that allows `makeuuid` to “remember” the last time value it used so it can increment that value for use in a new UUID. See the Internet Draft UUIDs and GUIDs, dated February 4, 1998, for details on the state file mechanism.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/management/product-registry

另请参见

[Intro\(3\)](#), [libwsreg\(3LIB\)](#), [attributes\(5\)](#)

附注

The formal UUID specification is in the OSF DCE specification, available at www.opengroup.org. As of the date of publication of this man page, a copy of the specification is available at:

```
http://www.opengroup.org/onlinepubs/9629399/apdxa.htm
```

Sun has no control over the availability of documents on the www.opengroup.org web site.

引用名

masfcnv – SNMP configuration migration script

用法概要

```
/usr/lib/sma_snmp/masfcnv [-cimnrs] [-l agentmaster]
    [-p enabledisableerror] [-t noneadd]
    [-u agentmastererror] [-y agentmastererror]
```

```
masfcnv [-V]
```

```
masfcnv [-?]
```

描述

The `masfcnv` script is used to assist the system administrator in migrating an existing set of configuration files for the Sun SNMP Management Agent for Sun Fire and Netra Systems (MASF) to the Systems Management Agent (SMA).

The script accepts as input the currently installed set of MASF and SMA configuration files and outputs a new set of SMA configuration files. Existing SMA configuration files are backed up by appending `.bak` to the filename. The administrator can choose to output the new configuration to standard output, instead of replacing the current configuration, by specifying the `-n` option.

The migration script must be run as the superuser. Failure to do so causes the script to exit with an error message. Before running the script you should ensure that both the SMA and MASF agents are not running. If the agents are running they will be shut down by the script.

The migration script installs a new startup script for the MASF agent in `/etc/init.d`, as well as a backup of the old script. During migration, MASF will be configured as an AgentX subagent of SMA. All migration settings will be migrated to the SMA configuration file.

The migration script aborts if any unrecognized directives are found in either the MASF configuration files or the SMA configuration files. This can be overridden with the `-i` option. If this option is selected, the behavior is to retain unrecognized directives that were present in the SMA configuration, but remove those present in the MASF configuration.

The migration script then proceeds to migrate access control and trap configuration. As a side effect of running the migration script, the following directives might be expanded by the script into multiple directives with an equivalent interpretation:

- `rwcommunity`
- `rocommunity`
- `rwuser`
- `rouser`
- `trapcommunity`
- `trapsink`
- `trap2sink`
- `informsink`

Access Control Migration

Access control directives are expanded into the equivalent `com2sec`, `group`, `access` and `view` directives. Existing group names are renamed by prepending a prefix to avoid conflict with any which may already be defined in SMA.

When migrating SNMPv1 or v2c access control, a conflict can occur if both MASF and SMA configuration files have defined access permissions for the same community and source address. The default behavior is to abort with a message, unless a use of the `-y` option specifies otherwise. If `-y agent` is specified then the MASF configuration takes precedence. If `-y master` is specified then the SMA configuration is retained.

When migrating USM configuration (SNMPv3), a conflict can occur if both SMA and MASF configurations define a user with the same `securityName`. If this occurs, the behavior of the script is determined by the `-u` option. If `-u agent` is specified, the configuration of the user defined in the MASF configuration files is the one that is retained. Otherwise, if the `-u master` option is specified, the use defined in the SMA configuration files is retained.

By default, the migration script attempts to migrate USM users from MASF to SMA. The script determines whether there are any SNMPv3 users present in the SMA configuration and whether the default `engineID` has been overridden in the SMA configuration files. If neither of these conditions obtain, then the any `usmUser` statements containing localized authentication keys can be migrated to SMA, along with the MASF `engineID`. This results in the `engineID` of the SMA master agent changing.

If the script determines that there are existing SNMPv3 users or a manually configured `engineID` present in the SMA configuration, only those users defined in `createUser` statements are transferred. Those users that were defined in `usmUser` statements are transferred but will have their passwords reset to a random value. You should notify your users of their new password or reset the password yourself by editing the newly-generated configuration file.

Trap/Inform Migration

The migration script performs a check to determine whether a trap destination defined for MASF is already specified in an existing SMA `trapsink`, `trap2sink` or `informsink` directive. If this is the case, then the directive in the MASF configuration will be discarded to avoid duplicate traps/informs being received.

`trapsink`, `trap2sink` and `informsink` directives specified in the existing SMA configuration are considered valid destinations for MASF traps/informs and will receive them from the MASF subagent after migration.

If the `-t none` option was specified on the command line, the migration script carries over any remaining MASF trap/inform directives without modification.

If the `-t add` option was specified (the default), the migration script expands any `trapsink`, `trap2sink`, or `informsink` directives to use the `TARGET-MIB` and `NOTIFICATION-MIB`. The `TARGET-MIB` specifies targets using IP addresses, so it might be desirable to use the `-t none` option if, for example, the network allocates IP addresses to hostnames dynamically by means of DHCP.

The expanded directives defines filters specific to the MASF agent so that traps from other subagents will not be received by migrated trap destinations. Existing filters present in the SMA configuration are, by default, not modified and might or might not receive MASF traps, depending upon the filters that were originally defined for them.

If the `-l` option is specified, any filters already defined in the `TARGET-MIB` and the `NOTIFICATION-MIB` for SMA are extended to include traps from MASF. In the event that a trap destination is already configured in the `TARGET-MIB` with the same target address and community as an existing MASF trap/inform sink, a conflict will arise.

If `-l` agent was specified and a conflict arises, the migration script uses the target SNMP parameters (that is, the SNMP version and choice of trap/inform) defined by the MASF `trap/informsink` directive to send traps to this destination. Otherwise, if the `-l master` option was specified, the conflict will be resolved using the target SNMP parameters specified in the SMA configuration.

Miscellaneous If the migration script encounters in the MASF configuration file any of the directives listed below and the directives are either not present or differ from the SMA configuration, the script will log a warning message.

- `syslocation`
- `syscontact`
- `sysname`
- `syssservices`
- `agentgroup`
- `agentuser`
- `authtrapeenable`

选项 The following options are supported:

<code>-?</code>	
<code>--help</code>	Displays usage information.
<code>-c</code>	
<code>--no-community</code>	Do not transfer v1/v2c communities.
<code>-i</code>	
<code>--ignore-unrecognized-directives</code>	Continue processing if unrecognized directives are present.
<code>-l agent master</code>	
<code>--master-trap-target=agent master</code>	If agent is specified, the existing SMA trap targets will be configured to receive traps that were previously sent to destinations for the Sun Fire SNMP agent. If master is specified, the targets will be configured to receive Sun Fire SNMP traps, but existing SNMP target parameters will be used.

<code>-m</code>	
<code>--no-usmuser</code>	Do not transfer usm (v3) users.
<code>-n</code>	
<code>--dry-run</code>	Run the migration without modifying any files. If an error arises, continue processing. This can be used to determine the likely migration issues.
<code>-p enable disable error</code>	
<code>--use-agent-port=enable disable error</code>	Indicates whether the port originally used by the Sun Fire SNMP agent should be used by the SMA agent after migration (if the two agents are using different ports). If <code>enable</code> is specified, then the port used by the Sun Fire SNMP agent will also be used by the SMA agent after migration. If <code>disable</code> is specified, the ports used by SMA will not be updated by the migration tool. If the <code>error</code> option is specified and the SMA agent is not already using the same ports as those used by the original Sun Fire SNMP agent, an error is reported and the migration process is terminated. If no option is specified the default behavior is equivalent to the <code>error</code> flag.
<code>-r</code>	
<code>--no-trap</code>	Do not transfer trap destinations.
<code>-s</code>	
<code>--skip-user</code>	If a user is found in the MASF configuration file that cannot be created in the new configuration because of a change in the engine ID, then output a message indicating that the user could not be migrated (needs to be manually recreated) and continue processing. If this option is not present, the migration tool will consider such a situation as an error and abort.
<code>-t none add</code>	
<code>--trap-filter=none add</code>	If <code>none</code> is specified then the script will copy trap directives directly. The administrator might need to manually update the

-u agent | master | error
--select-user=agent | master | error

configuration file to ensure traps are only delivered to their intended destinations. If add is specified, trap filters will be constructed so that traps originating from the original Sun Fire SNMP agent are delivered only to the destinations that originally received them. The default behavior is add.

Specifies that if a user with the same name is found in both configuration files that the conflict is to be resolved using the specified configuration file as input. Selecting a user from a particular file will also cause the group declaration for that user to be taken from the same file. If agent is specified then the user will be taken from the configuration file for the Sun Fire SNMP Agent. If master is specified, the user will be taken from the SMA configuration. Otherwise, if error is given, the script will terminate. If this option is not present, the default behavior is equivalent to the error flag.

-V
--version
-y agent | master | error
--select-community=agent | master | error

Display the version of this script.

Specifies that if a community with the same name is found in both configuration files that the conflict is to be resolved using the specified configuration file as input. If agent is specified then the community will be taken from the configuration file for the Sun Fire SNMP Agent. If master is specified, the community will be taken from the SMA configuration. Otherwise, if error is given, the script will terminate. If this option is not present, the default behavior is equivalent to the error flag.

示例 **示例 1 Simplest Case**

The command shown below is appropriate for a simple migration. The migration fails if there are any potential conflicts.

```
# masfcnv
```

示例 2 Migrating Such That MASF Settings Override

To migrate the MASF configuration such that it will always succeed, that MASF settings will override in the event of a conflict with SMA, and that access will still be provided on the original MASF port, enter:

```
# masfcnv -is -l agent -p enable -u agent -y agent
```

示例 3 Dry Run, Retaining SMA Settings

To attempt a dry run and migrate the configuration such that any conflicts will be resolved by retaining existing SMA settings, enter:

```
masfcnv -l master -u master -y master
```

退出状态 0 Success.

non-zero A problem occurred during migration.

文件

/etc/sma/snmp/snmpd.conf
/var/sma_snmp/snmpd.conf SMA configuration files

/etc/opt/SUNWmasf/conf/snmpd.conf
/var/opt/SUNWmasf/snmpd.dat MASF configuration files

/tmp/sma_migration.log masfcnv log file

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWsmcmd
Interface Stability	Committed

另请参见 [attributes\(5\)](#)

附注 The former path to this utility, /usr/sfw/lib, is now a link to /usr/lib.

引用名 mdlogd – Solaris Volume Manager daemon

用法概要 mdlogd

描述 mdlogd implements a simple daemon that watches the system console looking for messages written by the Solaris Volume Manger. When a Solaris Volume Manager message is detected, mdlogd sends a generic SNMP trap.

To enable traps, you must configure mdlogd into the SNMP framework. See [《Solaris Volume Manager Administration Guide》](#).

用法 mdlogd implements the following SNMP MIB:

```
SOLARIS-VOLUME-MGR-MIB DEFINITIONS ::= BEGIN
    IMPORTS
        enterprises FROM RFC1155-SMI
        DisplayString FROM SNMPv2-TC;

    -- Sun Private MIB for Solaris Volume Manager

    sun          OBJECT IDENTIFIER ::= { enterprises 42 }
    sunSVM       OBJECT IDENTIFIER ::= { sun 104 }

    -- this is actually just the string from /dev/log that
    -- matches the md: regular expressions.
    -- This is an interim SNMP trap generator to provide
    -- information until a more complete version is available.

    -- this definition is a formalization of the old
    -- Solaris DiskSuite mdlogd trap mib.

    svmOldTrapString OBJECT-TYPE
        SYNTAX DisplayString (SIZE (0..255))
        ACCESS read-only
        STATUS mandatory
        DESCRIPTION
            "This is the matched string that
            was obtained from /dev/log."
    ::= { sunSVM 1 }

    -- SVM Compatibility ( error trap )

    svmNotice     TrapTRAP-TYPE
        ENTERPRISE sunSVM
        VARIABLES { svmOldTrapString }
        DESCRIPTION
            "SVM error log trap for NOTICE.
            This matches 'NOTICE: md:'"
```

```
 ::= 1

svmWarningTrap TRAP-TYPE
    ENTERPRISE sunSVM
    VARIABLES { svmOldTrapString }
    DESCRIPTION
        "SVM error log trap for WARNING..
        This matches 'WARNING: md:'"

 ::= 2

svmPanicTrap TRAP-TYPE
    ENTERPRISE sunSVM
    VARIABLES { svmOldTrapString }
    DESCRIPTION
        "SVM error log traps for PANIC..
        This matches 'PANIC: md:'"

 ::= 3

END
```

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWlvma, SUNWlvmr
Interface Stability	Obsolete

另请参见 [snmpdx\(1M\)](#), [attributes\(5\)](#)
《Solaris Volume Manager Administration Guide》

引用名	mdmonitord – daemon to monitor metadevices				
用法概要	<code>/usr/sbin/mdmonitord [-t <i>time_interval</i>]</code>				
描述	The <code>mdmonitord</code> utility is part of Solaris Volume Manager. It monitors and checks RAID1 (mirrors), RAID5 and hot spares. There are two methods for checking: ▪ At fixed time intervals. ▪ When a RAID-1 (mirror), RAID-5, or hot spare fails. A failure generates an error event which triggers a check of these metadevices.				
选项	The following options are supported: -t Time interval in seconds. The default value is 0, which causes probes to occur only upon an error. If you want to run <code>mdmonitord</code> at a regular interval, a value of 1800 (seconds, every half hour) is recommended as a starting point.				
退出状态	The following exit values are returned: 0 Successful completion. >0 An error occurred.				
属性	See attributes(5) for descriptions of the following attributes:				
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>storage/svm</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	storage/svm
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	storage/svm				
另请参见	svcs(1), metaclear(1M), metadb(1M), metadetach(1M), metahs(1M), metainit(1M), metaoffline(1M), metaonline(1M), metaparam(1M), metarecover(1M), metarename(1M), metareplace(1M), metaset(1M), metassist(1M), metastat(1M), metasync(1M), metattach(1M), svcadm(1M), md.tab(4), md.cf(4), mddb.cf(4), md.tab(4), attributes(5), smf(5), md(7D) 《Solaris Volume Manager Administration Guide》				
附注	Since frequent probes can affect performance, it is recommended that the intervals between probes be limited. The <code>mdmonitord</code> service is managed by the service management facility, smf(5), under the service identifier: <pre>svc:/system/mdmonitor</pre> Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M). The service's status can be queried using the svcs(1) command.				

引用名	mdnsd – 多播 DNS 守护进程
用法概要	mdnsd [-debug]
描述	mdnsd 是多播 DNS 和 DNS 服务发现的守护程序。 mdnsd 守护进程在 UDP 端口 5353 侦听多播 DNS 查询数据包。当它接收到知道答案的查询时，mdnsd 将发出相应的多播 DNS 答复数据包。 mdnsd 守护进程还代表客户机进程执行多播 DNS 查询，并维护答复的高速缓存。 mdnsd 守护进程没有用户可指定的命令行参数，且用户不应手动运行 mdnsd。mdnsd 守护进程应由 Solaris 管理工具 (Solaris Management Facility, SMF) 和 svcadm(1M) 命令使用以下故障管理资源标识符 (Fault Management Resource Identifier, FMRI) 管理： <pre>svc:/network/dns/multicast:default</pre> 要检查 mdnsd 的内部状态以用于调试和诊断目的，可向其发送一个 SIGUSR1 信号，之后它将使用 syslog(3C) 工具记录其内部状态的快照摘要。mdnsd 使用系统日志工具代码守护进程和信息优先级。
选项	mdnsd 守护进程可识别以下选项： -debug 调试模式。mdnsd 守护进程将输出发送至标准错误，并且不在后台运行。此选项仅用于调试该守护进程。
配置	可以通过指定 mdns 作为名称服务转换配置文件 nsswitch.conf(4) 中的主机和 IP 节点的源，使用多播 DNS 查找主机名和主机地址。使用多播 DNS 查询主机名和主机地址的配置选项存储在 SMF 信息库中。此配置可通过 svccfg(1M) 命令使用以下故障管理资源标识符 (FMRI) 进行修改： <pre>svc:/network/dns/multicast:default</pre> 使用多播 DNS 查询主机名和主机地址的配置选项存储在名为 "nss_mdns_config" 的属性组中。构成配置选项的属性如下： search 主机名查找的搜索域列表。缺省情况下，配置中不包含任何搜索域。搜索列表当前限制为六个域。 valid 在使用多播 DNS 查找域的主机名或主机地址之前检查的有效域列表。搜索列表中指定的域始终包含在有效列表中。有效域列表当前限制为十个域。 请注意，以上配置选项是 "Volatile"（可变），并且在未来发行版中可能会发生更改。
文件	/usr/lib/inet/mdnsd
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	service/network/dns/mdns
接口稳定性	Volatile（可变）

另请参见

[dns-sd\(1M\)](#)、[svcadm\(1M\)](#)、[svccfg\(1M\)](#)、[syslog\(3C\)](#)、[nsswitch.conf\(4\)](#)、[attributes\(5\)](#)

有关多播 DNS 的信息，请参见 <http://www.multicastdns.org/>。

有关 DNS 服务发现的信息，请参见 <http://www.dns-sd.org/>。

引用名	medstat – check the status of mediator hosts for a given diskset
用法概要	<code>/usr/sbin/medstat [-q] -s setname</code>
描述	If a specified diskset has been configured for mediators, <code>medstat</code> attempts to contact these hosts to see if they are accessible and returns the results of the communication.
选项	-q This optional argument disables the printing of informative text. When used with <code>-q</code>, <code>medstat</code> still prints error messages and returns a result code. -s setname Specifies the name of a diskset on which <code>medstat</code> will work.
示例	示例 1 Checking diskset This example checks the mediator hosts for the selected diskset. # medstat -s relo-red The name of the diskset is <code>relo-red</code>. The <code>medstat</code> command prints the status for each mediator host. Additionally, if the mediator quorum is met, either through a “golden” mediator host or because half+1 of the mediator hosts respond, the exit code is 0. If the quorum is not met, then the exit code is 1. If no mediator hosts have been configured for the named diskset, the exit code is 2. The status field will contain one of the following values: Unreachable, Bad, Fatal, or Ok, where Unreachable indicates an RPC/communication problem, Bad indicates an error in the mediator data, Fatal indicates any other error condition, and Ok indicates no error conditions.
文件	/etc/lvm/meddb Contains the mediator data for a host that has been selected as a mediator host.
退出状态	The following exit values are returned: 0 Successful completion. >0 An error occurred.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见 [metaclear\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metahs\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metastat\(1M\)](#), [metasync\(1M\)](#), [metattach\(1M\)](#), [md.cf\(4\)](#), [md.tab\(4\)](#), [mddb.cf\(4\)](#), [meddb\(4\)](#), [mediator\(7D\)](#)

Sun Cluster documentation, Solaris Volume Manager Administration Guide

This command is designed for use in the high availability product.

引用名	metaclear – delete active metadevices and hot spare pools												
用法概要	<pre> /usr/sbin/metaclear -h /usr/sbin/metaclear [-s setname] -a [-f] /usr/sbin/metaclear component /usr/sbin/metaclear [-s setname] [-f] metadvice... hot_spare_pool... /usr/sbin/metaclear [-s setname] -r [-f] metadvice... hot_spare_pool... /usr/sbin/metaclear [-s setname] -p component /usr/sbin/metaclear [-s setname] -p metadvice </pre>												
描述	The metaclear command deletes the specified metadvice or <i>hot_spare_pool</i>, or purges all soft partitions from the designated component. Once a metadvice or hot spare pool is deleted, it must be re-created using metainit before it can be used again. Any metadvice currently in use (open) cannot be deleted.												
选项	Root privileges are required for all of the following options except -h. <table> <tr> <td>-a</td><td>Deletes all metadevices and configured hot spare pools in the set named by -s, or the local set by default.</td></tr> <tr> <td>-f</td><td>Deletes (forcibly) a metadvice that contains a subcomponent in an error state.</td></tr> <tr> <td>-h</td><td>Displays usage message.</td></tr> <tr> <td>-p</td><td>Deletes (purges) all soft partitions from the specified metadvice or component.</td></tr> <tr> <td>-r</td><td>Recursively deletes specified metadevices and hot spare pools, but does not delete metadevices on which others depend.</td></tr> <tr> <td>-s setname</td><td>Specifies the name of the diskset on which metaclear will work. Using the -s option causes the command to perform its administrative function within the specified diskset. Without this option, the command performs its function on local metadevices and/or hot spare pools.</td></tr> </table>	-a	Deletes all metadevices and configured hot spare pools in the set named by -s, or the local set by default.	-f	Deletes (forcibly) a metadvice that contains a subcomponent in an error state.	-h	Displays usage message.	-p	Deletes (purges) all soft partitions from the specified metadvice or component.	-r	Recursively deletes specified metadevices and hot spare pools, but does not delete metadevices on which others depend.	-s setname	Specifies the name of the diskset on which metaclear will work. Using the -s option causes the command to perform its administrative function within the specified diskset. Without this option, the command performs its function on local metadevices and/or hot spare pools.
-a	Deletes all metadevices and configured hot spare pools in the set named by -s, or the local set by default.												
-f	Deletes (forcibly) a metadvice that contains a subcomponent in an error state.												
-h	Displays usage message.												
-p	Deletes (purges) all soft partitions from the specified metadvice or component.												
-r	Recursively deletes specified metadevices and hot spare pools, but does not delete metadevices on which others depend.												
-s setname	Specifies the name of the diskset on which metaclear will work. Using the -s option causes the command to perform its administrative function within the specified diskset. Without this option, the command performs its function on local metadevices and/or hot spare pools.												
操作数	<table> <tr> <td>metadvice ...</td><td>Specifies the name(s) of the metadvice(s) to be deleted.</td></tr> <tr> <td>component</td><td>Specifies the c*d*t*s* name(s) of the components containing soft partitions to be deleted.</td></tr> <tr> <td>hot_spare_pool ...</td><td>Specifies the name(s) of the hot spare pools to be deleted. Names for hot spare pools can be any legal file name that is composed of alphanumeric characters, a dash (“-”), an underscore (“_”), or a period (“.”). Names must begin with a letter. The words “all” and “none” are reserved and cannot be used.</td></tr> </table>	metadvice ...	Specifies the name(s) of the metadvice(s) to be deleted.	component	Specifies the c*d*t*s* name(s) of the components containing soft partitions to be deleted.	hot_spare_pool ...	Specifies the name(s) of the hot spare pools to be deleted. Names for hot spare pools can be any legal file name that is composed of alphanumeric characters, a dash (“-”), an underscore (“_”), or a period (“.”). Names must begin with a letter. The words “all” and “none” are reserved and cannot be used.						
metadvice ...	Specifies the name(s) of the metadvice(s) to be deleted.												
component	Specifies the c*d*t*s* name(s) of the components containing soft partitions to be deleted.												
hot_spare_pool ...	Specifies the name(s) of the hot spare pools to be deleted. Names for hot spare pools can be any legal file name that is composed of alphanumeric characters, a dash (“-”), an underscore (“_”), or a period (“.”). Names must begin with a letter. The words “all” and “none” are reserved and cannot be used.												

示例

示例 1 Deleting Various Devices

The following example deletes a metadvice named d10.

```
# metaclear /dev/md/dsk/d10
```

The following example deletes all local metadevices and hot spare pools on the system.

```
# metaclear -a
```

The following example deletes a mirror, mymirror, with a submirror in an error state.

```
# metaclear -f mymirror
```

The following example deletes a hot spare pool, hsp001.

```
# metaclear hsp001
```

The following example deletes a soft partition, d23.

```
# metaclear d23
```

The following example purges all soft partitions on the slice c2t3d5s2 if those partitions are not being used by other metadevices or are not open.

```
# metaclear -p c2t3d5s2
```

The following example purges soft partitions from a metadvice.

```
# metaclear -p d2
d3: Soft Partition is cleared
d4: Soft Partition is cleared
d5: Soft Partition is cleared
```

退出状态

The following exit values are returned:

0 Successful completion.

>0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见

[mdmonitord\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metabs\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metarecover\(1M\)](#), [metarename\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metassist\(1M\)](#), [metastat\(1M\)](#), [metasync\(1M\)](#), [metattach\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [md.tab\(4\)](#), [attributes\(5\)](#), [md\(7D\)](#)

《Solaris Volume Manager Administration Guide》

引用名 metadb – create and delete replicas of the metadvice state database

用法概要

```
/usr/sbin/metadb -h

/usr/sbin/metadb [-s setname]

/usr/sbin/metadb [-s setname] -a [-f] [-k system-file] mddbnn

/usr/sbin/metadb [-s setname] -a [-f] [-k system-file]
    [-c number] [-l length] slice...

/usr/sbin/metadb [-s setname] -d [-f] [-k system-file] mddbnn

/usr/sbin/metadb [-s setname] -d [-f] [-k system-file] slice...

/usr/sbin/metadb [-s setname] -i

/usr/sbin/metadb [-s setname] -p [-k system-file]
    [mddb.cf-file]
```

描述

The metadb command creates and deletes replicas of the metadvice state database. State database replicas can be created on dedicated slices, or on slices that will later become part of a simple metadvice (concatenation or stripe) or RAID5 metadvice. Do not place state database replicas on fabric-attached storage, SANs, or other storage that is not directly attached to the system and available at the same point in the boot process as traditional SCSI or IDE drives. See NOTES.

The metadvice state database contains the configuration of all metadevices and hot spare pools in the system. Additionally, the metadvice state database keeps track of the current state of metadevices and hot spare pools, and their components. Solaris Volume Manager automatically updates the metadvice state database when a configuration or state change occurs. A submirror failure is an example of a state change. Creating a new metadvice is an example of a configuration change.

The metadvice state database is actually a collection of multiple, replicated database copies. Each copy, referred to as a replica, is subject to strict consistency checking to ensure correctness.

Replicated databases have an inherent problem in determining which database has valid and correct data. To solve this problem, Volume Manager uses a *majority consensus algorithm*. This algorithm requires that a majority of the database replicas be available before any of them are declared valid. This algorithm strongly encourages the presence of at least three initial replicas, which you create. A consensus can then be reached as long as at least two of the three replicas are available. If there is only one replica and the system crashes, it is possible that all metadvice configuration data can be lost.

The majority consensus algorithm is conservative in the sense that it will fail if a majority consensus cannot be reached, even if one replica actually does contain the most up-to-date data. This approach guarantees that stale data will not be accidentally used, regardless of the failure scenario. The majority consensus algorithm accounts for the following: the system will

stay running with exactly half or more replicas; the system will panic when less than half the replicas are available; the system will not reboot without one more than half the total replicas.

When used with no options, the `metadb` command gives a short form of the status of the metadvice state database. Use `metadb -i` for an explanation of the flags field in the output.

The initial state database is created using the `metadb` command with both the `-a` and `-f` options, followed by the slice where the replica is to reside. The `-a` option specifies that a replica (in this case, the initial) state database should be created. The `-f` option forces the creation to occur, even though a state database does not exist. (The `-a` and `-f` options should be used together only when no state databases exist.)

Additional replicas beyond those initially created can be added to the system. They contain the same information as the existing replicas, and help to prevent the loss of the configuration information. Loss of the configuration makes operation of the metadvice impossible. To create additional replicas, use the `metadb -a` command, followed by the name of the new slice(s) where the replicas will reside. All replicas that are located on the same slice must be created at the same time.

To delete all replicas that are located on the same slice, the `metadb -d` command is used, followed by the slice name.

When used with the `-i` option, `metadb` displays the status of the metadvice state databases. The status can change if a hardware failure occurs or when state databases have been added or deleted.

To fix a replica in an error state, delete the replica and add it back again.

The metadvice state database (`mddb`) also contains a list of the replica locations for this set (local or shared diskset).

The local set `mddb` can also contain host and drive information for each of the shared disksets of which this node is a member. Other than the diskset host and drive information stored in the local set `mddb`, the local and shared diskset `mddb`s are functionality identical.

The `mddb`s are written to during the `resync` of a mirror or during a component failure or configuration change. A configuration change or failure can also occur on a single replica (removal of a `mddb` or a failed disk) and this causes the other replicas to be updated with this failure information.

选项

Root privileges are required for all of the following options except `-h` and `-i`.

The following options can be used with the `metadb` command. Not all the options are compatible on the same command line. Refer to the SYNOPSIS to see the supported use of the options.

- a
Attach a new database device. The `/kernel/drv/md.conf` file is automatically updated with the new information and the `/etc/lvm/mddb.cf` file is updated as well. An alternate way to create replicas is by defining them in the `/etc/lvm/md.tab` file and specifying the assigned name at the command line in the form, `mddbnn`, where *nn* is a two-digit number given to the replica definitions. Refer to the [md.tab\(4\)](#) man page for instructions on setting up replicas in that file.
- c *number*
Specifies the number of replicas to be placed on each device. The default number of replicas is 1.
- d
Deletes all replicas that are located on the specified *slice*. The `/kernel/drv/md.conf` file is automatically updated with the new information and the `/etc/lvm/mddb.cf` file is updated as well.
- f
The -f option is used to create the initial state database. It is also used to force the deletion of replicas below the minimum of one. (The -a and -f options should be used together only when no state databases exist.)
- h
Displays a usage message.
- i
Inquire about the status of the replicas. The output of the -i option includes characters in front of the device name that represent the status of the state database. Explanations of the characters are displayed following the replica status and are as follows:
 - d
replica does not have an associated device ID.
 - o
replica active prior to last mddb configuration change
 - u
replica is up to date
 - l
locator for this replica was read successfully
 - c
replica's location was in `/etc/lvm/mddb.cf`
 - p
replica's location was patched in kernel
 - m
replica is master, this is replica selected as input

- r
replica does not have device relocation information
 - t
tagged data is associated with the replica
 - W
replica has device write errors
 - a
replica is active, commits are occurring to this
 - M
replica had problem with master blocks
 - D
replica had problem with data blocks
 - F
replica had format problems
 - S
replica is too small to hold current database
 - R
replica had device read errors
 - B
tagged data associated with the replica is not valid
- k *system-file***
Specifies the name of the kernel file where the replica information should be written. The default *system-file* is `/kernel/drv/md.conf`. This option is for use with the local diskset only.
- l *length***
Specifies the size of each replica. The default *length* is 8192 blocks, which should be appropriate for most configurations. “Replica” sizes of less than 128 blocks are not recommended.
- p**
Specifies updating the system file (`md.conf`) in the current working directory with entries from the `/etc/lvm/mddb.cf` file. This option is normally used to update a newly built system before it is booted for the first time. If the system has been built on a system other than the one where it will run, the location of the `mddb.cf` on the local machine can be passed as an argument. The system file to be updated can be changed using the `-k` option. This option is for use with the local diskset only.

-s *setname*

Specifies the name of the diskset on which the metadb command will work. Using the -s option will cause the command to perform its administrative function within the specified diskset. Without this option, the command will perform its function on local database replicas.

slice

Specifies the logical name of the physical slice (partition), such as /dev/dsk/c0t0d0s3.

示例**示例 1 Creating Initial State Database Replicas**

The following example creates the initial state database replicas on a new system.

```
# metadb -a -f c0t0d0s7 c0t1d0s3 c1t0d0s7 c1t1d0s3
```

The -a and -f options force the creation of the initial database and replicas. You could then create metadevices with these same slices, making efficient use of the system.

示例 2 Adding Two Replicas on Two New Disks

This example shows how to add two replicas on two new disks that have been connected to a system currently running Volume Manager.

```
# metadb -a c0t2d0s3 c1t1d0s3
```

示例 3 Deleting Two Replicas

This example shows how to delete two replicas from the system. Assume that replicas have been set up on /dev/dsk/c0t2d0s3 and /dev/dsk/c1t1d0s3.

```
# metadb -d c0t2d0s3 c1t1d0s3
```

Although you can delete all replicas, you should never do so while metadevices still exist. Removing all replicas causes existing metadevices to become inoperable.

文件

/etc/lvm/mddb.cf

Contains the location of each copy of the metadevice state database.

/etc/lvm/md.tab

Workspace file for metadevice database configuration.

/kernel/drv/md.conf

Contains database replica information for all metadevices on a system. Also contains Solaris Volume Manager configuration information.

退出状态

The following exit values are returned:

0

successful completion

>0

an error occurred

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见 [mdmonitord\(1M\)](#), [metaclear\(1M\)](#), [metadetach\(1M\)](#), [metahs\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metarecover\(1M\)](#), [metarename\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metassist\(1M\)](#), [metastat\(1M\)](#), [metasync\(1M\)](#), [metattach\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [md.tab\(4\)](#), [attributes\(5\)](#), [md\(7D\)](#)

 《Solaris Volume Manager Administration Guide》

附注 Replicas cannot be stored on fabric-attached storage, SANs, or other storage that is not directly attached to the system. Replicas must be on storage that is available at the same point in the boot process as traditional SCSI or IDE drives. A replica can be stored on a:

- Dedicated local disk partition
- Local partition that will be part of a volume
- Local partition that will be part of a UFS logging device

引用名	metadevadm – update metadvice information																				
用法概要	<pre>/usr/sbin/metadevadm [-h] [-n] [[-l]-r] [-s <i>setname</i>] [-u <i>disk_specifier</i>] [-v]</pre>																				
描述	The <code>metadevadm</code> command facilitates the administration of device ID entries in Solaris Volume Manager. Use this command when the pathname stored in the metadvice state database no longer correctly addresses the device or when a disk drive has had its device ID changed. This command requires root privileges.																				
选项	The following options are supported. <table><tr><td>-h</td><td>Provide a help display.</td></tr><tr><td>-l</td><td>Specify that <code>metadevadm</code> log to syslog(3C). <code>metadevadm</code> logs to the DAEMON facility at the ERR level by default. See syslog.conf(4) for additional information on changing logging levels.</td></tr><tr><td></td><td>Use this option anytime. It is most useful in startup scripts and less useful interactively.</td></tr><tr><td>-n</td><td>Emulate the effect of a command, without making any changes to the system.</td></tr><tr><td>-r</td><td>Recompute the pathname and disk specifier (including slice) associated with all devices in the metadvice state database if the device supports device IDs. If a device does not support device IDs or the device is not available, then no action is taken for that device.</td></tr><tr><td></td><td>Use this option when the disk has been moved or readdressed.</td></tr><tr><td></td><td>This option is run automatically at boot time to detect device ID changes and update the state database.</td></tr><tr><td>-s <i>setname</i></td><td>Specify the name of the disk set on which <code>metadevadm</code> works. This option causes the command to perform its administrative function within the specified disk set. Without this option, the command performs its function on devices in the local disk set.</td></tr><tr><td>-u <i>disk_specifier</i></td><td>Obtain the device ID associated with the <i>disk_specifier</i> (for example, <code>c1t2d0</code>) of a device and update the metadvice state database. If the device ID has not changed this option does nothing. Use this option when a disk drive has had its device ID changed during a firmware upgrade or due to changing the controller of a storage subsystem.</td></tr><tr><td>-v</td><td>Execute in verbose mode. This option has no effect when used with <code>-u</code>. Verbose is the default.</td></tr></table>	-h	Provide a help display.	-l	Specify that <code>metadevadm</code> log to syslog(3C) . <code>metadevadm</code> logs to the DAEMON facility at the ERR level by default. See syslog.conf(4) for additional information on changing logging levels.		Use this option anytime. It is most useful in startup scripts and less useful interactively.	-n	Emulate the effect of a command, without making any changes to the system.	-r	Recompute the pathname and disk specifier (including slice) associated with all devices in the metadvice state database if the device supports device IDs. If a device does not support device IDs or the device is not available, then no action is taken for that device.		Use this option when the disk has been moved or readdressed.		This option is run automatically at boot time to detect device ID changes and update the state database.	-s <i>setname</i>	Specify the name of the disk set on which <code>metadevadm</code> works. This option causes the command to perform its administrative function within the specified disk set. Without this option, the command performs its function on devices in the local disk set.	-u <i>disk_specifier</i>	Obtain the device ID associated with the <i>disk_specifier</i> (for example, <code>c1t2d0</code>) of a device and update the metadvice state database. If the device ID has not changed this option does nothing. Use this option when a disk drive has had its device ID changed during a firmware upgrade or due to changing the controller of a storage subsystem.	-v	Execute in verbose mode. This option has no effect when used with <code>-u</code> . Verbose is the default.
-h	Provide a help display.																				
-l	Specify that <code>metadevadm</code> log to syslog(3C) . <code>metadevadm</code> logs to the DAEMON facility at the ERR level by default. See syslog.conf(4) for additional information on changing logging levels.																				
	Use this option anytime. It is most useful in startup scripts and less useful interactively.																				
-n	Emulate the effect of a command, without making any changes to the system.																				
-r	Recompute the pathname and disk specifier (including slice) associated with all devices in the metadvice state database if the device supports device IDs. If a device does not support device IDs or the device is not available, then no action is taken for that device.																				
	Use this option when the disk has been moved or readdressed.																				
	This option is run automatically at boot time to detect device ID changes and update the state database.																				
-s <i>setname</i>	Specify the name of the disk set on which <code>metadevadm</code> works. This option causes the command to perform its administrative function within the specified disk set. Without this option, the command performs its function on devices in the local disk set.																				
-u <i>disk_specifier</i>	Obtain the device ID associated with the <i>disk_specifier</i> (for example, <code>c1t2d0</code>) of a device and update the metadvice state database. If the device ID has not changed this option does nothing. Use this option when a disk drive has had its device ID changed during a firmware upgrade or due to changing the controller of a storage subsystem.																				
-v	Execute in verbose mode. This option has no effect when used with <code>-u</code> . Verbose is the default.																				

示例

示例1 Updating Device ID of Disk

The following example updates the device c2t3d0:

```
# metadevadm -u c2t3d0
Updating SLVM device relocation information for c2t3d0.
Old device reloc information: id19280192391293123012012010012012091398
New device reloc information: id19380192391293123012012010012012091398
```

The following example is a variation of the preceding, using the full pathname.

```
# metadevadm -u /dev/dsk/c2t3d0
```

The following example uses the -n option, which means that the command is emulated, but does not take effect. Note that when the -v option is used with -u, -v has no effect (verbose is the default).

```
# metadevadm -u -v -n c2t3d0
Updating SLVM device relocation information for c2t3d0.
Old device reloc information: id19280192391293123012012010012012091398
New device reloc information: id19380192391293123012012010012012091398
```

示例2 Recomputing Pathnames

In the following example, all device names are valid.

```
# metadevadm -r
Disk movement detected.
Updating device names in SLVM.
```

In the following example, once again device names are valid.

```
# metadevadm -r -v
Disk movement detected.
Updating device names in SLVM.
c0t0d0s0 changed to c0t0d1s0 from device relocation information
id12098123lkmklsdjaaskfjadfjakds
```

In the following example, metadevadm detects an invalid device name.

```
# metadevadm -r
Invalid device relocation information detected in SLVM.
Please check status of following disk(s):
c3t0d0
```

返回值

The following exit values are returned:

- 0 Command was successful.
- 1 metadevadm encountered an error condition.

- 2 An invalid device ID was detected when using the `-r` option. This is for use in the `rc2.d` script. See `init.d(4)`.

属性

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见

`mdmonitord(1M)`, `metaclear(1M)`, `metadb(1M)`, `metadetach(1M)`, `metahs(1M)`, `metainit(1M)`, `metaoffline(1M)`, `metaonline(1M)`, `metaparam(1M)`, `metarecover(1M)`, `metarename(1M)`, `metareplace(1M)`, `metaset(1M)`, `metassist(1M)`, `metastat(1M)`, `metasync(1M)`, `metattach(1M)`, `md.tab(4)`, `md.cf(4)`, `mddb.cf(4)`, `md.tab(4)`, `attributes(5)`, `md(7D)`

《Solaris Volume Manager Administration Guide》

引用名

metahs – manage hot spares and hot spare pools

用法概要

```
/usr/sbin/metahs [-s setname] -a all component
/usr/sbin/metahs [-s setname] -a hot_spare_pool [component]
/usr/sbin/metahs [-s setname] -d hot_spare_pool [component]
/usr/sbin/metahs [-s setname] -d all component
/usr/sbin/metahs [-s setname] -e component
/usr/sbin/metahs [-s setname] -r hot_spare_pool component-old
/usr/sbin/metahs [-s setname] -r all component-old component-new
/usr/sbin/metahs [-s setname] -i [hot_spare_pool]....
```

描述

The metahs command manages existing hot spares and hot spare pools. It is used to add, delete, enable, and replace components (slices) in hot spare pools. Like the metainit command, the metahs command can also create an initial hot spare pool. The metahs command does not replace a component of a metadevice. This function is performed by the metareplace command.

Hot spares are always in one of three states: available, in-use, or broken. Available hot spares are running and ready to accept data, but are not currently being written to or read from. In-use hot spares are currently being written to and read from. Broken hot spares are out of service and should be repaired. The status of hot spares is displayed when metahs is invoked with the -i option.

Solaris Volume Manager supports storage devices and logical volumes, including hot spares, greater than 1 terabyte (TB) when Solaris 10 is running a 64-bit kernel.

If a system with large volumes or hot spares is rebooted under a 32-bit Solaris 10 kernel, the large volumes are visible through metastat output, but they cannot be accessed, modified or deleted, and no new large volumes can be created. Any volumes or file systems on a large volume in this situation are also unavailable. If a system with large volumes is rebooted under a version of Solaris prior to Solaris 10, Solaris Volume Manager will not start. All large volumes must be removed before Solaris Volume Manager runs under another version of the Solaris Operating Environment.

选项

Root privileges are required for any of the following options except -i.

The following options are supported:

-a *all component*

Add *component* to all hot spare pools. *all* is not case sensitive.

-a *hot_spare_pool [component]*

Add the *component* to the specified *hot_spare_pool*. *hot_spare_pool* is created if it does not already exist.

-d all <i>component</i>	Delete <i>component</i> from all the hot spare pools. The <i>component</i> cannot be deleted if it is in the in-use state.
-d <i>hot_spare_pool</i> [<i>component</i>]	Delete <i>hot_spare_pool</i> , if the <i>hot_spare_pool</i> is both empty and not referenced by a metadvice. If <i>component</i> is specified, it is deleted from the <i>hot_spare_pool</i> . Hot spares in the in-use state cannot be deleted.
-e <i>component</i>	Enable <i>component</i> to be available for use as a hot spare. The <i>component</i> can be enabled if it is in the broken state and has been repaired.
-i [<i>hot_spare_pool</i> . . .]	Display the status of the specified <i>hot_spare_pool</i> or for all hot spare pools if one is not specified.
-r all <i>component-old component-new</i>	Replace <i>component-old</i> with <i>component-new</i> in all hot spare pools which have the component associated. Components cannot be replaced from any hot spare pool if the old hot spare is in the in-use state.
-r <i>hot_spare_pool component-old component-new</i>	Replace <i>component-old</i> with <i>component-new</i> in the specified <i>hot_spare_pool</i> . Components cannot be replaced from a hot spare pool if the old hot spare is in the in-use state.
-s <i>setname</i>	Specify the name of the diskset on which metahs works. Using the -s option causes the command to perform its administrative function within the specified diskset. Without this option, the command performs its function on local hot spare pools.

操作数

The following operands are supported:

<i>component</i>	The logical name for the physical slice (partition) on a disk drive, such as /dev/dsk/c0t0d0s2.
<i>hot_spare_pool</i>	Names for hot spare pools can be any legal file name that is composed of alphanumeric characters, a dash (“-”), an underscore (“_”), or a period

(“.“). Names must begin with a letter. The words “all” and “none” are reserved and cannot be used.

示例

示例 1 Adding a Hot Spare to a Hot Spare Pool

The following example adds a hot spare `/dev/dsk/c0t0d0s7` to a hot spare pool `mirror1_pool`:

```
# metahs -a mirror1_pool c0t0d0s7
```

When the hot spare is added to the pool, the existing order of the hot spares already in the pool is preserved. The new hot spare is added at the end of the list of hot spares in the hot spare pool specified.

示例 2 Adding a Hot Spare to All Currently Defined Pools

This example adds a hot spare to the hot spare pools that are currently defined:

```
# metahs -a all c0t0d0s7
```

The keyword `all` in this example specifies adding the hot spare, `/dev/dsk/c0t0d0s7`, to all the hot spare pools.

示例 3 Deleting a Hot Spare

This example deletes a hot spare, `/dev/dsk/c0t0d0s7`, from a hot spare pool, `hsp003`:

```
# metahs -d hsp003 c0t0d0s7
```

When you delete a hot spare, the position of the remaining hot spares in the pool changes to reflect the new order. For instance, if in this example `/dev/dsk/c0t0d0s7` were the second of three hot spares, after deletion the third hot spare would move to the second position.

示例 4 Replacing a Hot Spare

This example replaces a hot spare that was previously defined:

```
# metahs -r hsp001 c0t1d0s0 c0t3d0s0
```

In this example, the hot spare `/dev/dsk/c0t1d0s0` is replaced by `/dev/dsk/c0t3d0s0`. The order of the hot spares does not change.

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见

[mdmonitor\(1M\)](#), [metaclear\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metarecover\(1M\)](#), [metarename\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metassist\(1M\)](#), [metastat\(1M\)](#), [metasync\(1M\)](#), [metattach\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [md.tab\(4\)](#), [attributes\(5\)](#), [md\(7D\)](#)

[《Solaris Volume Manager Administration Guide》](#)

警告

Do not create large (>1 TB) volumes if you expect to run the Solaris Operating Environment with a 32-bit kernel or if you expect to use a version of the Solaris Operating Environment prior to Solaris 10.

引用名	metainport – imports disk sets into existing Solaris Volume Manager configurations
用法概要	<pre>metainport -s setname [-n] [-f] [-v] [disks]... metainport -r [-v] [disks]... metainport -V metainport -?</pre>
描述	The <code>metainport</code> command allows the importing of disk sets, including replicated disk sets, into an existing Solaris Volume Manager configuration. Replicated disk sets are disk sets created using remote replication software. The default Solaris Volume Manager configuration specifies a maximum number of disk sets that can be configured. The <code>metainport</code> command fails if importing the disk set would result in exceeding the number of disk sets configured on the system. To increase the number of disk sets allowed on a system, see the 《Solaris Volume Manager Administration Guide》. Use metaset(1M) or metastat(1M) to view the configuration of the imported set. You must run <code>metainport</code> as root. <code>metainport</code> requires a functional Solaris Volume Manager configuration before it runs. The following options are supported: -f Force the import, even if a quorum of replicas from the imported disk set is not available. This option could result in corrupt configurations and should only be used when <code>metainport</code> fails with the “Insufficient quorum detected; exiting” error. If only a partial disk set is available, this option might be necessary to successfully import. Some or all data could be corrupted or unavailable when importing a partial set or a set lacking a replica quorum. -n Does not actually perform the operation, but shows the output or errors that would have resulted from the operation, had it been run. -r Report on the non-configured disk sets found on the system. If no disk device or LUN is specified, <code>metainport</code> reports on all non-configured disk sets attached to the system. When the name of one disk is specified, <code>metainport</code> reports on the disk set (or virtual LUN) containing the specified disk. If two or more disks are specified, <code>metainport</code> reports on the set (or sets, if they belong to different disk sets) containing the specified disks. If two or more disks are specified, <code>metainport</code> reports on the set (or sets, if they belong to different disk sets) containing the specified disks. This option can be used in conjunction with the <code>-v</code> option to give verbose output on each disk set reported.
选项	

- s *setname* Specify the disk set name to use when importing. The imported disk set will be called *setname*, without regard to the name it may have had on a different system.
- v Verbose. Provides detailed information about the metadb replica location and status. It also provides detailed information about the disk set configuration and status similar to the “metastat -c” output.
- V Version information.
- ? Display a help message.

示例

示例 1 Importing a Disk Set

The following example creates a disk set called blue and identifies c1t5d0 as a disk containing a state database replica from the disk set being imported.

```
# metainport -s blue c1t5d0
```

示例 2 Reporting Disk Sets to Import

The following example scans all disks and LUNs attached to the system and configured as part of the system. It scans for disks that could be part of a disk set to be imported. Components that are already part of the Solaris Volume Manager configuration are ignored.

This use of metainport provides suggested forms of the metainport command to use to actually import the disk sets that have been found. You can specify a component on the command line to reduce the scope of the scan and generate results more quickly.

```
# metainport -r
```

退出状态

- 0 Successful completion.
- >0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见

[mdmonitord\(1M\)](#), [metaclear\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metahs\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metarecover\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metastat\(1M\)](#), [metasync\(1M\)](#), [metattach\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [attributes\(5\)](#)

《Solaris Volume Manager Administration Guide》

引用名

metainit – configure metadevices

用法概要

```

/usr/sbin/metainit -h

/usr/sbin/metainit [generic options] concat/stripe numstripes width
                    component... [-i interlace]

/usr/sbin/metainit [width component... [-i interlace]]
                    [-h hot_spare_pool]

/usr/sbin/metainit [generic options] mirror -m submirror
                    [read_options] [write_options]
                    [pass_num]

/usr/sbin/metainit [generic options] RAID -r component...
                    [-i interlace]
                    [-h hot_spare_pool] [-k] [-o original_column_count]

/usr/sbin/metainit [generic options] hot_spare_pool
                    [hotspare...]

/usr/sbin/metainit [generic options] metadvice-name

/usr/sbin/metainit [generic options] -a

/usr/sbin/metainit [generic options] softpart -p [-e] component
                    [-A alignment] size

/usr/sbin/metainit -r

```

描述

The `metainit` command configures metadevices and hot spares according to the information specified on the command line. Alternatively, you can run `metainit` so that it uses configuration entries you specify in the `/etc/lvm/md.tab` file (see [md.tab\(4\)](#)). All metadevices must be set up by the `metainit` command before they can be used.

Solaris Volume Manager supports storage devices and logical volumes greater than 1 terabyte (TB) when a system runs a 64-bit Solaris kernel. Support for large volumes is automatic. If a device greater than 1 TB is created, Solaris Volume Manager configures it appropriately and without user intervention.

If a system with large volumes is rebooted under a 32-bit Solaris kernel, the large volumes are visible through `metastat` output. Large volumes cannot be accessed, modified or deleted, and no new large volumes can be created. Any volumes or file systems on a large volume in this situation are unavailable. If a system with large volumes is rebooted under a version of Solaris prior to the Solaris 9 4/03 release, Solaris Volume Manager does not start. You must remove all large volumes before Solaris Volume Manager runs under an earlier version of the Solaris Operating System.

If you edit the `/etc/lvm/md.tab` file to configure metadevices, specify one complete configuration entry per line. You then run the `metainit` command with either the `-a` option, to activate all metadevices you entered in the `/etc/lvm/md.tab` file, or with the metadvice name corresponding to a specific configuration entry.

metainit does not maintain the state of the volumes that would have been created when metainit is run with both the `-a` and `-n` flags. Any volumes in `md.tab` that have dependencies on other volumes in `md.tab` are reported as errors when `metainit -a -n` is run, although the operations might succeed when `metainit -a` is run. See [md.tab\(4\)](#).

Solaris Volume Manager never updates the `/etc/lvm/md.tab` file. Complete configuration information is stored in the metadvice state database, not `md.tab`. The only way information appears in `md.tab` is through editing it by hand.

When setting up a disk mirror, the *first* step is to use metainit create a one-on-one concatenation for the named slice. See EXAMPLES.

选项

The following options are supported:

Generic Options

Root privileges are required for all of the following options except `-h`.

The following generic options are supported:

`-f`

Forces the `metainit` command to continue even if one of the slices contains a mounted file system or is being used as swap, or if the stripe being created is smaller in size than the underlying soft partition. This option is required when configuring mirrors on swap.

`-h`

Displays usage message.

`-n`

Checks the syntax of your command line or `md.tab` entry without actually setting up the metadvice. If used with `-a`, all devices are checked but not initialized.

`-r`

Only used in a shell script at boot time. Sets up all metadvice devices that were configured before the system crashed or was shut down. The information about previously configured metadvice devices is stored in the metadvice state database (see [metadb\(1M\)](#)).

`-s setname`

Specifies the name of the diskset on which `metainit` works. Without the `-s` option, the `metainit` command operates on your local metadvice devices and/or hotspares.

Concat/Stripe Options

The following concat/stripe options are supported:

concat/stripe

Specifies the metadvice name of the concatenation, stripe, or concatenation of stripes being defined.

numstripes

Specifies the number of individual stripes in the metadvice. For a simple stripe, *numstripes* is always 1. For a concatenation, *numstripes* is equal to the number of slices. For a concatenation of stripes, *numstripes* varies according to the number of stripes.

width

Specifies the number of slices that make up a stripe. When *width* is greater than 1, the slices are striped.

component

The logical name for the physical slice (partition) on a disk drive, such as `/dev/dsk/c0t0d0s0`. For RAID level 5 metadevices, a minimum of three slices is necessary to enable striping of the parity information across slices.

-i interlace

Specifies the interlace size. This value tells Solaris Volume Manager how much data to place on a slice of a striped or RAID level 5 metadvice before moving on to the next slice. *interlace* is a specified value, followed by either 'k' for kilobytes, 'm' for megabytes, or 'b' for blocks. The characters can be either uppercase or lowercase. The *interlace* specified cannot be less than 16 blocks, or greater than 100 megabytes. If *interlace* is not specified, it defaults to 512 kilobytes.

-h hot_spare_pool

Specifies the *hot_spare_pool* to be associated with the metadvice. If you use the command line, the hot spare pool must have been previously created by the `metainit` command before it can be associated with a metadvice. Use `/ -h hspnnn` when the concat/stripe being created is to be used as a submirror.

Names for hot spare pools can be any legal file name that is composed of alphanumeric characters, a dash ("-"), an underscore ("_"), or a period ("."). Names must begin with a letter. The words "all" and "none" are reserved and cannot be used.

Mirror Options

The following mirror options are supported:

mirror -m submirror

Specifies the metadvice name of the mirror. The `-m` indicates that the configuration is a mirror. *submirror* is a metadvice (stripe or concatenation) that makes up the initial one-way mirror. Solaris Volume Manager supports a maximum of four-way mirroring. When defining mirrors, first create the mirror with the `metainit` command as a one-way mirror. Then attach subsequent submirrors using the `metattach` command. This method ensures that Solaris Volume Manager properly syncs the mirrors. (The second and any subsequent submirrors are first created using the `metainit` command.)

read_options

The following read options for mirrors are supported:

-g

Enables the geometric read option, which results in faster performance on sequential reads.

-r

Directs all reads to the first submirror. This should only be used when the devices comprising the first submirror are substantially faster than those of the second mirror. This flag cannot be used with the `-g` flag.

If neither the `-g` nor `-r` flags are specified, reads are made in a round-robin order from all submirrors in the mirror. This enables load balancing across the submirrors.

write_options

The following write options for mirrors are supported:

`-S`

Performs serial writes to mirrors. The first submirror write completes before the second is started. This can be useful if hardware is susceptible to partial sector failures. If `-S` is not specified, writes are replicated and dispatched to all mirrors simultaneously.

pass_num

A number in the range 0-9 at the end of an entry defining a mirror that determines the order in which that mirror is resynced during a reboot. The default is 1. Smaller pass numbers are resynced first. Equal pass numbers are run concurrently. If 0 is used, the resync is skipped. 0 should be used only for mirrors mounted as read-only, or as swap.

RAID Level 5 Options

The following RAID level 5 options are available:

RAID - r

Specifies the name of the RAID level 5 metadvice. The `-r` specifies that the configuration is RAID level 5.

`-k`

For RAID level 5 metadevices, informs the driver that it is not to initialize (zero the disk blocks) due to existing data. Only use this option to recreate a previously created RAID level 5 device.

Use the `-k` option with extreme caution. This option sets the disk blocks to the OK state. If any errors exist on disk blocks within the metadvice, Solaris Volume Manager might begin fabricating data. Instead of using the `-k` option, you might want to initialize the device and restore data from tape.

-o original_column_count

For RAID level 5 metadevices, used with the `-k` option to define the number of original slices in the event the originally defined metadvice was grown. This is necessary since the parity segments are not striped across concatenated devices.

Use the `-o` option with extreme caution. This option sets the disk blocks to the OK state. If any errors exist on disk blocks within the metadvice, Solaris Volume Manager might begin fabricating data. Instead of using the `-o` option, you might want to initialize the device and restore data from tape.

Soft Partition Options

The following soft partition options are supported:

softpart -p [-e] component [-A alignment] size

The *softpart* argument specifies the name of the soft partition. The `-p` specifies that the configuration is a soft partition.

The `-e` specifies that the entire disk specified by *component* as `c*t*d*` should be repartitioned and reserved for soft partitions. The specified component is repartitioned such that slice 7 reserves space for system (state database replica) usage and slice 0 contains all remaining space on the disk. Slice 7 is a minimum of 4MB, but can be larger, depending on the disk geometry. The newly created soft partition is placed on slice 0 of the device.

The *component* argument specifies the disk (`c*t*d*`), slice (`c*t*d*s*`), or meta device (`d*`) from which to create the soft partition. The *size* argument determines the space to use for the soft partition and can be specified in K or k for kilobytes, M or m for megabytes, G or g for gigabytes, T or t for terabyte (one terabyte is the maximum size), and B or b for blocks (sectors). All values represent powers of 2, and upper and lower case options are equivalent. Only integer values are permitted.

The `-A` alignment option sets the value of the soft partition extent alignment. This option is used when it is important to specify a starting offset for the soft partition. It preserves the data alignment between the metadvice address space and the address space of the underlying physical device. For example, a hardware device that does checksumming should not have its I/O requests divided by Solaris Volume Manager. In this case, use a value from the hardware configuration as the value for the alignment. When you use this option in conjunction with a software I/O load, the alignment value corresponds to the I/O load of the application. This prevents I/O from being divided unnecessarily and affecting performance.

The literal `all`, used instead of specifying size, specifies that the soft partition should occupy all available space on the device.

Hot Spare Pool Options The following hot spare pool options are supported:

hot_spare_pool [*hotspare...*]

When used as arguments to the `metainit` command, *hot_spare_pool* defines the name for a hot spare pool, and *hotspare...* is the logical name for the physical slice(s) for availability in that pool. Names for hot spare pools can be any legal file name that is composed of alphanumeric characters, a dash (“-”), an underscore (“_”), or a period (“.”). Names must begin with a letter. The words “all” and “none” are reserved and cannot be used.

md.tab File Options The following `md.tab` file options are supported:

metadvice-name

When the `metainit` command is run with a *metadvice-name* as its only argument, it searches the `/etc/lvm/md.tab` file to find that name and its corresponding entry. The order in which entries appear in the `md.tab` file is unimportant. For example, consider the following `md.tab` entry:

```
d0 2 1 c1t0d0s0 1 c2t1d0s0
```

When you run the command `metainit d0`, it configures metadvice `d0` based on the configuration information found in the `md.tab` file.

-a

Activates all metadevices defined in the `md.tab` file.

`metainit` does not maintain the state of the volumes that would have been created when `metainit` is run with both the `-a` and `-n` flags. If a device `d0` is created in the first line of the `md.tab` file, and a later line in `md.tab` assumes the existence of `d0`, the later line fails when `metainit -an` runs (even if it would succeed with `metainit -a`).

示例

示例 1 Creating a One-on-One Concatenation

The following command creates a one-on-one concatenation for the root slice. This is the first step you take when setting up a mirror for any other slice that cannot be unmounted. The `-f` option is required to create a volume with an existing file system.

```
# metainit -f dl 1 1 c0t0d0s0
```

The preceding command makes `d1` a one-on-one concatenation, using the named slice. You can then enter:

```
# metainit d0 -m d1
```

...to make a one-way mirror of the named slice.

示例 2 Concatenation

All drives in the following examples have the same size of 525 Mbytes.

This example shows a metadevice, `/dev/md/dsk/d7`, consisting of a concatenation of four slices.

```
# metainit d7 4 1 c0t1d0s0 1 c0t2d0s0 1 c0t3d0s0 1 /dev/dsk/c0t4d0s0
```

The number 4 indicates there are four individual stripes in the concatenation. Each stripe is made of one slice, hence the number 1 appears in front of each slice. The first disk sector in all of these devices contains a disk label. To preserve the labels on devices `/dev/dsk/c0t2d0s0`, `/dev/dsk/c0t3d0s0`, and `/dev/dsk/c0t4d0s0`, the metadisk driver must skip at least the first sector of those disks when mapping accesses across the concatenation boundaries. Because skipping only the first sector would create an irregular disk geometry, the entire first cylinder of these disks is skipped. This allows higher level file system software to optimize block allocations correctly.

示例 3 Stripe

This example shows a metadevice, `/dev/md/dsk/d15`, consisting of two slices.

```
# metainit d15 1 2 c0t1d0s0 c0t2d0s0 -i 32k
```

The number 1 indicates that one stripe is being created. Because the stripe is made of two slices, the number 2 follows next. The optional `-i` followed by `32k` specifies the interlace size as 32 Kbytes. If the interlace size were not specified, the stripe would use the default value of 16

示例3 Stripe (续)

Kbytes.

示例4 Concatentation of Stripes

This example shows a metadvice, /dev/md/dsk/d75, consisting of a concatenation of two stripes of three disks.

```
# metainit d75 2 3 c0t1d0s0 c0t2d0s0 \
    c0t3d0s0 -i 16k \
    3 c1t1d0s0 c1t2d0s0 c1t3d0s0 -i 32k
```

On the first line, the -i followed by 16k specifies that the stripe interlace size is 16 Kbytes. The second set specifies the stripe interlace size as 32 Kbytes. If the second set did not specify 32 Kbytes, the set would use the default interlace value of 16 Kbytes. The blocks of each set of three disks are interlaced across three disks.

示例5 Mirroring

This example shows a two-way mirror, /dev/md/dsk/d50, consisting of two submirrors. This mirror does not contain any existing data.

```
# metainit d51 1 1 c0t1d0s0
# metainit d52 1 1 c0t2d0s0
# metainit d50 -m d51
# metattach d50 d52
```

In this example, two submirrors, d51 and d52, are created with the metainit command. These two submirrors are simple concatenations. Next, a one-way mirror, d50, is created using the -m option with d51. The second submirror is attached later using the metattach command. When creating a mirror, any combination of stripes and concatenations can be used. The default read and write options in this example are a round-robin read algorithm and parallel writes to all submirrors.

示例6 Creating a metadvice in a diskset

This example shows a metadvice, /dev/md/dsk/d75, consisting of a concatenation of two stripes within a diskset called set1.

```
# metainit -s set1 d75 2 3 c2t1d0s0 c2t2d0s0 \
    c2t3d0s0 -i 32k
# metainit -s set1 d51 1 1 c2t1d0s0
# metainit -s set1 d52 1 1 c3t1d0s0
# metainit -s set1 d50 -m d51
# metattach -s set1 d50 d52
```

In this example, a diskset is created using the metaset command. Metadevices are then created within the diskset using the metainit command. The two submirrors, d51 and d52, are simple concatenations. Next, a one-way mirror, d50, is created using the -m option with

示例 6 Creating a metadvice in a diskset (续)

d51. The second submirror is attached later using the `metattach` command. When creating a mirror, any combination of stripes and concatenations can be used. The default read and write options in this example are a round-robin read algorithm and parallel writes to all submirrors.

示例 7 RAID Level 5

This example shows a RAID level 5 device, `d80`, consisting of three slices:

```
# metainit d80 -r c1t0d0s0 c1t1d0s0 c1t3d0s0 -i 20k
```

In this example, a RAID level 5 metadvice is defined using the `-r` option with an interlace size of 20 Kbytes. The data and parity segments are striped across the slices, `c1t0d0s0`, `c1t2d0s0`, and `c1t3d0s0`.

示例 8 Soft Partition

The following example shows a soft partition device, `d1`, built on metadvice `d100` and 100 Mbytes (indicated by `100M`) in size:

```
# metainit d1 -p d100 100M
```

The preceding command creates a 100 Mbyte soft partition on the `d100` metadvice. This metadvice could be a RAID level 5, stripe, concatenation, or mirror.

示例 9 Soft Partition on Full Disk

The following example shows a soft partition device, `d1`, built on disk `c3t4d0`:

```
# metainit d1 -p -e c3t4d0 9G
```

In this example, the disk is repartitioned and a soft partition is defined to occupy all 9 Gbytes of disk `c3t4d0s0`.

示例 10 Soft Partition Taking All Available Space

The following example shows a soft partition device, `d1`, built on disk `c3t4d0`:

```
# metainit d1 -p -e c3t4d0 all
```

In this example, the disk is repartitioned and a soft partition is defined to occupy all available disk space on slice `c3t4d0s0`.

示例 11 Hot Spare

This example shows a two-way mirror, `/dev/md/dsk/d10`, and a hot spare pool with three hot spare components. The mirror does not contain any existing data.

```
# metainit hsp001 c2t2d0s0 c3t2d0s0 c1t2d0s0
# metainit d41 1 1 c1t0d0s0 -h hsp001
# metainit d42 1 1 c3t0d0s0 -h hsp001
```

示例 11 Hot Spare (续)

```
# metainit d40 -m d41
# metattach d40 d42
```

In this example, a hot spare pool, hsp001, is created with three slices from three different disks used as hot spares. Next, two submirrors are created, d41 and d42. These are simple concatenations. The `metainit` command uses the `-h` option to associate the hot spare pool hsp001 with each submirror. A one-way mirror is then defined using the `-m` option. The second submirror is attached using the `metattach` command.

示例 12 Setting the Value of the Soft Partition Extent Alignment

This example shows how to set the alignment of the soft partition to 1 megabyte.

```
# metainit -s red d13 -p clt3d0s4 -A 1m 4m
```

In this example the soft partition, d13, is created with an extent alignment of 1 megabyte. The `metainit` command uses the `-A` option with an alignment of 1m to define the soft partition extent alignment.

文件

/etc/lvm/md.tab

Contains list of metadvice and hot spare configurations for batch-like creation.

警告

This section contains information on different types of warnings.

Devices and Volumes Greater Than 1 TB

Do not create large (>1 TB) volumes if you expect to run the Solaris Operating Environment with a 32-bit kernel or if you expect to use a version of the Solaris Operating Environment prior to Solaris 10.

Multi-Way Mirror

Do not use the `metainit` command to create a multi-way mirror. Rather, create a one-way mirror with `metainit` then attach additional submirrors with `metattach`. When the `metattach` command is not used, no resync operations occur and data could become corrupted.

If you use `metainit` to create a mirror with multiple submirrors, the following message is displayed:

```
WARNING: This form of metainit is not recommended.
The submirrors may not have the same data.
Please see ERRORS in metainit(1M) for additional information.
```

Truncation of Soft Partitions

When creating stripes on top of soft partitions it is possible for the size of the new stripe to be less than the size of the underlying soft partition. If this occurs, `metainit` fails with an error indicating the actions required to overcome the failure.

If you use the `-f` option to override this behavior, the following message is displayed:

```
WARNING: This form of metainit is not recommended.
The stripe is truncating the size of the underlying device.
Please see ERRORS in metainit(1M) for additional information.
```

Write-On-Write Problem

When mirroring data in Solaris Volume Manager, transfers from memory to the disks do not all occur at exactly the same time for all sides of the mirror. If the contents of buffers are changed while the data is in-flight to the disk (called write-on-write), then different data can end up being stored on each side of a mirror.

This problem can be addressed by making a private copy of the data for mirror writes, however, doing this copy is expensive. Another approach is to detect when memory has been modified across a write by looking at the dirty-bit associated with the memory page. Solaris Volume Manager uses this dirty-bit technique when it can. Unfortunately, this technique does not work for raw I/O or direct I/O. By default, Solaris Volume Manager is tuned for performance with the liability that mirrored data might be out of sync if an application does a “write-on-write” to buffers associated with raw I/O or direct I/O. Without mirroring, you were not guaranteed what data would actually end up on media, but multiple reads would return the same data. With mirroring, multiple reads can return different data. The following line can be added to /etc/system to cause a stable copy of the buffers to be used for all raw I/O and direct I/O write operations.

```
set md_mirror:md_mirror_wow_flg=0x20
```

Setting this flag degrades performance.

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见

[mdmonitord\(1M\)](#), [metaclear\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metahs\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metarecover\(1M\)](#), [metarename\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metassist\(1M\)](#), [metastat\(1M\)](#), [metasync\(1M\)](#), [metattach\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [md.tab\(4\)](#), [attributes\(5\)](#), [md\(7D\)](#)

Limitations

Recursive mirroring is not allowed; that is, a mirror cannot appear in the definition of another mirror.

Recursive logging is not allowed; that is, a trans metadvice cannot appear in the definition of another metadvice.

Stripes, concatenations, and RAID level 5 metadevices must consist of slices only.

Mirroring of RAID level 5 metadevices is not allowed.

Soft partitions can be built on raw devices, or on stripes, RAID level 5, or mirrors.

RAID level 5 or stripe metadevices can be built directly on soft partitions.

附注

Trans metadevices have been replaced by UFS logging. Existing trans devices are *not* logging--they pass data directly through to the underlying device. See [mount_ufs\(1M\)](#) for more information about UFS logging.

引用名 metaoffline, metaonline – place submirrors offline and online

用法概要

```
/usr/sbin/metaoffline -h
/usr/sbin/metaoffline [-s setname] [-f] mirror submirror
/usr/sbin/metaonline -h
/usr/sbin/metaonline [-s setname] mirror submirror
```

描述

The `metaoffline` command prevents Solaris Volume Manager from reading and writing to the submirror that has been taken offline. While the submirror is offline, all writes to the mirror will be kept track of (by region) and will be written when the submirror is brought back online. The `metaoffline` command can also be used to perform online backups: one submirror is taken offline and backed up while the mirror remains accessible. (However, if this is a two-way mirror, data redundancy is lost while one submirror is offline.) The `metaoffline` command differs from the `metadetach` command because it does not sever the logical association between the submirror and the mirror. To completely remove a submirror from a mirror, use the `metadetach` command.

A submirror that has been taken offline will only remain offline until the `metaonline` command is invoked or the system is rebooted.

When the `metaonline` command is used, reading from and writing to the submirror resumes. A `resync` is automatically invoked to resync the regions written while the submirror was offline. Writes are directed to the submirror during `resync`. Reads, however, will come from a different submirror. Once the `resync` operation completes, reads and writes are performed on that submirror. The `metaonline` command is only effective on a submirror of a mirror that has been taken offline.

The `metaoffline` and `metaonline` commands can not be used on RAID 1 volumes in application-based recovery (ABR) mode.

A submirror that has been taken offline with the `metaoffline` command can only be mounted as read-only.

选项 Root privileges are required for all of the following options except `-h`.

- `-f` Forces offlining of submirrors that have slices requiring maintenance.
- `-h` Displays usage message.
- `-s setname` Specifies the name of the diskset on which `metaoffline` and `metaonline` will work. Using the `-s` option will cause the command to perform its administrative function within the specified diskset. Without this option, the command will perform its function on local metadevices.
- mirror* Specifies the metadvice name of the mirror from which the submirror will be either taken offline or put online.

submirror Specifies the metadevice name of the submirror to be either taken offline or put online.

示例

示例 1 Taking a Submirror Offline
This example takes one submirror, *mymirror_sub1*, offline from mirror *mymirror*.

```
# metaoffline mymirror mymirror_sub1
```

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见

[mdmonitord\(1M\)](#), [metaclear\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metahs\(1M\)](#), [metainit\(1M\)](#), [metaparam\(1M\)](#), [metarecover\(1M\)](#), [metarename\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metassist\(1M\)](#), [metastat\(1M\)](#), [metasync\(1M\)](#), [metattach\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [md.tab\(4\)](#), [attributes\(5\)](#), [md\(7D\)](#)

[《Solaris Volume Manager Administration Guide》](#)

附注

The *metaonline* and *metaoffline* commands are not applicable to mirrors in application-based recovery (ABR) mode.

引用名 metaparam – modify parameters of metadevices

用法概要

```
/usr/sbin/metaparam -h  
  
/usr/sbin/metaparam [-s setname]  
                    [concat/stripe or RAID5 options] concat/stripe RAID  
  
/usr/sbin/metaparam [-s setname] [mirror options] mirror
```

描述

The `metaparam` command is used to display or modify current parameters of metadevices.

If just the metadevice is specified as an argument to the `metaparam` command, the current settings are displayed.

The `metaparam` command enables most metadevice (volume) parameters to be changed. Only the interlace value cannot be changed by `metaparam`, because it is established when the metadevice is created and cannot be changed thereafter.

选项

Root privileges are required for all of the options.

The following options are supported:

`-h`

Displays usage message.

`-s setname`

Specify the name of the diskset on which `metaparam` works. Using the `-s` option causes the command to perform its administrative function within the specified diskset. Without this option, the command performs its function on local metadevices.

**Concat/STRIPE Or
Raid5 Options**

`-hhot_spare_pool | none`

Specifies the hot spare pool to be used by a metadevice. If `none` is specified, the metadevice is disassociated with the hot spare pool assigned to it. If the metadevice is currently using a hot spare, then `metaparam` cannot replace the hot spare pool.

`concat/stripe | RAID`

Specifies the metadevice name of the concatenation, stripe, or concatenation of stripes, or of the RAID5 metadevice.

Mirror Options

`-r roundrobin | geometric | first`

Modifies the read option for a mirror. The `-r` option must be followed by either `roundrobin`, `geometric`, or `first`. `roundrobin`, which is the default action under the `metainit` command, specifies reading the disks in a round-robin (load balancing) method. `geometric` allows for faster performance on sequential reads. `first` specifies reading only from the first submirror.

`-w parallel | serial`

Modifies the write option for a mirror. The `-w` option must be followed by either `parallel` or `serial`. `parallel`, the default action under the `metainit` command, specifies that all writes are parallel. `serial` specifies that all writes are serial.

-p pass_number
A number from 0-to-9 that specifies the order in which a mirror is resynced during reboot. The default is 1. Smaller pass numbers are resynced first. Equal pass numbers are run concurrently. If 0 is used, the mirror resync is skipped. 0 should only be used for mirrors mounted as read-only, or as swap.

mirror
Specifies the metadevice name of the mirror.

示例

示例 1 Associating Hot Spare Pool with RAID5 Metadevice
This example associates a hot spare pool, `user_pool`, with a RAID5 metadevice, `user_raid`.
`# metaparam -h user_pool user_raid`

示例 2 Changing Read Option to Geometric
This example changes the read option on a mirror `d50` from the default of `roundrobin` to `geometric`.
`# metaparam -r geometric d50`

退出状态

The following exit values are returned:
`0`
Successful completion.
`>0`
An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见

[mdmonitord\(1M\)](#), [metaclear\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metahs\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metarecover\(1M\)](#), [metarename\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metassist\(1M\)](#), [metastat\(1M\)](#), [metasync\(1M\)](#), [metattach\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [md.tab\(4\)](#), [attributes\(5\)](#), [md\(7D\)](#)

[《Solaris Volume Manager Administration Guide》](#)

引用名	metarecover – recover soft partition information
用法概要	<pre>/usr/sbin/metarecover [-n] [-v] [-s <i>setname</i>] <i>component</i> -p /usr/sbin/metarecover [-n] [-v] [-s <i>setname</i>] <i>component</i> -p {-d} /usr/sbin/metarecover [-n] [-v] [-s <i>setname</i>] <i>component</i> -p {-m}</pre>
描述	The metarecover command scans a specified component to look for soft partition configuration information and to regenerate the configuration.
选项	The following options are supported: -d Recover soft partitions in the metadvice state database from the extent headers on the device. Options -d and -m are mutually exclusive. -m Regenerate the extent headers and reapplies them to the underlying device based on the soft partitions listed in the metadvice state database. Options -d and -m are mutually exclusive. -n Do not actually perform the operation. Show the output or errors that would have resulted from the operation, had it been run. -p Regenerate soft partitions based on the metadvice state database or extent headers on the underlying device. If neither -d nor -m are specified, this option compares the soft partition information in the metadvice state database to the extent headers. -s <i>setname</i> Specify the name of the diskset on which metarecover works. Using the s option causes the command to perform its function within the specified diskset. Without the -s option, the metarecover command operates on the metadevices and/or hot spare pools in the local diskset. This option is required to recover former sps from a diskset component or raw-device. <i>setname</i> must be identical to the former <i>setname</i> in which the sps were created. The set numbers, however, seem irrelevant. -v Verbose mode, displaying the changes being made.
操作数	The following operand is supported: <i>component</i> Specifies the c*t*d*s* number of the disk or slice containing the partitions, or the device name (for example, d10) of the metadvice containing the partitions. <i>component</i> can be a slice name, component name, /dev/dsk path, or /dev/rdisk path.

示例

示例 1 Updating Metadevice State Database Based on Disk Extent Headers

A disk containing soft partitions is moved from one system to another. The system administrator would like to use the existing soft partitions. `metarecover` updates the metadevice state database based on the extent headers on the disk.

```
# metarecover -v c0t3d0s2 -p -d
```

示例 2 Updating Metadevice State Database Based on Incomplete Soft Partition Creation

A system crashes in the middle of creating a new soft partition. The soft partition is in the creating state and the driver does not let that device be opened. `metarecover` rewrites the extent headers for the partially created soft partition and mark it as Okay.

```
# metarecover -v c0t3d0s2 -p -m
```

示例 3 Updating Extent Headers Based on Metadevice State Database

Someone accidentally overwrote a portion of a disk leaving extent headers destroyed. `metarecover` rewrites the extent headers to ensure a valid soft partition configuration, though user data is not recovered.

```
# metarecover -v d5 -m
```

The following example implements the same command using a descriptive name.

```
# metarecover -v myvolume -m
```

示例 4 Validating Soft Partition Configuration

To validate the existing soft partition configuration, use `metarecover` with only the `-p` flag.

```
# metarecover c0t3d0s2 -p
```

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见

[mdmonitord\(1M\)](#), [metaclear\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metahs\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metarename\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metassist\(1M\)](#), [metastat\(1M\)](#), [metasync\(1M\)](#), [metattach\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [md.tab\(4\)](#), [attributes\(5\)](#), [md\(7D\)](#)

《Solaris Volume Manager Administration Guide》

引用名	metarename – rename metadvice or switch layered metadvice names
用法概要	<pre> /usr/sbin/metarename [-s setname] metadvice1 metadvice2 /usr/sbin/metarename [-s setname] [-f] -x metadvice1 metadvice2 /usr/sbin/metarename -h </pre>
描述	There are two ways to use <code>metarename</code>, one with and one without the <code>-x</code> option. The first method (without <code>-x</code>) renames an existing metadvice to a new name. This makes managing the metadvice namespace easier. The metadvice being renamed cannot be mounted or open, nor can the new name already exist. For example, to rename a metadvice that contains a mounted file system, you would first need to unmount the file system. With the second way to use <code>metarename</code>, using the <code>-x</code> option, <code>metarename</code> switches (exchanges) the names of an existing layered metadvice and one of its subdevices. In Solaris Volume Manager terms, a layered metadvice can be either a mirror or a trans metadvice. The <code>-x</code> option enables you to switch the metadvice names of a mirror and one of its submirrors, or a trans metadvice and its master device. <code>metarename -x</code> makes it easier to mirror or unmirror an existing stripe or concatenation, and to remove a trans device. When used to mirror an existing stripe or concatenation, you must stop access to the device. For example, if the device contains a mounted file system, you must first unmount the file system before doing the rename. You can also use the <code>metarename -x</code> command to untrans a trans metadvice from an existing device. This applies only to the master device. You cannot remove a logging device with <code>metarename</code>. Before you can rename a trans device, you must detach the logging device. Then you must stop access to the trans metadvice itself. You cannot rename or switch metadvice devices that are in an error state or that have subcomponents in an error state, or metadvice devices actively using a hot spare replacement. You can only switch metadvice devices that have a direct child/parent relationship. You could not, for example, directly exchange a stripe in a mirror that is a master device with the trans metadvice. You must use the <code>-f</code> flag when switching members of a trans metadvice. Only metadvice devices can be switched, not slices.
选项	The following options are supported: -f Force the switching of trans metadvice members. -h Display a help message.

- `-s setname` Specifies the name of the diskset on which `metarename` will work. Using the `-s` option will cause the command to perform its administrative function within the specified diskset. Without this option, the command will perform its function on the local metadevices.
- `-x` Exchange the metadevice names *metadevice1* and *metadevice2*.
- metadevice1* Specifies the metadevice to be renamed or switched.
- metadevice2* Specifies the target metadevice name for the rename or switch operation.

示例

示例 1 Renaming a Metadevice

This example renames a metadevice named `d10` to `account_records`. Note that `account_records` must not exist for the rename to succeed.

```
# metarename d10 account_records
```

示例 2 Creating a Two-Way Mirror

This example creates a two-way mirror from an existing stripe named `d1` with a mounted file system, `/home2`.

```
# metainit d2 1 1 c13d0s1
# metainit -f d20 -m d1
# umount /home2
# metarename -x d20 d1
# metattach d1 d2
# mount /home2
```

First, a second concatenation `d2`, is created. (`d1` already exists.) The `metainit` command creates a one-way mirror, `d20`, from `d1`. Next, you unmount the file system and switch `d1` for `d20`, making `d1` the top-level device (mirror). You attach the second submirror, `d2`, to create a two-way mirror. Lastly, you remount the file system.

示例 3 Mounting a Mirrored File System on Stripe

This example takes an existing mirror named `d1` with a mounted file system, and ends up with the file system mounted on a stripe `d1`.

```
# umount /fs2
# metarename -x d1 d20
# metadetach d20 d1
# metaclear -r d20
# mount /fs2
```

First, you unmount the file system, then switch the mirror `d1` and its submirror `d20`. This makes the mirror into `d20`. Next, you detach `d1` from `d20`, then delete the mirror `d20` and its other submirror. You then remount the file system.

示例 4 Deleting a Trans Metadevice

This example deletes a trans metadevice named `d10` while its mount point is `/myhome`. The master device, which is a stripe, is named `d2`. The logging device, also a stripe, is named `d5`.

```
# umount /myhome
# metadetach d10
# metarename -f -x d10 d2
# metaclear d2
# metaclear d5
# fsck /dev/md/dsk/d10
# mount /myhome
```

You unmount the file system first, then detach the trans metadevice's logging device. The trans metadevice is switched with the master device, making the trans metadevice `d2` and the underlying stripe `d10`. You clear the trans metadevice `d2` and the logging device `d5`. `d10` must be `fsck`'d, and then the file system is remounted.

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见

[mdmonitord\(1M\)](#), [metaclear\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metahs\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metarecover\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metassist\(1M\)](#), [metastat\(1M\)](#), [metasync\(1M\)](#), [metattach\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [md.tab\(4\)](#), [attributes\(5\)](#), [md\(7D\)](#)

《Solaris Volume Manager Administration Guide》

Limitations

Renaming and exchanging metadevice names can only be used for metadevices. A physical slice cannot be renamed to a metadevice, nor can a metadevice be exchanged with a physical slice name.

Metadevice names are strings of the pattern `d<xyz>` where `xyz` is a value between 0 and 8192. You cannot use logical names for metadevices.

附注

Trans metadevices have been replaced by UFS logging. Existing trans devices are *not* logging--they pass data directly through to the underlying device. See [mount_ufs\(1M\)](#) for more information about UFS logging.

引用名	metareplace – enable or replace components of submirrors or RAID5 metadevices								
用法概要	<pre>/usr/sbin/metareplace -h /usr/sbin/metareplace [-s setname] -e mirror component /usr/sbin/metareplace [-s setname] mirror component-old component-new /usr/sbin/metareplace [-s setname] -e RAID component /usr/sbin/metareplace [-s setname] [-f] RAID component-old component-new</pre>								
描述	The <code>metareplace</code> command is used to enable or replace components (slices) within a submirror or a RAID5 metadevice. When you replace a component, the <code>metareplace</code> command automatically starts resyncing the new component with the rest of the metadevice. When the resync completes, the replaced component becomes readable and writable. If the failed component has been hot spare replaced, the hot spare is placed in the available state and made available for other hot spare replacements. Note that the new component must be large enough to replace the old component. A component may be in one of several states. The <code>Last Erred</code> and the <code>Maintenance</code> states require action. Always replace components in the <code>Maintenance</code> state first, followed by a resync and validation of data. After components requiring maintenance are fixed, validated, and resynced, components in the <code>Last Erred</code> state should be replaced. To avoid data loss, it is always best to back up all data before replacing <code>Last Erred</code> devices.								
选项	Root privileges are required for all of the following options except <code>-h</code>. <table><tr><td>-e</td><td>Transitions the state of <i>component</i> to the available state and resyncs the failed component. If the failed component has been hot spare replaced, the hot spare is placed in the available state and made available for other hot spare replacements. This command is useful when a component fails due to human error (for example, accidentally turning off a disk), or because the component was physically replaced. In this case, the replacement component must be partitioned to match the disk being replaced before running the <code>metareplace</code> command.</td></tr><tr><td>-f</td><td>Forces the replacement of an errored component of a metadevice in which multiple components are in error. The component determined by the <code>metastat display</code> to be in the “Maintenance” state must be replaced first. This option may cause data to be fabricated since multiple components are in error.</td></tr><tr><td>-h</td><td>Display help message.</td></tr><tr><td>-s <i>setname</i></td><td>Specifies the name of the diskset on which <code>metareplace</code> will work. Using the <code>-s</code> option will cause the command to perform its administrative</td></tr></table>	-e	Transitions the state of <i>component</i> to the available state and resyncs the failed component. If the failed component has been hot spare replaced, the hot spare is placed in the available state and made available for other hot spare replacements. This command is useful when a component fails due to human error (for example, accidentally turning off a disk), or because the component was physically replaced. In this case, the replacement component must be partitioned to match the disk being replaced before running the <code>metareplace</code> command.	-f	Forces the replacement of an errored component of a metadevice in which multiple components are in error. The component determined by the <code>metastat display</code> to be in the “Maintenance” state must be replaced first. This option may cause data to be fabricated since multiple components are in error.	-h	Display help message.	-s <i>setname</i>	Specifies the name of the diskset on which <code>metareplace</code> will work. Using the <code>-s</code> option will cause the command to perform its administrative
-e	Transitions the state of <i>component</i> to the available state and resyncs the failed component. If the failed component has been hot spare replaced, the hot spare is placed in the available state and made available for other hot spare replacements. This command is useful when a component fails due to human error (for example, accidentally turning off a disk), or because the component was physically replaced. In this case, the replacement component must be partitioned to match the disk being replaced before running the <code>metareplace</code> command.								
-f	Forces the replacement of an errored component of a metadevice in which multiple components are in error. The component determined by the <code>metastat display</code> to be in the “Maintenance” state must be replaced first. This option may cause data to be fabricated since multiple components are in error.								
-h	Display help message.								
-s <i>setname</i>	Specifies the name of the diskset on which <code>metareplace</code> will work. Using the <code>-s</code> option will cause the command to perform its administrative								

function within the specified diskset. Without this option, the command will perform its function on local metadevices.

- mirror* The metadvice name of the mirror.
- component* The logical name for the physical slice (partition) on a disk drive, such as /dev/dsk/c0t0d0s2.
- component-old* The physical slice that is being replaced.
- component-new* The physical slice that is replacing *component-old*.
- RAID* The metadvice name of the RAID5 device.

示例

示例 1 Recovering from Error Condition in RAID5 Metadvice

This example shows how to recover when a single component in a RAID5 metadvice is errored.

```
# metareplace d10 c3t0d0s2 c5t0d0s2
```

In this example, a RAID5 metadvice d10 has an errored component, c3t0d0s2, replaced by a new component, c5t0d0s2.

示例 2 Use of -e After Physical Disk Replacement

This example shows the use of the -e option after a physical disk in a submirror (a submirror of mirror mymirror1, in this case) has been replaced.

```
# metareplace -e mymirror1 c1t4d0s2
```

Note: The replacement disk must be partitioned to match the disk it is replacing before running the metareplace command.

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见

[mdmonitord\(1M\)](#), [metaclear\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metahs\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metarecover\(1M\)](#),

metarename(1M), metaset(1M), metassist(1M), metastat(1M), metasync(1M),
metattach(1M), md.tab(4), md.cf(4), mddb.cf(4), md.tab(4), attributes(5), md(7D)

《Solaris Volume Manager Administration Guide》

引用名

metaset – configure disk sets

用法概要

```

/usr/sbin/metaset -s setname [-M-a -h hostname]
/usr/sbin/metaset -s setname -A {enable | disable}
/usr/sbin/metaset -s setname [-A {enable | disable}] -a -h hostname...
/usr/sbin/metaset -s setname -a [-l length] [-L] drivename...
/usr/sbin/metaset -s setname -C {take | release | purge}
/usr/sbin/metaset -s setname -d [-f] -h hostname...
/usr/sbin/metaset -s setname -d [-f] drivename...
/usr/sbin/metaset -s setname -j
/usr/sbin/metaset -s setname -r
/usr/sbin/metaset -s setname -w
/usr/sbin/metaset -s setname -t [-f] [-u tagnumber] [y]
/usr/sbin/metaset -s setname -b
/usr/sbin/metaset -s setname -P
/usr/sbin/metaset -s setname -q
/usr/sbin/metaset -s setname -o [-h hostname]
/usr/sbin/metaset [-s setname]
/usr/sbin/metaset [-s setname] -a | -d
    [ [m] mediator_host_list]

```

描述

The metaset command administers sets of disks in named disk sets. Named disk sets include any disk set that is not in the local set. While disk sets enable a high-availability configuration, Solaris Volume Manager itself does not actually provide a high-availability environment.

A single-owner disk set configuration manages storage on a SAN or fabric-attached storage, or provides namespace control and state database replica management for a specified set of disks.

In a shared disk set configuration, multiple hosts are physically connected to the same set of disks. When one host fails, another host has exclusive access to the disks. Each host can control a shared disk set, but only one host can control it at a time.

When you add a new disk to any disk set, Solaris Volume Manager checks the disk format. If necessary, it repartitions the disk to ensure that the disk has an appropriately configured reserved slice 7 (or slice 6 on an EFI labelled device) with adequate space for a state database replica. The precise size of slice 7 (or slice 6 on an EFI labelled device) depends on the disk geometry. For traditional disk sets, the slice is no less than 4 Mbytes, and probably closer to 6 Mbytes, depending on where the cylinder boundaries lie. For multi-owner disk sets, the slice is

a minimum of 256 Mbytes. The minimal size for slice 7 might change in the future. This change is based on a variety of factors, including the size of the state database replica and information to be stored in the state database replica.

For use in disk sets, disks must have a dedicated slice (six or seven) that meets specific criteria:

- The slice must start at sector 0
- The slice must include enough space for disk label
- The state database replicas cannot be mounted
- The slice does not overlap with any other slices, including slice 2

If the existing partition table does not meet these criteria, or if the `-L` flag is specified, Solaris Volume Manager repartitions the disk. A small portion of each drive is reserved in slice 7 (or slice 6 on an EFI labelled device) for use by Solaris Volume Manager. The remainder of the space on each drive is placed into slice 0. Any existing data on the disks is lost by repartitioning.

After you add a drive to a disk set, it can be repartitioned as necessary, with the exception that slice 7 (or slice 6 on an EFI labelled device) is not altered in any way.

After a disk set is created and metadevices are set up within the set, the metadvice name is in the following form:

```
/dev/md/setname{/dsk,rdisk}/dnumber
```

where *setname* is the name of the disk set, and *number* is the number of the metadvice (0-127).

If you have disk sets that you upgraded from Solstice DiskSuite software, the default state database replica size on those sets is 1034 blocks, not the 8192 block size from Solaris Volume Manager. Also, slice 7 on the disks that were added under Solstice DiskSuite are correspondingly smaller than slice 7 on disks that were added under Solaris Volume Manager.

If disks you add to a disk set have acceptable slice 7s (that start at cylinder 0 and that have sufficient space for the state database replica), they are not reformatted.

Hot spare pools within local disk sets use standard Solaris Volume Manager naming conventions. Hot spare pools with shared disk sets use the following convention:

```
setname/hot_spare_pool
```

where *setname* is the name of the disk set, and *hot_spare_pool* is the name of the hot spare pool associated with the disk set.

Multi-node Environment

To create and work with a disk set in a multi—node environment, `root` must be a member of Group 14 on all hosts, or the `/etc/rhosts` file must contain an entry for all other host names. This is not required in a SunCluster 3.x environment.

Tagged data Tagged data occurs when there are different versions of a disk set's replicas. This tagged data consists of the set owner's nodename, the hardware serial number of the owner, and the time it was written out to the available replicas. The system administrator can use this information to determine which replica contains the correct data.

When a disk set is configured with an even number of storage enclosures and has replicas balanced across them evenly, it is possible that up to half of the replicas can be lost (for example, through a power failure of half of the storage enclosures). After the enclosure that went down is rebooted, half of the replicas are not recognized by SVM. When the set is retaken, the `metaset` command returns an error of “stale databases”, and all of the metadevices are in a read-only state.

Some of the replicas that are not recognized need to be deleted. The action of deleting the replicas also causes updates to the replicas that are not being deleted. In a dual hosted disk set environment, the second node can access the deleted replicas instead of the existing replicas when it takes the set. This leads to the possibility of getting the wrong replica record on a disk set take. An error message is displayed, and user intervention is required.

Use the `-q` to query the disk set and the `-t`, `-u`, and `-y`, options to select the tag and take the disk set. See `OPTIONS`.

Mediator Configuration SVM provides support for a low-end HA solution consisting of two hosts that share only two strings of drives. The hosts in this type of configuration, referred to as *mediators* or mediator hosts, run a special daemon, `rpc.metamedd(1M)`. The mediator hosts take on additional responsibilities to ensure that data is available in the case of host or drive failures.

A mediator configuration can survive the failure of a single host or a single string of drives, without administrative intervention. If both a host and a string of drives fail (multiple failures), the integrity of the data cannot be guaranteed. At this point, administrative intervention is required to make the data accessible. See [mediator\(7D\)](#) for further details.

Use the `-m` option to add or delete a mediator host. See `OPTIONS`.

选项

The following options are supported:

`-a drivename`

Add drives or hosts to the named set. For a drive to be accepted into a set, the drive must not be in use within another metadevice or disk set, mounted on, or swapped on. When the drive is accepted into the set, it is repartitioned and the metadevice state database replica (for the set) can be placed on it. However, if a slice 7 (or slice 6 on an EFI labelled device), starts at cylinder 0, and is large enough to hold a state database replica, then the disk is not repartitioned. Also, a drive is not accepted if it cannot be found on all hosts specified as part of the set. This means that if a host within the specified set is unreachable due to network problems, or is administratively down, the add fails.

Specify a drive name in the form `cnumtnumdnum`. Do not specify a slice number (`snum`). For drives in a Sun Cluster, you must specify a complete pathname for each drive. Such a name has the form:

/dev/did/[r]dsk/dnum

-a | -d | -m *mediator_host_list*

Add (-a) or delete (-d) mediator hosts to the specified disk set. A *mediator_host_list* is the nodename (established by `svc:/system/identity:node smf(5)` service) of the mediator host to be added and (for adding) up to two other aliases for the mediator host. The nodename and aliases for each mediator host are separated only by commas. Up to three mediator hosts can be specified for the named disk set. Specify only the nodename of that host as the argument to -m to delete a mediator host.

In a single metaset command you can add or delete up to three mediator hosts. See EXAMPLES.

-A {enable | disable}

Specify auto-take status for a disk set. If auto-take is enabled for a set, the disk set is automatically taken at boot, and file systems on volumes within the disk set can be mounted through `/etc/vfstab` entries. Only a single host can be associated with an auto-take set, so attempts to add a second host to an auto-take set or attempts to configure a disk set with multiple hosts as auto-take fails with an error message. Disabling auto-take status for a specific disk set causes the disk set to revert to normal behavior. That is, the disk set is potentially shared (non-concurrently) among hosts, and unavailable for mounting through `/etc/vfstab`.

-b

Insure that the replicas are distributed according to the replica layout algorithm. This can be invoked at any time, and does nothing if the replicas are correctly distributed. In cases where the user has used the `metadb` command to manually remove or add replicas, this command can be used to insure that the distribution of replicas matches the replica layout algorithm.

-C {take | release | purge}

Do not interact with the Cluster Framework when used in a Sun Cluster 3 environment. In effect, this means do not modify the Cluster Configuration Repository. These options should only be used to fix a broken disk set configuration.

take

Take ownership of the disk set but do not inform the Cluster Framework that the disk set is available. This option is not for use with a multi-owner disk set.

release

Release ownership of the disk set without informing the Cluster Framework. This option should only be used if the disk set ownership was taken with the corresponding -C take option. This option is not for use with a multi-owner disk set.

purge

Remove the disk set without informing the Cluster Framework that the disk set has been purged. This option should only be used when the disk set is not accessible and requires rebuilding.

-d *drivename*

Delete drives or hosts from the named disk set. For a drive to be deleted, it must not be in use within the set. The last host cannot be deleted unless all of the drives within the set are deleted. Deleting the last host in a disk set destroys the disk set.

Specify a drive name in the form *cnumt numdnum*. Do not specify a slice number (*snum*). For drives in a Sun Cluster, you must specify a complete pathname for each drive. Such a name has the form:

```
/dev/did/[r]dsk/dnum
```

This option fails on a multi-owner disk set if attempting to withdraw the master node while other nodes are in the set.

-f

Force one of three actions to occur: takes ownership of a disk set when used with **-t**; deletes the last disk drive from the disk set; or deletes the last host from the disk set. Deleting the last drive or host from a disk set requires the **-d** option.

When used to forcibly take ownership of the disk set, this causes the disk set to be grabbed whether or not another host owns the set. All of the disks within the set are taken over (reserved) and fail fast is enabled, causing the other host to panic if it had disk set ownership. The metadvice state database is read in by the host performing the take, and the shared metadvice contained in the set are accessible.

You can use this option to delete the last drive in the disk set, because this drive would implicitly contain the last state database replica.

You can use **-f** option to delete hosts from a set. When specified with a partial list of hosts, it can be used for one-host administration. One-host administration could be useful when a host is known to be non-functional, thus avoiding timeouts and failed commands. When specified with a complete list of hosts, the set is completely deleted. It is generally specified with a complete list of hosts to clean up after one-host administration has been performed.

-h *hostname...*

Specify one or more host names to be added to or deleted from a disk set. Adding the first host creates the set. The last host cannot be deleted unless all of the drives within the set have been deleted. The host name is not accepted if all of the drives within the set cannot be found on the specified host. The host name is the nodename established by the `svc:/system/identity:node` [smf\(5\)](#) service.

-j

Join a host to the owner list for a multi-owner disk set. The concepts of take and release, used with traditional disk sets, do not apply to multi-owner sets, because multiple owners are allowed.

As a host boots and is brought online, it must go through three configuration levels to be able to use a multi-owner disk set:

1. It must be included in the cluster nodelist, which happens automatically in a cluster or single-node situation.
2. It must be added to the multi-owner disk set with the `-a -h` options documented elsewhere in this man page
3. It must join the set. When the host is first added to the set, it is automatically joined.

On manual restarts, the administrator must manually issue

```
metaset -s multinodesetname -j
```

to join the host to the owner list. After the cluster reconfiguration, when the host reenters the cluster, the node is automatically joined to the set. The `metaset -j` command joins the host to all multi-owner sets that the host has been added to. In a single node situation, joining the node to the disk set starts any necessary resynchronizations.

-L

When adding a disk to a disk set, force the disk to be repartitioned using the standard Solaris Volume Manager algorithm. See DESCRIPTION.

-l *length*

Set the size (in blocks) for the metadevice state database replica. The length can only be set when adding a new drive; it cannot be changed on an existing drive. The default (and maximum) size is 8192 blocks, which should be appropriate for most configurations. Replica sizes of less than 128 blocks are not recommended.

-M

Specify that the disk set to be created or modified is a multi-owner disk set that supports multiple concurrent owners.

This option is required when creating a multi-owner disk set. Its use is optional on all other operations on a multi-owner disk set and has no effect. Existing disk sets cannot be converted to multi-owner sets.

-o

Return an exit status of 0 if the local host or the host specified with the `-h` option is the owner of the disk set.

-P

Purge the named disk set from the node on which the `metaset` command is run. The disk set must not be owned by the node that runs this command. If the node does own the disk set, the command fails.

If you need to delete a disk set but cannot take ownership of the set, use the `-P` option.

-q

Displays an enumerated list of tags pertaining to "tagged data" that can be encountered during a take of the ownership of a disk set.

This option is not for use with a multi-owner disk set.

-r

Release ownership of a disk set. All of the disks within the set are released. The metadevices set up within the set are no longer accessible.

This option is not for use with a multi-owner disk set.

-s *setname*

Specify the name of a disk set on which `metaset` works. If no *setname* is specified, all disk sets are returned.

-t

Take ownership of a disk set safely. If `metaset` finds that another host owns the set, this host is not be allowed to take ownership of the set. If the set is not owned by any other host, all the disks within the set are owned by the host on which `metaset` was executed. The metadvice state database is read in, and the shared metadevices contained in the set become accessible. The -t option takes a disk set that has stale databases. When the databases are stale, `metaset` exits with code 66, and prints a message. At that point, the only operations permitted are the addition and deletion of replicas. Once the addition or deletion of the replicas has been completed, the disk set should be released and retaken to gain full access to the data.

This option is not for use with a multi-owner disk set.

-u *tagnumber*

Once a tag has been selected, a subsequent take with -u *tagnumber* can be executed to select the data associated with the given *tagnumber*.

w

Withdraws a host from the owner list for a multi-owner disk set. The concepts of take and release, used with traditional disk sets, do not apply to multi-owner sets, because multiple owners are allowed.

Instead of releasing a set, a host can issue

```
metaset -s multinodesetname -w
```

to withdraw from the owner list. A host automatically withdraws on a reboot, but can be manually withdrawn if it should not be able to use the set, but should be able to rejoin at a later time. A host that withdrew due to a reboot can still appear joined from other hosts in the set until a reconfiguration cycle occurs.

`metaset -w` withdraws from ownership of all multi-owner sets of which the host is a member. This option fails if you attempt to withdraw the master node while other nodes are in the disk set owner list. This option cancels all resyncs running on the node. A cluster reconfiguration process that is removing a node from the cluster membership list effectively withdraws the host from the ownership list.

-y

Execute a subsequent take. If the take operation encounters “tagged data,” the take operation exits with code 2. You can then run the `metaset` command with the `-q` option to see an enumerated list of tags.

示例

示例 1 Defining a Disk Set

This example defines a disk set.

```
# metaset -s relo-red -a -h red blue
```

The name of the disk set is `relo-red`. The names of the first and second hosts added to the set are `red` and `blue`, respectively. (The hostname is the nodename established by `svc:/system/identity:node` [smf\(5\)](#) service.) Adding the first host creates the disk set. A disk set can be created with just one host, with the second added later. The last host cannot be deleted until all of the drives within the set have been deleted.

示例 2 Adding Drives to a Disk Set

This example adds drives to a disk set.

```
# metaset -s relo-red -a c2t0d0 c2t1d0 c2t2d0 c2t3d0 c2t4d0 c2t5d0
```

The name of the previously created disk set is `relo-red`. The names of the drives are `c2t0d0`, `c2t1d0`, `c2t2d0`, `c2t3d0`, `c2t4d0`, and `c2t5d0`. There is no slice identifier (“`sx`”) at the end of the drive names.

示例 3 Adding Multiple Mediator Hosts

The following command adds three mediator hosts to the specified disk set.

```
# metaset -s mydiskset -a -m myhost1,alias1 myhost2,alias2 myhost3,alias3
```

示例 4 Purging a Disk Set from the Node

The following command purges the disk set `relo-red` from the node:

```
# metaset -s relo-red -P
```

示例 5 Querying a Disk Set for Tagged Data

The following command queries the disk set `relo-red` for a list of the tagged data:

```
# metaset -s relo-red -q
```

This command produces the following results:

The following tag(s) were found:

```
1 - vha-1000c - Fri Sep 20 17:20:08 2002
2 - vha-1000c - Mon Sep 23 11:01:27 2002
```

示例 6 Selecting a tag and taking a Disk set

The following command selects a tag and takes the disk set `relo-red`:

```
# metaset -s relo-red -t -u 2
```

示例 7 Defining a Multi-Owner Disk Set

The following command defines a multi-owner disk set:

```
# metaset -s blue -M -a -h hahost1 hahost2
```

The name of the disk set is `blue`. The names of the first and second hosts added to the set are `hahost1` and `hahost2`, respectively. The hostname is the nodename established by `svc:/system/identity:node` [smf\(5\)](#) service. Adding the first host creates the multi-owner disk set. A disk set can be created with just one host, with additional hosts added later. The last host cannot be deleted until all of the drives within the set have been deleted.

文件

`/etc/lvm/md.tab`
Contains list of metadvice configurations.

退出状态

The following exit values are returned:

0
Successful completion.

>0
An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见

[mdmonitord\(1M\)](#), [metaclear\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metahs\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metarecover\(1M\)](#), [metarename\(1M\)](#), [metareplace\(1M\)](#), [metassist\(1M\)](#), [metastat\(1M\)](#), [metasync\(1M\)](#), [metattach\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [md.tab\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [md\(7D\)](#)

《Solaris Volume Manager Administration Guide》

附注

Disk set administration, including the addition and deletion of hosts and drives, requires all hosts in the set to be accessible from the network.

引用名

metassist – automated volume creation utility to support Solaris Volume Manager

用法概要

```
metassist -V
metassist -?

metassist create [-v n] [-c] -F config_file
metassist create [-v n] [-c | -d] -F request_file
metassist create [-v n] [-c | -d] [-f] [-n name]
    [-p datapaths] [-r redundancy]
    [-a available [,available,...]]
    [-u unavailable [,unavailable,...]] -s setname -S size

metassist create -?
```

描述

The `metassist` command provides assistance, through automation, with common Solaris Volume Manager tasks.

SUBCOMMANDS

The following subcommands are supported:

create The `create` subcommand creates one or more Solaris Volume Manager volumes. You can specify this request on the command line or in a file specified on the command line.

If you create a volume using the command line, you can specify the characteristics of the volume in terms of the desired quality of service it will provide - its size, the number of redundant copies of the data it contains, the number of data paths by which it is accessible, and whether faulty components are replaced automatically. The diskset in which the volume will reside and the volume's size must be specified on the command line in this form of the command.

If you create a volume using a request in a file, you can specify the characteristics of the volume in terms of the quality of service they provide, as on the command line. Alternatively, the file can specify the types and component parts of the volume, (for example, mirrors, stripes, concatenations, and their component slices). The file may also specify volumes partly in terms of their types and partly in terms of their component parts, and may specify the characteristics of more than one volume. All volumes specified in a file must reside in the same diskset, whose name must be specified in the file.

If you specify the `-c` or `-d` option on the command line, the command runs without creating an actual volume or volumes. Instead, it outputs either a Bourne shell command script (`-c` option) or a volume configuration (`-d` option). The command script, when run, creates the specified volume or volumes. The volume configuration specifies the volume or volumes in complete detail, naming all their components.

The input file given on the command line can take one of the following forms:

选项

- a volume request, which specifies a request for a volume with explicit attributes and components, or matching a given quality of service
- a volume configuration, produced by a previous execution of the command

The following option is mandatory if you specify a volume request or volume configuration in a file:

-F *config_file* | *request_file* Specify the volume request or volume configuration file to process. If *config_file* or *request_file* is -, it is read from standard input.

The -d option cannot be specified when *inputfile* is a volume configuration file.

The following options are mandatory if you specify a volume request on the command line:

- s *set*** Specify the disk set to use when creating volumes. All the volumes and hot spare pools are created in this disk set. If necessary, disks are moved into the diskset for use in the volumes and hot spare pools. If the diskset doesn't exist the command creates it. This option is required. *metassist* works entirely within a named disk set. Use of the local, or unnamed disk set, is not allowed.
- S *size*** Specify the size of the volume to be created. The size argument consists of a numeric value (a decimal can be specified) followed by KB, MB, GB, or TB, indicating kilobytes, megabytes, gigabytes, or terabytes, respectively. Case is ignored when interpreting this option. This option is required.

The following options are optional command line parameters:

- a *device1, device2, . . .*** Explicitly specify the devices that can be used in the creation of this volume. Named devices may be controllers or disks. Only used when specifying a volume on the command line.
- c** Output the command script that would implement the specified or generated volume configuration. The command script is not run, and processing stops at this stage.
- d** Output the volume configuration that satisfies the specified or generated volume request. No command script is generated or executed, and processing stops at this stage.
- f** Specify whether the volume should support automatic component replacement after a fault. If this option is specified, a mirror is created and its submirrors are associated with a hot spare.
- n *name*** Specify the name of the new volume. See [metainit\(1M\)](#) for naming guidelines.

- `-p n` Specify the number of required paths to the storage volume. The value of *n* cannot be greater than the number of different physical paths and logical paths to attached storage. Only used when specifying a volume on the command line.
- `-r n` Specify the redundancy level (0-4) of the data. The default is 0. Only used when specifying a volume on the command line. If redundancy is 0, a stripe is created. If redundancy is 1 or greater, a mirror with this number of submirrors is created. In this case, the volume can suffer a disk failure on *n* - 1 copies without data loss. With the use of hot spares (see the `-f` option), a volume can suffer a disk failure on *n*+*hsp*s - 1 volumes without data loss, assuming non-concurrent failures.
- `-u device1, device2, . . .` Explicitly specify devices to exclude in the creation of this volume. Named devices can be controllers or disks. You can use this option alone, or to exclude some of the devices listed as available with the `-a` option, Only used when specifying a volume on the command line.
- `-v value` Specify the level of verbosity. Values from 0 to 2 are available, with higher numbers specifying more verbose output when the command is run. `-v 0` indicates silent output, except for errors or other critical messages. The default level is 1.
- `-V` Display program version information.
- `-?` Display help information. This option can follow a subcommand for subcommand-specific help.

示例

示例 1 Creating a Mirror

The following example creates a two-way, 36Gb mirror on available devices from controller 1 and controller 2. It places the volume in `diskset mirrorset`.

```
# metassist create -r 2 -a c1,c2 -s mirrorset -S 36GB
```

示例 2 Creating a Mirror with Additional Fault Tolerance

The following example creates a two-way, 36Gb mirror on available devices from controller 1 and controller 2. It provides additional fault tolerance in the form of a hot spare. It places the volume in `diskset mirrorset`.

```
# metassist create -f -r 2 -a c1,c2 -s mirrorset -S 36GB
```

示例 3 Creating a Three-way Mirror and Excluding Devices

The following example creates a three-way, 180Gb mirror from storage devices on controller 1 or controller 2. It excludes the disks c1t2d0 and c2t2d1 from the volume. It places the volume in diskset mirrorset.

```
metassist create -r 3 -a c1,c2 -u c1t2d0, c2t2d1 \
    -s mirrorset -S 180GB
```

示例 4 Determining and Implementing a Configuration

The following example determines and implements a configuration satisfying the request specified in a request file:

```
# metassist create -F request.xml
```

示例 5 Determining a Configuration and Saving It in a volume-config File

The following example determines a configuration which satisfies the given request. It saves the configuration in a volume-config file without implementing it:

```
# metassist create -d -F request.xml > volume-config
```

示例 6 Determining a Configuration and Saving It in a Shell Script

The following example determines a configuration which satisfies the given request. It saves the configuration in a shell script without implementing it:

```
# metassist create -c -F request.xml > setupvols.sh
```

示例 7 Implementing the Given volume-config

The following example implements the given volume-config:

```
# metassist create -F config.xml
```

示例 8 Converting the Given volume-config to a Shell Script

The following example converts the given volume-config to a shell script that you can run later:

```
# metassist create -c -F config.xml > setupvols.sh
```

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

文件

- /usr/share/lib/xml/dtd/volume-request.dtd
- /usr/share/lib/xml/dtd/volume-defaults.dtd
- /usr/share/lib/xml/dtd/volume-config.dtd

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见

[mdmonitord\(1M\)](#), [metaclear\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metahs\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metarecover\(1M\)](#), [metarename\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metastat\(1M\)](#), [metasync\(1M\)](#), [metattach\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [md.tab\(4\)](#), [volume-config\(4\)](#), [volume-request\(4\)](#), [attributes\(5\)](#), [md\(7D\)](#)

附注

The quality of service arguments are mutually exclusive with the *-F inputfile* argument.

When specifying a request file or quality of service arguments on the command line, the */etc/default/metassist.xml* file is read for global and per-disk set defaults.

Characteristics of this file are specified in the DTD, in */usr/share/lib/xml/dtd/volume-defaults.dtd*.

Characteristics of the XML request file are specified in the DTD, in */usr/share/lib/xml/dtd/volume-request.dtd*.

Characteristics of the XML configuration file are specified in the DTD, in */usr/share/lib/xml/dtd/volume-config.dtd*.

This command must be run as root.

This command requires a functional Solaris Volume Manager configuration before it runs.

引用名 metastat – display status for metadevice or hot spare pool

用法概要

```
/usr/sbin/metastat -h

/usr/sbin/metastat [-a] [-B] [-D] [-c] [-i] [-p] [-q]
                  [-s setname] [-t] [metadevice...] [hot_spare_pool...]

/usr/sbin/metastat [-a] [-B] [-D] [-c] [-i] [-p] [-q]
                  [-s setname] component...
```

描述

The `metastat` command displays the current status for each metadevice (including stripes, concatenations, concatenations of stripes, mirrors, RAID5, soft partitions, and trans devices) or hot spare pool, or of specified metadevices, components, or hot spare pools.

It is helpful to run the `metastat` command after using the `metattach` command to view the status of the metadevice.

`metastat` displays the state of each Solaris Volume Manager RAID-1 volume on the system. The possible states include:

Okay	The device reports no errors.
Needs maintenance	A problem has been detected. This requires that the system administrator replace the failed physical device. Volumes displaying Needs maintenance have incurred no data loss, although additional failures could risk data loss. Take action as quickly as possible.
Last erred	A problem has been detected. Data loss is a possibility. This might occur if a component of a submirror fails and is not replaced by a hot spare, therefore going into Needs maintenance state. If the corresponding component also fails, it would go into Last erred state and, as there is no remaining valid data source, data loss could be a possibility.
Unavailable	A device cannot be accessed, but has not incurred errors. This might occur if a physical device has been removed with Solaris Dynamic Reconfiguration (DR) features, thus leaving the Solaris Volume Manager volume unavailable. It could also occur if an array or disk is powered off at system initialization, or if a >1TB volume is present when the system is booted in 32-bit mode.

After the storage has been made available, run the `metastat` command with the `-i` option to update the status of the metadevices. This clears the unavailable state for accessible devices.

See the [《Solaris Volume Manager Administration Guide》](#) for instructions on replacing disks and handling volumes in Needs maintenance or Last erred states.

- a Display all disk sets. Only metadevices in disk sets that are owned by the current host are displayed.
- B Display the current status of all of the 64-bit metadevices and hot spares.
- c Display concise output.

The -c output format is distinct from the -p output format. The -p option does not display metadvice status and is not intended as human-readable output.

-h Display usage message.

-p Display the list of active metadevices and hot spare pools in the same format as `md . tab`. See [md . tab\(4\)](#).

-q Display the status for metadevices without the device relocation information.

- t Display the current status and timestamp for the specified metadevices and hot spare pools. The timestamp provides the date and time of the last state change.

操作数

<i>component</i>	Display the status of the component hosting a soft partition, including extents, starting blocks, and block count.

1386

metadevice Display the status of the specified metadevice(s). If a trans metadevice is specified, the status of the master and log devices is also displayed. Trans metadevices have been replaced by UFS logging. See NOTES.

示例

示例 1 Output Showing Mirror with Two Submirrors

The following example shows the partial output of the `metastat` command after creating a mirror, `opt_mirror`, consisting of two submirrors, `opt_sub1` and `opt_sub2`.

```
# metastat opt_mirror
opt_mirror: Mirror
  Submirror 0: opt_sub1
    State: Okay
  Submirror 1: opt_sub2
    State: Resyncing
  Resync in progress: 15 % done
  Pass: 1
  Read option: roundrobin (default)
  Write option: parallel (default)
  Size: 2006130 blocks
  .
  .
  .
```

示例 2 Soft Partition on Mirror with Submirror

The following example shows the partial output of the `metastat` command after creating a soft partition, `d3`, on concat `d2`, which is built on a soft partition.

```
# metastat
d2: Concat/Stripe
  Size: 204800 blocks
  Stripe 0:
    Device          Start Block  Dbase State      Hot Spare
    d0              0          No   Okay
d0: Soft Partition
  Component: c0t3d0s0
  Status: Okay
  Size: 204800 blocks
    Extent          Start Block  Block count
    0              129          204800
d3: Soft Partition
  Component: d2
  Status: Okay
  Size: 202752 blocks
    Extent          Start Block  Block count
    0              129          202752
```

示例3 Trans Metadevice

The following example shows the output of the metastat command after creating a trans metadevice.

```
# metastat
d2: Concat/Stripe
  Size: 204800 blocks
  Stripe 0:
    Device          Start Block  Dbase State      Hot Spare
    d0              0          No    Okay
d0: Soft Partition
  Component: c0t3d0s0
  Status: Okay
  Size: 204800 blocks
    Extent          Start Block  Block count
    0              129          204800
d3: Soft Partition
  Component: d2
  Status: Okay
  Size: 202752 blocks
    Extent          Start Block  Block count
    0              129          202752
```

示例4 Multi-owner disk set

The following example shows the output of the metastat command with a multi-owner disk set and application-based mirror resynchronization option. Application-based resynchronization is set automatically if needed.

```
# metastat -s oban
oban/d100: Mirror
  Submirror 0: oban/d10
  State: Okay
  Submirror 1: oban/d11
  State: Okay
  Pass: 1
  Read option: roundrobin (default)
  Write option: parallel (default)
  Resync option: application based
  Owner: None
  Size: 1027216 blocks (501 MB)
oban/d10: Submirror of oban/d100
  State: Okay
  Size: 1027216 blocks (501 MB)
  Stripe 0:
```


示例 4 Multi-owner disk set (续)

```

Device      Start Block  Dbase    State Reloc Hot Spare
c1t3d0s0      0        No      Okay

oban/d11: Submirror of oban/d100
State: Okay
Size: 1027216 blocks (501 MB)
Stripe 0:
Device      Start Block  Dbase    State Reloc Hot Spare
c1t4d0s0      0        No      Okay
```

警告

metastat displays states as of the time the command is entered. It is unwise to use the output of the metastat -p command to create a [md.tab\(4\)](#) file for a number of reasons:

- The output of metastat -p might show hot spares being used.
- It might show mirrors with multiple submirrors. See [metainit\(1M\)](#) for instructions for creating multi-way mirrors using metainit and metattach.
- A slice may go into an error state after metastat -p is issued.

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见

[mdmonitord\(1M\)](#), [metaclear\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metahs\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metarecover\(1M\)](#), [metarename\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metassist\(1M\)](#), [metasync\(1M\)](#), [metattach\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [md.tab\(4\)](#), [attributes\(5\)](#), [md\(7D\)](#)

《Solaris Volume Manager Administration Guide》

附注

Trans metadevices have been replaced by UFS logging. Existing trans devices are *not* logging--they pass data directly through to the underlying device. See [mount_ufs\(1M\)](#) for more information about UFS logging.

引用名 metasync – handle metadvice resync during reboot

用法概要

```
/usr/sbin/metasync -h
/usr/sbin/metasync [-s setname] [buffer_size] metadvice
/usr/sbin/metasync [-s setname] -r [buffer_size]
/usr/sbin/metasync -p metadvice
/usr/sbin/metasync -c metadvice
```

描述

The `metasync` command starts a resync operation on the specified *metadvice*. All components that need to be resynced are resynced. If the system crashes during a RAID5 initialization, or during a RAID5 resync, either an initialization or resync restarts when the system reboots.

Applications are free to access a metadvice at the same time that it is being resynced by `metasync`. Also, `metasync` performs the copy operations from inside the kernel, which makes the utility more efficient.

Use the `-r` option in boot scripts to resync all possible submirrors.

选项

The following options are supported:

- `-c metadvice` Cancels the resync that is in progress on the specified metadvice. The resync will be stopped at its current point and can be resumed by running the “`metasync metadvice`” command. This option only applies to RAID1 volumes.
- `-h` Displays usage message.
- `-p metadvice` Regenerates parity information for RAID5 metadvicees.
- `-s setname` Specifies the name of the diskset on which `metasync` will work. Using the `-s` option will cause the command to perform its administrative function within the specified diskset. Without this option, the command will perform its function on local metadvicees.
- `-r` Specifies that the `metasync` command handle special resync requirements during a system reboot. `metasync -r` should only be invoked from the `svc:/system/mdmonitor` service. The `metasync` command only resyncs those metadvicees that need to be resynced. `metasync` schedules all the mirror resyncs according to their pass numbers.

To override the default `buffer_size` value used by the `svc:/system/mdmonitor` service, you can edit `/etc/system` to specify:

```
set md_mirror:md_resync_bufsz = 2048
```

so that resyncs occur as quickly as possible.

操作数 *buffer_size* Specifies the size (number of 512-byte disk blocks) of the internal copy buffer for the mirror resync. The size defaults to 1024 512-byte disk blocks. It can be no more than 2048 blocks. For best performance (quickest completion of the resync), 2048 blocks is the recommended size.

退出状态 The following exit values are returned:

0 Successful completion.

>0 An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm

另请参见 [mdmonitord\(1M\)](#), [metaclear\(1M\)](#), [metadb\(1M\)](#), [metadetach\(1M\)](#), [metahs\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metarecover\(1M\)](#), [metarename\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metassist\(1M\)](#), [metastat\(1M\)](#), [metattach\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [md.tab\(4\)](#), [attributes\(5\)](#), [md\(7D\)](#)

[《Solaris Volume Manager Administration Guide》](#)

附注 The metasync service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/system/mdmonitor`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名

metattach, metadetach – attach or detach a metadevice

用法概要

```
/usr/sbin/metattach [-h]
/usr/sbin/metattach [-s setname] mirror [metadevice]
/usr/sbin/metattach [-s setname] [-i interlace] concat/stripe component...
/usr/sbin/metattach [-s setname] RAID component...
/usr/sbin/metattach [-s setname] [-A alignment] softpart size | all
/usr/sbin/metadetach [-s setname] [-f] mirror submirror
/usr/sbin/metadetach [-s setname] [-f] trans
```

描述

metattach adds submirrors to a mirror, grows metadevices, or grows soft partitions. Growing metadevices can be done without interrupting service. To grow the size of a mirror or trans, the slices must be added to the submirrors or to the master devices.

Solaris Volume Manager supports storage devices and logical volumes greater than 1 terabyte (TB) when a system runs a 64-bit Solaris kernel. Support for large volumes is automatic. If a device greater than 1 TB is created, Solaris Volume Manager configures it appropriately and without user intervention.

If a system with large volumes is rebooted under a 32-bit Solaris kernel, the large volumes are visible through `metastat` output. Large volumes cannot be accessed, modified or deleted, and no new large volumes can be created. Any volumes or file systems on a large volume in this situation are also unavailable. If a system with large volumes is rebooted under a version of Solaris prior to the Solaris 9 4/03 release, Solaris Volume Manager does not start. You must remove all large volumes before Solaris Volume Manager runs under an earlier version of the Solaris Operating System.

Solaris Volume Manager supports one-to-four-way mirrors. You can only attach a metadevice to a mirror if there are three or fewer submirrors beneath the mirror. Once a new metadevice is attached to a mirror, `metattach` automatically starts a resync operation to the new submirror.

`metadetach` detaches submirrors from mirrors and logging devices from trans metadevices.

When a submirror is detached from a mirror, it is no longer part of the mirror, thus reads and writes to and from that metadevice by way of the mirror are no longer performed through the mirror. Detaching the only existing submirror is not allowed. Detaching a submirror that has slices reported as needing maintenance (by `metastat`) is not allowed unless the `-f` (force) flag is used.

`metadetach` also detaches the logging device from a trans. This step is necessary before you can clear the trans volume. Trans metadevices have been replaced by UFS logging. Existing trans devices are not logging. They pass data directly through to the underlying device. See [mount_ufs\(1M\)](#) for more information about UFS logging.

Detaching the logging device from a busy trans device is not allowed unless the `-f` (force) flag is used. Even so, the logging device is not actually detached until the trans is idle. The trans is in the *Detaching* state (`metastat`) until the logging device is detached.

选项

Root privileges are required for all of the following options except `-h`.

The following options are supported:

- `-A alignment` Set the value of the soft partition extent alignment. Use this option when it is important specify a starting offset for the soft partition. It preserves the data alignment between the metadvice address space and the address space of the underlying physical device.

For example, a hardware device that does checksumming should not have its I/O requests divided by Solaris Volume Manager. In this case, use a value from the hardware configuration as the value for the alignment. When using this option in conjunction with a software I/O load, the alignment value corresponds to the I/O load of the application. This prevents I/O from being divided unnecessarily and affecting performance.
- `-f` Force the detaching of metadvice devices that have components that need maintenance or are busy. You can use this option only when a mirror is in a maintenance state that can be fixed with `metareplace(1M)`. If the mirror is in a maintenance state that can only be fixed with `metasync(1M)` (as shown by the output of `metastat(1M)`), `metadvice -f` has no effect, because the mirrors must be resynchronized before one of them can be detached.
- `-h` Display a usage message.
- `-i interlace` Specify the interlace value for stripes, where `size` is a specified value followed by either `k` for kilobytes, `m` for megabytes, or `b` for blocks. The units can be either uppercase or lowercase. If `size` is not specified, the size defaults to the interlace size of the last stripe of the metadvice. When an interlace size change is made on a stripe, it is carried forward on all stripes that follow.
- `-s setname` Specify the name of the diskset on which the `metattach` command or the `metadvice` command works.. Using the `-s` option causes the command to perform its administrative function within the specified diskset. Without this option, the command performs its function on local metadvice devices.

操作数

The following operands are supported:

- `component` The logical name for the physical slice (partition) on a disk drive, such as `/dev/dsk/c0t0d0s2`, being added to the concatenation, stripe, concatenation of stripes, or RAID5 metadvice.
- `concat/stripe` The metadvice name of the concatenation, stripe, or concatenation of stripes.

<i>log</i>	The metadvice name of the logging device to be attached to the trans metadvice.
<i>metadvice</i>	The metadvice name to be attached to the mirror as a submirror. This metadvice must have been previously created by the <code>metainit</code> command.
<i>mirror</i>	The name of the mirror.
<i>RAID</i>	The metadvice name of the RAID5 metadvice.
<i>size</i> <i>all</i>	The amount of space to add to the soft partition in K or k for kilobytes, M or m for megabytes, G or g for gigabytes, T or t for terabytes, and B or b for blocks (sectors). All values represent powers of 2, and upper and lower case options are equivalent. Only integer values are permitted. The literal <code>all</code> specifies that the soft partition should grow to occupy all available space on the underlying volume.
<i>softpart</i>	The metadvice name of the existing soft partition.
<i>submirror</i>	The metadvice name of the submirror to be detached from the mirror.
<i>trans</i>	The metadvice name of the trans metadvice (not the master or logging device).

示例

示例 1 Concatenating a New Slice to a Metadvice

This example concatenates a single new slice to an existing metadvice, `Volume.1`. Afterwards, you would use the `growfs(1M)` command to expand the file system.

```
# metattach Volume.1 /dev/dsk/c0t1d0s2
```

示例 2 Detaching Logging Device from Trans Metadvice

This example detaches the logging device from a trans metadvice `d9`. Notice that you do not have to specify the logging device itself, as there can only be one.

```
# metadetach d9
```

示例 3 Expanding a RAID5 Metadvice

This example expands a RAID5 metadvice, `d45`, by attaching another slice.

```
# metattach d45 /dev/dsk/c3t0d0s2
```

When you add additional slices to a RAID5 metadvice, the additional space is devoted to data. No new parity blocks are allocated. The data on the added slices is, however, included in the overall parity calculations, so it is protected against single-device failure.

示例 4 Expanding a Soft Partition

The following example expands a soft partition, `d42`, attaching all space available on the underlying device.

示例 4 Expanding a Soft Partition (续)

```
# metattach d42 all
```

When you add additional space to a soft partition, the additional space is taken from any available space on the slice and might not be contiguous with the existing soft partition.

示例 5 Adding Space to Two-Way Mirror

This example adds space to a two-way mirror by adding a slice to each submirror. Afterwards, you would use the [growfs\(1M\)](#) command to expand the file system.

```
# metattach d9 /dev/dsk/c0t2d0s5
# metattach d10 /dev/dsk/c0t3d0s5
```

This example tells the mirror to grow to the size of the underlying devices

```
# metattach d11
```

This example increases the size of the UFS on the device so the space can be used.

```
# growfs -M /export /dev/md/rdisk/d11
```

示例 6 Detaching a Submirror from a Mirror

This example detaches a submirror, d2, from a mirror, d4.

```
# metadetach d4 d2
```

示例 7 Adding Four Slices to Metadevice

This example adds four slices to an existing metadevice, d9. Afterwards, you would use the [growfs\(1M\)](#) command to expand the file system.

```
# metattach d9 /dev/dsk/c0t1d0s2 /dev/dsk/c0t2d0s2 \\  
/dev/dsk/c0t3d0s2 /dev/dsk/c0t4d0s2
```

示例 8 Setting the Value of the Soft Partition Extent Alignment

This example shows how to set the alignment of the soft partition to 1mb when the soft partition is expanded.

```
# metattach -s red -A 2m d13 1m
```

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见 [mdmonitord\(1M\)](#), [metaclear\(1M\)](#), [metadb\(1M\)](#), [metahs\(1M\)](#), [metainit\(1M\)](#), [metaoffline\(1M\)](#), [metaonline\(1M\)](#), [metaparam\(1M\)](#), [metarecover\(1M\)](#), [metarename\(1M\)](#), [metareplace\(1M\)](#), [metaset\(1M\)](#), [metassist\(1M\)](#), [metastat\(1M\)](#), [metasync\(1M\)](#), [md.tab\(4\)](#), [md.cf\(4\)](#), [mddb.cf\(4\)](#), [md.tab\(4\)](#), [attributes\(5\)](#), [md\(7D\)](#)

《Solaris Volume Manager Administration Guide》

警告 This section provides information regarding warnings for devices greater than 1 TB and for multi-way mirrors.

Devices and Volumes Greater Than 1 TB Do not create large (>1 TB) volumes if you expect to run the Solaris Operating System with a 32-bit kernel or if you expect to use a version of the Solaris Operating System prior to Solaris 9 4/03.

Multi-Way Mirrors When a submirror is detached from its mirror, the data on the metadvice might not be the same as the data that existed on the mirror prior to running metadetach. In particular, if the -f option was needed, the metadvice and mirror probably do not contain the same data.

附注 Trans metadevices have been replaced by UFS logging. Existing trans devices are *not* logging. They pass data directly through to the underlying device. See [mount_ufs\(1M\)](#) for more information about UFS logging.

引用名

mib2mof – generate MOF file(s) from input SNMP MIB file(s)

用法概要

/usr/sadm/bin/mib2mof [-n] [-d *directory*] [-q] [-c] [-a] [-h] *files*

描述

The `mib2mof` utility reads input Management Information Base (MIB) files and produces one or more Managed Object Format (MOF) files. MOF files contain a Common Information Model (CIM) class declaration that represents the MIB for the Solaris Simple Network Management Protocol (SNMP) provider. The SNMP provider allows Web-Based Enterprise Management (WBEM) applications to access SNMP device information.

SNMP scalar variables map to properties in the CIM class. Qualifiers on each property convey the following MIB information for each scalar variable:

- syntax
- read/write access
- OID (Object IDentifier)
- description (optional)
- index (if the variable is within a group [sequence] that defines a row)

The syntax of an SNMP scalar variable is represented in a CIM class by the property's CIM datatype. All properties are marked with write access (true or false).

The following table shows how a Solaris SNMP datatype in a MIB maps to a Web-Based Enterprise Management (WBEM) CIM datatype and then to an SNMP datatype used by the WBEM SNMP API:

SNMP SMI Datatype	SNMP Ver.	CIM Datatype	SNMP API Object type
INTEGER	v1	sint32	SnmpInt
OCTET STRING	v1	string	SnmpString
OBJECT IDENTIFIER	v1	string	SnmpOid
IpAddress	v1	string	SnmpIpAddress
Counter	v1	uint32	SnmpCounter
Gauge	v1	uint32	SnmpGauge
TimeTicks	v1	uint32	SnmpTimeTicks
Opaque	v1	sint8[]	SnmpOpaque
DisplayString - see OCTET STRING	v1		
NetworkAddress - see IpAddress	v1		
Counter32 - see Counter	v2		
Counter64	v2	uint64	SnmpCounter64
Integer32	v2	sint32	SnmpInt
Gauge32 - see Gauge	v2		
Unsigned32	v2	uint32	SnmpGauge
TruthValue	v2	sint32	SnmpInt
BITS - see OCTET STRING	v2		

The `mib2mof` utility includes its required `Solaris_SNMPmib_core.txt` file (containing core MIB definitions), installed in `/usr/sadm/mof`. The `mib2mof` utility looks first for `mib` core file in local directory. If this file is not found in the local directory, `mib2mof` looks in `/usr/sadm/mof`.

A MOF file is generated for each SNMP group and table row sequence (that is, the columns in one row) found in the supplied MIBs. (This does not include the core MIB definitions contained in the `Solaris_SNMPmib_core.txt` file.)

There is no MOF file or property for an SNMP table - all table access is through the rows and columns of the table, and the SNMP variable for the table is marked as inaccessible in the MIB.

The MOF file created contains a CIM class that represents an SNMP group or row and a CIM class to represent a CIM association. The output file name (and CIM class) is of the format `<SNMP_><MIB name><Group name>.mof`.

选项

The following options are supported:

- a Generate MOF files for all of the input MIB files. If -a is not given, a MOF file is generated only for the last file of the input list.
- c Do not use the default `Solaris_SNMPmib_core.txt` definitions file shipped with the Solaris SNMP Provider for WBEM. If this option is specified, you must specify another `MIB_CORE` definitions file as one of the input files.
- d *directory* Generate output MOF files in the specified directory.
- h Show how to invoke `mib2mof` and list its arguments.
- n Parse the input MIB files without generating any output.
- q Include the DESCRIPTION clause of SNMP OBJECT-TYPE as a qualifier in the generated MOF file.

操作数

The following operands are supported:

files List of SNMP MIB files to be converted.

退出状态

The `mib2mof` utility terminates with exit status 0.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWwbco

另请参见

[init.wbem\(1M\)](#), [mofcomp\(1M\)](#), [wbemadmin\(1M\)](#), [attributes\(5\)](#)

引用名	mibiisa – Sun SNMP Agent
用法概要	mibiisa [-ar] [-c <i>config-dir</i>] [-d <i>debug-level</i>] [-p <i>port</i>] [-t <i>cache-timer</i>]
描述	The mibiisa utility is an RFC 1157-compliant SNMP agent. It supports MIB-II as defined in RFC 1213, with Sun extensions under Sun's enterprise number. The MIB (Management Information Base) is both readable and writable. The mibiisa utility supports all SNMP protocol operations including GET-REQUEST, GETNEXT-REQUEST, SET-REQUEST, GET-REPLY, and TRAP. The SMA (Systems Management Agent) is the default SNMP agent in Solaris. MIB-II subagent mibiisa does not run by default. To enable mibiisa, rename the configuration file from /etc/snmp/conf/mibiisa.rsrc - to /etc/snmp/conf/mibiisa.rsrc. SMA has the capability to handle any MIB-II requests. The mibiisa utility supports the coldStart, linkUp, linkDown, and authentication traps. The authentication trap may be disabled by a command-line switch, which itself may be overridden by a management station writing to a MIB variable in the standard SNMP MIB group. The mibiisa utility supports four distinct views of the MIB. The view used for any request is determined by the community string contained in that request. To enhance security, mibiisa supports an option to block all writes to the MIB. You can also limit the set of management stations from which the agent will accept requests in the configuration file used when starting the mibiisa. See the 安全 section for more information. Unless overridden, mibiisa uses UDP port 161, the standard SNMP port. The mibiisa utility issues traps through the same port on which it receives SNMP requests. The mibiisa utility must run with super-user privileges and is typically started at system startup via /etc/rc3.d. mibiisa may not be started using inetd(1M). When started, mibiisa detaches itself from the keyboard, disables all signals except SIGKILL, SIGILL, SIGUSR1, and SIGUSR2, and places itself in the background.
选项	The following options are supported by mibiisa: -a Disable the generation of authentication traps. However, an SNMP manager may write a value into snmpEnableAuthenTraps to enable or disable authentication traps. -c <i>config-dir</i> Specify a directory where it expects snmpd.conf file, on startup. The default directory is /etc/snmp/conf. -d <i>debug-level</i> Debug. A value of 0 disables all debug and is the default. Levels 1 through 3 represent increasing levels of debug output. When mibiisa receives the signal SIGUSR1, it resets the debug-level to 0. When mibiisa receives the signal SIGUSR2, it increments the debug-level by one.

Debug output is sent to the standard output in effect at the time *mibiisa* is started. No matter what debug level is in effect, certain significant events are logged in the system log.

- p *port* Define an alternative UDP port on which *mibiisa* listens for incoming requests. The default is UDP port 161.
- r Place the MIB into read-only mode.
- t *cache-timer* By default, information fetched from the kernel is considered to be valid for 45 seconds from the time it is retrieved. This cache lifetime may be altered with this parameter. You cannot set *cache-timer* to any value less than 1.

Configuration File The *snmpd.conf* file is used for configuration information. Each entry in the file consists of a keyword followed by a parameter string. The keyword must begin in the first position. Parameters are separated from the keyword and from one another by white space. Case in keywords is ignored. Each entry must be contained on a single line. All text following (and including) a pound sign (#) is ignored. Keywords currently supported are:

- sysdescr* The value to be used to answer queries for *sysDescr*.
- syscontact* The value to be used to answer queries for *sysContact*.
- syslocation* The value to be used to answer queries for *sysLocation*.
- trap* The parameter names one or more hosts to receive traps. Only five hosts may be listed.
- system-group-read-community* The community name to get read access to the system group and Sun's extended system group.
- system-group-write-community* The community name to get write access to the system group and Sun's extended system group.
- read-community* The community name to get read access to the entire MIB.
- write-community* The community name to get write access to the entire MIB (implies read access).
- trap-community* The community name to be used in traps.
- kernel-file* The name of the file to use for kernel symbols.
- managers* The names of hosts that may send SNMP queries. Only five hosts may be listed on any one line. This keyword may be repeated for a total of 32 hosts.
- newdevice* The additional devices which are not built in *SNMPD*. The format is as follows: *newdevice* type *speed* *name* where *newdevice* is the keyword, type is an integer

which has to match your schema file, *speed* is the new device's speed, and *name* is this new device's name.

An example `snmpd.conf` file is shown below:

```
sysdescr      Sun SNMP Agent, Sun Fire 4800, Company
               Property Number 123456
syscontact    Cliff Claven
sysLocation   Room 1515, building 1
#
system-group-read-community    public
system-group-write-community   private
#
read-community  all_public
write-community all_private
#
trap            localhost
trap-community  SNMP-trap
#
#kernel-file    /vmunix
#
managers        lvs golden
managers        swap
```

Installation

The `mibiisa` utility and its configuration file, `snmpd.conf`, may be placed in any directory. However for Solaris 2.4 and subsequent releases, use `/usr/lib/snmp` for `mibiisa` itself and `/etc/snmp/conf` for the configuration file. You can modify the configuration file as appropriate. If you make any changes to `snmpd.conf` file keyword values, you must kill and restart `mibiisa` for the changes to take effect.

Your `/etc/services` file (or NIS equivalent) should contain the following entries:

snmp	161/udp		# Simple Network Mgmt Protocol
snmp-trap	162/udp	snmptrap	# SNMP trap (event) messages

The following is an example for Solaris 2.x and releases compatible with Solaris 2.x, such as Solaris 9:

```
#
# Start the SNMP agent
#
if [ -f /etc/snmp/conf/snmpd.conf -a -x
    /usr/lib/snmp/mibiisa ];
then
```

```
/opt/SUNWconn/snm/agents/snmpd
echo 'Starting SNMP-agent.'
```

安全

SNMP, as presently defined, offers relatively little security. The `mibiisa` utility accepts requests from other machines, which can have the effect of disabling the network capabilities of your computer. To limit the risk, the configuration file lets you specify a list of up to 32 manager stations from which `mibiisa` will accept requests. If you do not specify any such manager stations, `mibiisa` accepts requests from anywhere.

The `mibiisa` utility also allows you to mark the MIB as “read-only” by using the `-r` option.

`mibiisa` supports four different community strings. These strings, however, are visible in the configuration file and within the SNMP packets as they flow on the network.

The configuration file should be owned by, and readable only by super-user. In other words the mode should be:

```
-rw----- 1 root          2090 Oct 17 15:04 /etc/snmp/conf/snmpd.conf
```

Managers can be restricted based on the community strings. This can be configured by creating an optional secondary configuration file `/etc/snmp/conf/mibiisa.acl`. To enable such a restriction, add the security line in the `/etc/snmp/conf/mibiisa.rsrc` file.

An example `mibiisa.acl` file is as follows:

```
acl = {
    {
        communities = public
        access = read-only
        managers = xyz
    }
    {
        communities = private
        access = read-write
        managers = abc,pqrs
    }
}
```

An example `mibiisa.rsrc` file is as follows:

```
resource =
{
    {
        registration_file = "/etc/snmp/conf/mibiisa.reg"
        security = "/etc/snmp/conf/mibiisa.acl"
        policy = "spawn"
        type = "legacy"
        command = "/usr/lib/snmp/mibiisa -r -p $PORT"
    }
}
```

Mib This section discusses some of the differences between the mibiisa MIB and the standard MIB-II (as defined in RFC 1213).

The following variables are read-only in the mibiisa MIB:

sysName
atIfIndex
ipDefaultTTL

These variables are read-write in the standard MIB-II.

The mibiisa MIB Address Translation tables support limited write access: only atPhysAddress may be written, either to change the physical address of an existing entry or to delete an entire ARP table entry.

The mibiisa MIB IP Net to Media table supports limited write access: only ipNetToMediaPhysAddress and ipNetToMediaType may be written, either to change the physical address of an existing entry or to delete an entire ARP table entry.

The following variables are read-write in the mibiisa MIB; however, these variables have fixed values. Any new values “set” to them are accepted, but have no effect:

ipRouteIfIndex
ipRouteMetric1
ipRouteMetric2
ipRouteMetric3
ipRouteMetric4
ipRouteType
ipRouteAge
ipRouteMask
ipRouteMetric5

The following mibiisa MIB variable reflects the actual state of the related table entry. “Sets” are accepted but have no effect:

tcpConnState

The following mibiisa MIB variables are readable, but return a fixed value:

icmpInDestUnreachs	Returns 1
icmpInTimeExcds	Returns 1
icmpInParmProbs	Returns 1
icmpInSrcQuenchs	Returns 1
icmpInRedirects	Returns 1
icmpInEchos	Returns 1

icmpInEchoReps	Returns 1
icmpInTimestamps	Returns 1
icmpInTimestampReps	Returns 1
icmpInAddrMasks	Returns 1
icmpInAddrMaskReps	Returns 1
icmpOutDestUnreachs	Returns 1
icmpOutTimeExcds	Returns 1
icmpOutParmProbs	Returns 1
icmpOutSrcQuenchs	Returns 1
icmpOutRedirects	Returns 1
icmpOutEchos	Returns 1
icmpOutEchoReps	Returns 1
icmpOutTimestamps	Returns 1
icmpOutTimestampReps	Returns 1
icmpOutAddrMasks	Returns 1
icmpOutAddrMaskReps	Returns 1
ifInUnknownProtos	Returns 0
ipAdEntBcastAddr	Returns 1
ipAdEntReasmMaxSiz	Returns 65535
ipRouteMetric1	Returns -1
ipRouteMetric2	Returns -1
ipRouteMetric3	Returns -1
ipRouteMetric4	Returns -1
ipRouteAge	Returns 0
ipRouteMetric5	Returns -1
ipNetToMediaType	Returns (3) dynamic
ipRoutingDiscards	Returns 0

The following variables return a fixed value of 0 for drivers not conforming to the GLD framework (see [gld\(7D\)](#)), including the old LAN drivers on SPARC machines:

ifInOctets	Returns 0
ifInNUcastPkts	Returns 0
ifInDiscards	Returns 0
ifOutOctets	Returns 0
ifOutNUcastPkts	Returns 0
ifOutDiscards	Returns 0

Schema Attributes The following describes the attributes in the group and table definitions in the /var/snmp/mib/sun.mib file.

system The system group reports statistics about a particular system (for example, a workstation or a printer).

- sysDescr – A textual description of the entity. This value should include the full name and version identification of the system's hardware type, software operating-system, and networking software. This value must only contain printable ASCII characters. (string[255])
- sysObjectID – The vendor's authoritative identification of the network management subsystem contained in the entity. This value is allocated within the SMI enterprises subtree (1.3.6.1.4.1) and provides an easy and unambiguous means for determining what type of equipment is being managed. For example, if vendor “Flintstones, Inc.” was assigned the subtree 1.3.6.1.4.1.4242, it could assign the identifier 1.3.6.1.4.1.4242.1.1 to its “Fred Router.” (objectid)
- sysUpTime – Time (in hundredths of a second) since the network management portion of the system was last reinitialized. (timeticks)
- sysContact – The textual identification of the contact person for this managed node, together with information on how to contact this person. (string[255])
- sysName – An administratively-assigned name for this managed node. By convention, this is the node's fully-qualified domain name. (string[255])
- sysLocation – The physical location of this node (for example, “telephone closet, 3rd floor” (string[255]))
- sysServices – A value indicating the set of services that this entity primarily offers. (int) The value is a sum. This sum initially takes the value zero. Then, for each layer L in the range 1 through 7 for which this node performs transactions, 2 raised to (L - 1) is added to the sum. For example, a node that performs primarily routing functions would have a value of 4 (2*(3-1)). In contrast, a node that is a host offering application services would have a value of 72 (2*(4-1) + 2*(7-1)). Note that in the context of the Internet suite of protocols, values should be calculated accordingly:

Layer	Functionality
1	physical (such as repeaters)
2	datalink/subnetwork (such as bridges)
3	internet (such as IP gateways)
4	end-to-end (such as IP hosts)
7	applications (such as mail relays)

For systems including OSI protocols, Layers 5 and 6 may also be counted.

interfaces

The `interfaces` group reports the number of interfaces handled by the agent.

`ifNumber` – The number of network interfaces, regardless of their current state, present on this system. (int)

ifTable

The `ifTable` is a table of interface entries. The number of entries is given by the value of `ifNumber`.

`ifIndex` – A unique value for each interface. Its value ranges between 1 and the value of `ifNumber`. The value for each interface must remain constant at least from one reinitialization of the entity's network management system to the next reinitialization. (int)

`ifDescr` – A textual string containing information about the interface. This string should include the name of the manufacturer, the product name, and the version of the hardware interface. (string[255])

`ifType` – The type of interface, distinguished according to the physical/link protocol(s) immediately below the network layer in the protocol stack. (enum)

`ifMtu` – The size of the largest datagram that can be sent/received on the interface, specified in octets. For interfaces used for transmitting network datagrams, this is the size of the largest network datagram that can be sent on the interface. (int)

`ifSpeed` – An estimate of the interface's current bandwidth in bits-per-second. For interfaces that do not vary in bandwidth, or for those where no accurate estimation can be made, this object should contain the nominal bandwidth. (gauge)

`ifPhysAddress` – The interface's address at the protocol layer immediately below the network layer in the protocol stack. For interfaces without such an address (for example, a serial line), this object should contain an octet string of zero length. (octet[128])

`ifAdminStatus` – The desired state of the interface. The `testing(3)` state indicates that no operational packets can be passed. (enum)

`ifOperStatus` – The current operational state of the interface. The `testing(3)` state indicates that no operational packets can be passed. (enum)

ifLastChange – The value of sysUpTime at the time the interface entered its current operational state. If the current state was entered prior to the last reinitialization of the local network management subsystem, then this object contains a zero value. (timeticks)

ifInOctets – The total number of octets received on the interface, including framing characters. (counter) Returns a fixed value of 0.

ifInUcastPkts – The number of subnetwork-unicast packets delivered to a higher-layer protocol. (counter)

ifInNUcastPkts – The number of non-unicast (that is, subnetwork- broadcast or subnetwork-multicast) packets delivered to a higher-layer protocol. (counter) Returns a fixed value of 0.

ifInDiscards – The number of inbound packets chosen to be discarded, even though no errors had been detected to prevent their being deliverable to a higher-layer protocol. One possible reason for discarding such a packet could be to free up buffer space. (counter) Returns a fixed value of 0.

ifInErrors – The number of inbound packets that contained errors preventing them from being deliverable to a higher-layer protocol. (counter)

ifInUnknownProtos – The number of packets received via the interface that were discarded because of an unknown or unsupported protocol. (counter) Returns a fixed value of 0.

ifOutOctets – The total number of octets transmitted out of the interface, including framing characters. (counter) Returns a fixed value of 0.

ifOutUcastPkts – The total number of packets that higher-level protocols requested be transmitted to a subnetwork-unicast address, including those that were discarded or not sent. (counter)

ifOutNUcastPkts – The total number of packets that higher-level protocols requested be transmitted to a non- unicast (that is, a subnetwork-broadcast or subnetwork-multicast) address, including those that were discarded or not sent. (counter) Returns a fixed value of 0.

ifOutDiscards – The number of outbound packets that were chosen to be discarded even though no errors had been detected to prevent their being transmitted. One possible reason for discarding such a packet could be to free up buffer space. (counter) Returns a fixed value of 0.

ifOutErrors – The number of outbound packets that could not be transmitted because of errors. (counter)

ifOutQLen – The length of the output packet queue (in packets). (gauge)

ifSpecific – A reference to MIB definitions specific to the particular media being used to realize the interface. For example, if the interface is realized by an Ethernet, then the value of

this object refers to a document defining objects specific to Ethernet. If this information is not present, its value should be set to the OBJECT IDENTIFIER { 0 0 }, which is a syntactically valid object identifier. Any conformant implementation of ASN.1 and BER must be able to generate and recognize this value. (objectid)

atTable

atTable Address Translation tables contain the NetworkAddress to physical address equivalences. Some interfaces do not use translation tables for determining address equivalences (for example, DDN-X.25 has an algorithmic method). If all interfaces are of this type, then the Address Translation table is empty, that is, has zero entries.

atIfIndex – The interface on which this entry's equivalence is effective. The interface identified by a particular value of this index is the same interface as identified by the same value of ifIndex. (int)

atPhysAddress – The media-dependent physical address. (octet[128]) Setting this object to a null string (one of zero length) has the effect of invalidating the corresponding entry in the atTable object. That is, it effectively dissociates the interface identified with said entry from the mapping identified with said entry. It is an implementation-specific matter as to whether the agent removes an invalidated entry from the table. Accordingly, management stations must be prepared to receive tabular information from agents that corresponds to entries not currently in use. Proper interpretation of such entries requires examination of the relevant atPhysAddress object.

atNetAddress – The NetworkAddress (that is, the IP address) corresponding to the media-dependent physical address. (netaddress)

ip

The ip group reports statistics about the Internet Protocol (IP) group.

ipForwarding – The indication of whether this entity is acting as an IP gateway in respect to the forwarding of datagrams received by, but not addressed to, this entity. IP gateways forward datagrams. IP hosts do not— except those source-routed via the host. (enum)

Note that for some managed nodes, this object may take on only a subset of the values possible. Accordingly, it is appropriate for an agent to return a “badValue” response if a management station attempts to change this object to an inappropriate value.

ipDefaultTTL – The default value inserted into the Time-To-Live field of the IP header of datagrams originated at this entity, whenever a TTL value is not supplied by the transport layer protocol. (int)

ipInReceives – The total number of input datagrams received from interfaces, including those received in error. (counter)

ipInHdrErrors – The number of input datagrams discarded due to errors in their IP headers, including bad checksums, version number mismatch, other format errors, time-to-live exceeded, errors discovered in processing their IP options, and so on. (counter)

ipInAddrErrors – The number of input datagrams discarded because the IP address in their IP header's destination field was not a valid address to be received at this entity. This count includes invalid addresses (for example, 0.0.0.0) and addresses of unsupported Classes (for example, Class E). For entities that are not IP Gateways and therefore do not forward datagrams, this counter includes datagrams discarded because the destination address was not a local address. (counter)

ipForwDatagrams – The number of input datagrams for which this entity was not their final IP destination, as a result of which an attempt was made to find a route to forward them to that final destination. In entities that do not act as IP Gateways, this counter will include only those packets that were Source-Routed via this entity, and the Source-Route option processing was successful. (counter)

ipInUnknownProtos – The number of locally-addressed datagrams received successfully but discarded because of an unknown or unsupported protocol. (counter)

ipInDiscards – The number of input IP datagrams for which no problems were encountered to prevent their continued processing, but which were discarded, for example, for lack of buffer space. Note that this counter does not include any datagrams discarded while awaiting reassembly. (counter)

ipInDelivers – The total number of input datagrams successfully delivered to IP user-protocols (including ICMP). (counter)

ipOutRequests – The total number of IP datagrams that local IP user-protocols (including ICMP) supplied to IP in requests for transmission. Note that this counter does not include any datagrams counted in ipForwDatagrams. (counter)

ipOutDiscards – The number of output IP datagrams for which no problem was encountered to prevent their transmission to their destination, but which were discarded (for example, for lack of buffer space). Note that this counter would include datagrams counted in ipForwDatagrams if any such packets met this (discretionary) discard criterion. (counter)

ipOutNoRoutes – The number of IP datagrams discarded because no route could be found to transmit them to their destination. Note that this counter includes any packets counted in ipForwDatagrams which meet this “no-route” criterion. Note that this includes any datagrams that a host cannot route because all its default gateways are down. (counter)

ipReasmTimeout – The maximum number of seconds that received fragments are held while they are awaiting reassembly at this entity. (int)

ipReasmReqds – The number of IP fragments received that needed to be reassembled at this entity. (counter)

ipReasmOKs – The number of IP datagrams successfully reassembled. (counter)

ipReasmFails – The number of failures detected by the IP reassembly algorithm, for whatever reason: timed out, errors, and the like. Note that this is not necessarily a count of discarded IP fragments since some algorithms (notably the algorithm in RFC 815) can lose track of the number of fragments by combining them as they are received. (counter)

ipFragOKs – The number of IP datagrams that have been successfully fragmented at this entity. (counter)

ipFragFails – The number of IP datagrams that have been discarded because they needed to be fragmented at this entity but could not be, for example, because their “Don't Fragment” flag was set. (counter)

ipFragCreates – The number of IP datagram fragments that have been generated as a result of fragmentation at this entity. (counter)

ipRoutingDiscards – The number of routing entries that were chosen to be discarded even though they were valid. One possible reason for discarding such an entry could be to free-up buffer space for other routing entries. (counter) Returns a fixed value of 0.

ipAddrTable **ipAddrTable** is a table of addressing information relevant to this entity's IP addresses.

ipAdEntAddr – The IP address to which this entry's addressing information pertains. (netaddress)

ipAdEntIfIndex – The index value that uniquely identifies the interface to which this entry is applicable. The interface identified by a particular value of this index is the same interface as identified by the same value of ifIndex. (int)

ipAdEntNetMask – The subnet mask associated with the IP address of this entry. The value of the mask is an IP address with all the network bits set to 1, and all the hosts bits set to 0. (netaddress)

ipAdEntBcastAddr – The value of the least-significant bit in the IP broadcast address used for sending datagrams on the (logical) interface associated with the IP address of this entry. For example, when the Internet standard all-ones broadcast address is used, the value will be 1. This value applies to both the subnet and network broadcasts addresses used by the entity on this (logical) interface. (int) Returns a fixed value of 1.

ipAdEntReasmMaxSize – The size of the largest IP datagram that this entity can reassemble from incoming IP fragmented datagrams received on this interface. (int) Returns a fixed value of 65535.

ipRouteTable The **ipRouteTable** is this entity's IP Routing table.

ipRouteDest – The destination IP address of this route. An entry with a value of 0.0.0.0 is considered a default route. Multiple routes to a single destination can appear in the table, but access to such multiple entries is dependent on the table- access mechanisms defined by the network management protocol in use. (netaddress)

ipRouteIfIndex – The index value that uniquely identifies the local interface through which the next hop of this route should be reached. The interface identified by a particular value of this index is the same interface as identified by the same value of **ifIndex**. (int)

ipRouteMetric1 – The primary routing metric for this route. The semantics of this metric are determined by the routing-protocol specified in the route's **ipRouteProto** value. If this metric is not used, its value should be set to -1. (int) Returns a fixed value of -1.

ipRouteMetric2 – An alternate routing metric for this route. The semantics of this metric are determined by the routing-protocol specified in the route's **ipRouteProto** value. If this metric is not used, its value should be set to -1. (int) Returns a fixed value of -1.

ipRouteMetric3 – An alternate routing metric for this route. The semantics of this metric are determined by the routing-protocol specified in the route's **ipRouteProto** value. If this metric is not used, its value should be set to -1. (int) Returns a fixed value of -1.

ipRouteMetric4 – An alternate routing metric for this route. The semantics of this metric are determined by the routing-protocol specified in the route's **ipRouteProto** value. If this metric is not used, its value should be set to -1. (int) Returns a fixed value of -1.

ipRouteNextHop – The IP address of the next hop of this route. (In the case of a route bound to an interface that is realized via a broadcast media, the value of this field is the agent's IP address on that interface.) (netaddress)

ipRouteType – The type of route. Note that the values **direct** (3) and **indirect** (4) refer to the notion of direct and indirect routing in the IP architecture. (enum)

Setting this object to the value **invalid** (2) has the effect of invalidating the corresponding entry in the **ipRouteTable** object. That is, it effectively dissociates the destination identified with said entry from the route identified with said entry. It is an implementation-specific matter as to whether the agent removes an invalidated entry from the table. Accordingly, management stations must be prepared to receive tabular information from agents that corresponds to entries not currently in use. Proper interpretation of such entries requires examination of the relevant **ipRouteType** object.

ipRouteProto – The routing mechanism through which this route was learned. Inclusion of values for gateway routing protocols is not intended to imply that hosts should support those protocols. (enum)

ipRouteAge – The number of seconds since this route was last updated or otherwise determined to be correct. Note that no semantics of “too old” can be implied except through knowledge of the routing protocol by which the route was learned. (int) Returns a fixed value of 0.

ipRouteMask – Indicate the mask to be logical-ANDed with the destination address before being compared to the value in the **ipRouteDest** field. For those systems that do not support

arbitrary subnet masks, an agent constructs the value of the ipRouteMask by determining whether the value of the correspondent ipRouteDest field belongs to a class-A, B, or C network, and then using one of:

Mask	Network
255.0.0.0	class-A
255.255.0.0	class-B
255.255.255.0	class-C

If the value of the ipRouteDest is 0.0.0.0 (a default route), then the mask value is also 0.0.0.0. It should be noted that all IP routing subsystems implicitly use this mechanism. (netaddress)

ipRouteMetric5 – An alternate routing metric for this route. The semantics of this metric are determined by the routing-protocol specified in the route's ipRouteProto value. If this metric is not used, its value should be set to -1. (int) Returns a fixed value of -1.

ipRouteInfo – A reference to MIB definitions specific to the particular routing protocol responsible for this route, as determined by the value specified in the route's ipRouteProto value. If this information is not present, its value should be set to the OBJECT IDENTIFIER { 0 0 }, which is a syntactically valid object identifier. Any conformant implementation of ASN.1 and BER must be able to generate and recognize this value. (objectid)

ipNetToMediaTable The ipNetToMediaTable is the IP Address Translation table used for mapping from IP addresses to physical addresses.

ipNetToMediaIfIndex – The interface on which this entry's equivalence is effective. The interface identified by a particular value of this index is the same interface as identified by the same value of ifIndex. (int)

ipNetToMediaPhysAddress – The media-dependent physical address. (octet[128])

ipNetToMediaNetAddress – The IPAddress corresponding to the media- dependent physical address. (netaddress)

ipNetToMediaType – The type of mapping. (enum) Returns a fixed value of (3)dynamic. Setting this object to the value invalid(2) has the effect of invalidating the corresponding entry in the ipNetToMediaTable. That is, it effectively dissociates the interface identified with said entry from the mapping identified with said entry. It is an implementation-specific matter as to whether the agent removes an invalidated entry from the table. Accordingly, management stations must be prepared to receive tabular information from agents that corresponds to entries not currently in use. Proper interpretation of such entries requires examination of the relevant ipNetToMediaType object.

icmp

The icmp group reports statistics about the ICMP group.

`icmpInMsgs` – The total number of ICMP messages that the entity received. Note that this counter includes all those counted by `icmpInErrors`. (counter)

`icmpInErrors` – The number of ICMP messages that the entity received but determined as having ICMP-specific errors (bad ICMP checksums, bad length, and the like.). (counter)

`icmpInDestUnreachs` – The number of ICMP Destination Unreachable messages received. (counter)

`icmpInTimeExcds` – The number of ICMP Time Exceeded messages received. (counter)

`icmpInParmProbs` – The number of ICMP Parameter Problem messages received. (counter)

`icmpInSrcQuenchs` – The number of ICMP Source Quench messages received. (counter)

`icmpInRedirects` – The number of ICMP Redirect messages received. (counter)

`icmpInEchos` – The number of ICMP Echo (request) messages received. (counter)

`icmpInEchoReps` – The number of ICMP Echo Reply messages received. (counter)

`icmpInTimestamps` – The number of ICMP Timestamp (request) messages received. (counter)

`icmpInTimestampReps` – The number of ICMP Timestamp Reply messages received. (counter)

`icmpInAddrMasks` – The number of ICMP Address Mask Request messages received. (counter)

`icmpInAddrMaskReps` – The number of ICMP Address Mask Reply messages received. (counter)

`icmpOutMsgs` – The total number of ICMP messages that this entity attempted to send. Note that this counter includes all those counted by `icmpOutErrors`. (counter)

`icmpOutErrors` – The number of ICMP messages that this entity did not send due to problems discovered within ICMP, such as a lack of buffers. This value should not include errors discovered outside the ICMP layer, such as the inability of IP to route the resultant datagram. In some implementations there may be no types of errors that contribute to this counter's value. (counter)

`icmpOutDestUnreachs` – The number of ICMP Destination Unreachable messages sent. (counter)

`icmpOutTimeExcds` – The number of ICMP Time Exceeded messages sent. (counter)

`icmpOutParmProbs` – The number of ICMP Parameter Problem messages sent. (counter)

`icmpOutSrcQuenchs` – The number of ICMP Source Quench messages sent. (counter)

`icmpOutRedirects` – The number of ICMP Redirect messages sent. For a host, this object will always be zero, since hosts do not send redirects. (counter)

`icmpOutEchos` – The number of ICMP Echo (request) messages sent. (counter)

`icmpOutEchoReps` – The number of ICMP Echo Reply messages sent. (counter)

`icmpOutTimestamps` – The number of ICMP Timestamp (request) messages sent. (counter)

`icmpOutTimestampReps` – The number of ICMP Timestamp Reply messages sent. (counter)

`icmpOutAddrMasks` – The number of ICMP Address Mask Request messages sent. (counter)

`icmpOutAddrMaskReps` – The number of ICMP Address Mask Reply messages sent. (counter)

tcp

The `tcp` group reports statistics about the TCP group.

`tcpRtoAlgorithm` – The algorithm used to determine the timeout value used for retransmitting unacknowledged octets. (enum)

`tcpRtoMin` – The minimum value permitted by a TCP implementation for the retransmission timeout, measured in milliseconds. More refined semantics for objects of this type depend upon the algorithm used to determine the retransmission timeout. In particular, when the timeout algorithm is `rsre(3)`, an object of this type has the semantics of the `LBOUND` quantity described in RFC 793. (int)

`tcpRtoMax` – The maximum value permitted by a TCP implementation for the retransmission timeout, measured in milliseconds. More refined semantics for objects of this type depend upon the algorithm used to determine the retransmission timeout. In particular, when the timeout algorithm is `rsre(3)`, an object of this type has the semantics of the `UBOUND` quantity described in RFC 793. (int)

`tcpMaxConn` – The limit on the total number of TCP connections that the entity can support. In entities where the maximum number of connections is dynamic, this object should contain the value `-1`. (int)

`tcpActiveOpens` – The number of times that TCP connections have made a direct transition to the `SYN-SENT` state from the `CLOSED` state. (counter)

`tcpPassiveOpens` – The number of times that TCP connections have made a direct transition to the `SYN-RCVD` state from the `LISTEN` state. (counter)

`tcpAttemptFails` – The number of times that TCP connections have made a direct transition to the `CLOSED` state from either the `SYN-SENT` state or the `SYN-RCVD` state, plus the number of times TCP connections have made a direct transition to the `LISTEN` state from the `SYN-RCVD` state. (counter)

`tcpEstabResets` – The number of times TCP connections have made a direct transition to the CLOSED state from either the ESTABLISHED state or the CLOSE-WAIT state. (counter)

`tcpCurrEstab` – The number of TCP connections for which the current state is either ESTABLISHED or CLOSE-WAIT. (gauge)

`tcpInSegs` – The total number of segments received, including those received in error. This count includes segments received on currently established connections. (counter)

`tcpOutSegs` – The total number of segments sent, including those on current connections but excluding those containing only retransmitted octets. (counter)

`tcpRetransSegs` – The total number of segments retransmitted - that is, the number of TCP segments transmitted containing one or more previously transmitted octets. (counter)

`tcpInErrs` – The total number of segments received in error (for example, bad TCP checksums). (counter)

`tcpOutRsts` – The number of TCP segments sent containing the RST flag. (counter)

`tcpConnTable`

The `tcpConnTable` is a table containing TCP connection-specific information.

`tcpConnState` – The state of this TCP connection. (enum)

The only value that may be set by a management station is `deleteTCB(12)`. Accordingly, it is appropriate for an agent to return a “badValue” response if a management station attempts to set this object to any other value.

If a management station sets this object to the value `deleteTCB(12)`, then this has the effect of deleting the TCB (as defined in RFC 793) of the corresponding connection on the managed node. This results in immediate termination of the connection.

As an implementation-specific option, an RST segment may be sent from the managed node to the other TCP endpoint. (Note, however, that RST segments are not sent reliably.)

`tcpConnLocalAddress` – The local IP address for this TCP connection. For a connection in the listen state that is willing to accept connections for any IP interface associated with the node, the value 0.0.0.0 is used. (netaddress)

`tcpConnLocalPort` – The local port number for this TCP connection. (int)

`tcpConnRemAddress` – The remote IP address for this TCP connection. (netaddress)

`tcpConnRemPort` – The remote port number for this TCP connection. (int)

`udp`

The `udp` group reports statistics about the UDP group.

`udpInDatagrams` – The total number of UDP datagrams delivered to UDP users. (counter)
Returns a fixed value of 0.

`udpNoPorts` – The total number of received UDP datagrams for which there was no application at the destination port. (counter) Returns a fixed value of 0.

`udpInErrors` – The number of received UDP datagrams that could not be delivered for reasons other than the lack of an application at the destination port. (counter)

`udpOutDatagrams` – The total number of UDP datagrams sent from this entity. (counter) Returns a fixed value of 0.

`udpTable` The `udpTable` is a table containing UDP listener information.

`udpLocalAddress` – The local IP address for this UDP listener. For a UDP listener that is willing to accept datagrams for any IP interface associated with the node, the value 0.0.0.0 is used. (netaddress)

`udpLocalPort` – The local port number for this UDP listener. (int)

`snmp` The `snmp` group reports statistics about the SNMP group.

`snmpInPkts` – The total number of Messages delivered to the SNMP entity from the transport service. (counter)

`snmpOutPkts` – The total number of SNMP Messages passed from the SNMP protocol entity to the transport service. (counter)

`snmpInBadVersions` – The total number of SNMP Messages delivered to the SNMP protocol entity that were for an unsupported SNMP version. (counter)

`snmpInBadCommunityNames` – The total number of SNMP Messages delivered to the SNMP protocol entity that used a SNMP community name not known to said entity. (counter)

`snmpInBadCommunityUses` – The total number of SNMP Messages delivered to the SNMP protocol entity, which represented an SNMP operation not allowed by the SNMP community named in the Message. (counter)

`snmpInASNParseErrs` – The total number of ASN.1 or BER errors encountered by the SNMP protocol entity when decoding received SNMP Messages. (counter)

`snmpInTooBigs` – The total number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is “tooBig.” (counter)

`snmpInNoSuchNames` – The total number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is “noSuchName.” (counter)

`snmpInBadValues` – The total number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is “badValue.” (counter)

snmpInReadOnlys – The total number valid SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is “readOnly.” It should be noted that it is a protocol error to generate an SNMP PDU that contains the value “readOnly” in the error-status field. This object is provided as a means of detecting incorrect implementations of the SNMP. (counter)

snmpInGenErrs – The total number of SNMP PDUs delivered to the SNMP protocol entity for which the value of the error-status field is “genErr.” (counter)

snmpInTotalReqVars – The total number of MIB objects successfully retrieved by the SNMP protocol entity as the result of receiving valid SNMP Get-Request and Get-Next PDUs. (counter)

snmpInTotalSetVars – The total number of MIB objects successfully altered by the SNMP protocol entity as the result of receiving valid SNMP Set-Request PDUs. (counter)

snmpInGetRequests – The total number of SNMP Get-Request PDUs accepted and processed by the SNMP protocol entity. (counter)

snmpInGetNexts – The total number of SNMP Get-Next PDUs accepted and processed by the SNMP protocol entity. (counter)

snmpInSetRequests – The total number of SNMP Set-Request PDUs accepted and processed by the SNMP protocol entity. (counter)

snmpInGetResponses – The total number of SNMP Get-Response PDUs accepted and processed by the SNMP protocol entity. (counter)

snmpInTraps – The total number of SNMP Trap PDUs accepted and processed by the SNMP protocol entity. (counter)

snmpOutTooBigs – The total number of SNMP PDUs generated by the SNMP protocol entity for which the value of the error-status field is “tooBig.” (counter)

snmpOutNoSuchNames – The total number of SNMP PDUs generated by the SNMP protocol entity for which the value of the error-status is “noSuchName.” (counter)

snmpOutBadValues – The total number of SNMP PDUs generated by the SNMP protocol entity for which the value of the error-status field is “badValue.” (counter)

snmpOutGenErrs – The total number of SNMP PDUs generated by the SNMP protocol entity for which the value of the error-status field is “genErr.” (counter)

snmpOutGetRequests – The total number of SNMP Get-Request PDUs which have been generated by the SNMP protocol entity. (counter)

snmpOutGetNexts – The total number of SNMP Get-Next PDUs generated by the SNMP protocol entity. (counter)

- snmpOutSetRequests – The total number of SNMP Set-Request PDUs generated by the SNMP protocol entity. (counter)
- snmpOutGetResponses – The total number of SNMP Get-Response PDUs generated by the SNMP protocol entity. (counter)
- snmpOutTraps – The total number of SNMP Trap PDUs generated by the SNMP protocol entity. (counter)
- snmpEnableAuthenTraps – Indicates whether the SNMP agent process is permitted to generate authentication-failure traps. The value of this object overrides any configuration information. As such, it provides a means whereby all authentication-failure traps may be disabled. (enum)

Note that this object must be stored in non-volatile memory, so that it remains constant between reinitializations of the network management system.

The following are Sun-specific group and table definitions.

sunSystem

- The sunSystem group reports general system information.
- agentDescr – The SNMP agent's description of itself. (string[255])
- hostID – The unique Sun hardware identifier. The value returned is four byte binary string. (octet[4])
- motd – The first line of /etc/motd. (string[255])
- unixTime – The UNIX system time. Measured in seconds since January 1, 1970 UTC. (counter)

sunProcessTable

- The sunProcessTable table reports UNIX process table information.
- psProcessID – The process identifier for this process. (int)
- psParentProcessID – The process identifier of this process's parent. (int)
- psProcessSize – The combined size of the data and stack segments (in kilobytes.) (int)
- psProcessCpuTime – The CPU time (including both user and system time) consumed so far. (int)
- psProcessState – The run-state of the process. (octet[4])

R	Runnable
T	Stopped
P	In page wait

D	Non-interruptable wait
S	Sleeping (less than 20 seconds)
I	Idle (more than 20 seconds)
Z	Zombie

psProcessWaitChannel – Reason process is waiting. (octet[16])

psProcessTTY – Terminal, if any, controlling this process. (octet[16])

psProcessUserName – Name of the user associated with this process. (octet[16])

psProcessUserID – Numeric form of the name of the user associated with this process. (int)

psProcessName – Command name used to invoke this process. (octet[64])

psProcessStatus – Setting this variable will cause a signal of the set value to be sent to the process. (int)

sunHostPerf

The sunHostPerf group reports hostperf information.

rsUserProcessTime – Total number of timeticks used by user processes since the last system boot. (counter)

rsNiceModeTime – Total number of timeticks used by “nice” mode since the last system boot. (counter)

rsSystemProcessTime – Total number of timeticks used by system processes since the last system boot. (counter)

rsIdleModeTime – Total number of timeticks in idle mode since the last system boot. (counter)

rsDiskXfer1 – Total number of disk transfers since the last boot for the first of four configured disks. (counter)

rsDiskXfer2 – Total number of disk transfers since the last boot for the second of four configured disks. (counter)

rsDiskXfer3 – Total number of disk transfers since the last boot for the third of four configured disks. (counter)

rsDiskXfer4 – Total number of disk transfers since the last boot for the fourth of four configured disks. (counter)

rsVPagesIn – Number of pages read in from disk. (counter)

rsVPagesOut – Number of pages written to disk. (counter)

rsVSwapIn – Number of pages swapped in. (counter)
rsVSwapOut – Number of pages swapped out. (counter)
rsVIntr – Number of device interrupts. (counter)
rsIfInPackets – Number of input packets. (counter)
rsIfOutPackets – Number of output packets. (counter)
rsIfInErrors – Number of input errors. (counter)
rsIfOutErrors – Number of output errors. (counter)
rsIfCollisions – Number of output collisions. (counter)

文件

/etc/snmp/conf/snmpd.conf configuration information
/etc/snmp/conf/mibiisa.acl access control file
/var/snmp/mib/sun.mib standard SNMP MIBII file

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/snmp/mibiisa
Interface Stability	Obsolete

另请参见 [inetd\(1M\)](#), [select\(3C\)](#), [recvfrom\(3SOCKET\)](#), [sendto\(3SOCKET\)](#), [attributes\(5\)](#), [gld\(7D\)](#)

诊断

cannot dispatch request	The proxy cannot dispatch the request. The rest of the message indicates the cause of the failure.
select(3C) failed	A select(3C) call failed. The rest of the message indicates the cause of the failure.
sendto(3SOCKET) failed	A sendto(3SOCKET) call failed. The rest of the message indicates the cause of the failure.
recvfrom(3SOCKET) failed	A recvfrom(3SOCKET) call failed. The rest of the message indicates the cause of the failure.
no response from system	The SNMP agent on the target system does not respond to SNMP requests. This error might indicate that the

	SNMP agent is not running on the target system, the target system is down, or the network containing the target system is unreachable.
response too big	The agent could not fit the results of an operation into a single SNMP message. Split large groups or tables into smaller entities.
missing attribute	An attribute is missing from the requested group.
bad attribute type	An object attribute type received from the SNMP agent that does not match the attribute type specified by the proxy agent schema. The rest of the message indicates the expected type and received type.
cannot get sysUpTime	The proxy agent cannot get the variable <i>sysUpTime</i> from the SNMP agent.
sysUpTime type bad	The variable <i>sysUpTime</i> received from the SNMP agent has the wrong data type.
unknown SNMP error	An unknown SNMP error was received.
bad variable value	The requested specified an incorrect syntax or value for a set operation.
variable is read only	The SNMP agent did not perform the set request because a variable to set may not be written.
general error	A general error was received.
cannot make request PDU	An error occurred building a request PDU.
cannot make request varbind list	An error occurred building a request variable binding list.
cannot parse response PDU	An error occurred parsing a response PDU.

request ID - response ID mismatch	The response ID does not match the request ID.
string contains non-displayable characters	A displayable string contains non-displayable characters.
cannot open schema file	An error occurred opening the proxy agent schema file.
cannot parse schema file	The proxy agent couldn't parse the proxy agent schema file.
cannot open host file	An error occurred opening the file associated with the <i>na.snmp.hostfile</i> keyword in <i>/etc/snmp/conf/snmpd.conf</i>
cannot parse host file	The proxy agent was unable to parse the file associated with the <i>na.snmp.hostfile</i> keyword in <i>/etc/snmp/conf/snmp.conf</i> .
attribute unavailable for set operations	The set could not be completed because the attribute was not available for set operations.

已知问题

The mibiisa utility returns the wrong interface speed for the SBUS FDDI interface (for example, “bf0”).

The mibiisa utility does not return a MAC address for the SBUS FDDI interface (for example, “bf0”).

Process names retrieved from mibiisa contain a leading blank space.

When you change attribute values in the system group with an SNMP set request, the change is effective only as long as mibiisa is running. mibiisa does not save the changes to */etc/snmp/conf/snmpd.conf*.

引用名	mkbootmedia – create bootable Solaris ISO image						
用法概要	<code>/usr/bin/mkbootmedia -v [-l label] media-root iso</code>						
描述	The <code>mkbootmedia</code> utility takes <i>media-root</i> (the root of an on-disk Solaris install media) as input and creates a bootable Solaris ISO image in the file <i>iso</i>, using <code>mkisofs(8)</code>. The file can then be burned onto a CD/DVD with utilities such as <code>cdwrw(1)</code> or <code>cdrecord(1)</code>. (Neither <code>mkisofs(8)</code> nor <code>cdrecord(1)</code> are SunOS man pages.) 注意 – The directory tree <i>media-root</i> must contain the file <code>boot/grub/stage2_eltorito</code>, which will be written to the media boot sectors. This file will be modified with some boot information, thus it must be writable. If necessary, first save a copy prior to running this utility.						
选项	The following options are supported: <code>-l label</code> Sets <i>label</i> as the label/volume name of the ISO image. <code>-v</code> Verbose. Multiple <code>-v</code> options increase verbosity.						
操作数	The following operands are supported: <i>media-root</i> Top-level directory of an on-disk Solaris install media. <i>iso</i> Name of the output file which will contain the resulting ISO image.						
示例	示例 1 Creating an ISO Image and Burning a CD/DVD The following commands create an ISO image from the content of <code>s10u1</code> and burn the image to a CD/DVD. <pre># /usr/bin/mkbootmedia s10u1 s10u1.iso # /usr/bin/cdrw -i s10u1.iso</pre>						
属性	See attributes(5) for descriptions of the following attributes:						
<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>system/core-os</td></tr><tr><td>Interface Stability</td><td>Committed</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	system/core-os	Interface Stability	Committed
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	system/core-os						
Interface Stability	Committed						
另请参见	cdwrw(1), attributes(5) <code>mkisofs(8)</code>, (<code>/usr/share/man/man8/mkisofs.8</code>), in the <code>SUNWfsman</code> package (not a SunOS man page)						

引用名 mkdevalloc – Make device_allocate entries

用法概要 /usr/sbin/mkdevalloc

描述 The mkdevalloc command writes to standard out a set of [device_allocate\(4\)](#) entries describing the system's frame buffer, audio, and removable media devices.

The mkdevalloc command is used by the device allocation service:

svc:/system/device/allocate:default

...to create or update the [device_allocate\(4\)](#) file. The device allocation service is managed by [smf\(5\)](#) and described in [device_allocate\(1M\)](#).

Entries are generated based on the device special files found in /dev. For the different categories of devices, the mkdevalloc command checks for the following files under /dev:

```
audio          /dev/audio, /dev/audioctl, /dev/sound/...
tape           /dev/rst*, /dev/nrst*, /dev/rmt/...
removable disk /dev/sr*, /dev/nsr*, /dev/dsk/c0t?d0s?, /dev/rdisk/c0t?d0s?
frame buffer   /dev/fb
```

All entries set the *device-minimum* and *device-maximum* fields to the hex representations of ADMIN_LOW and ADMIN_HIGH, respectively. The *device-authorization* field is set to solaris.device.allocate, except for the framebuffer entry, where it is set to *. The *device-name*, *device-type* and *device-clean* fields are set to the following values:

	device-name	device-type	device-clean
audio	audio	audio	audio_clean_wrapper
tape	mag_tape_0,1,...	st	st_clean
floppy	floppy_0,1,...	fd	disk_clean
removable disk	cdrom_0,1,...	sr	disk_clean
frame buffer	framebuffer	fb	/bin/true

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Obsolete

另请参见 [allocate\(1\)](#), [device_allocate\(1M\)](#), [device_allocate\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 mkdevalloc might not be supported in a future release of the Solaris operating system.

引用名	mkdevmaps – make device_maps entries						
用法概要	/usr/sbin/mkdevmaps						
描述	The mkdevmaps command writes to standard out a set of device_maps(4) entries describing the system's frame buffer, audio, and removable media devices. The mkdevmaps command is used by the device allocation service: <pre>svc:/system/device/allocate:default</pre> ...to create or update the device_maps(4) file. The device allocation service is managed by smf(5) and described in device_allocate(1M). Entries are generated based on the device special files found in /dev. For the different categories of devices, the mkdevmaps command checks for the following files under /dev: <pre>audio /dev/audio, /dev/audioctl, /dev/sound/... tape /dev/rst*, /dev/nrst*, /dev/rmt/... removable disk /dev/dsk/c0t?d0s?, /dev/rdsk/c0t?d0s? frame buffer /dev/fb</pre>						
属性	See attributes(5) for descriptions of the following attributes: <table><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>system/core-os</td></tr><tr><td>Interface Stability</td><td>Obsolete</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	system/core-os	Interface Stability	Obsolete
ATTRIBUTE TYPE	ATTRIBUTE VALUE						
Availability	system/core-os						
Interface Stability	Obsolete						
另请参见	allocate(1) , device_allocate(1M) , device_maps(4) , attributes(5) , smf(5)						
附注	mkdevmaps might not be supported in a future release of the Solaris operating system.						

引用名	mkfifo – make FIFO special file
用法概要	<code>/usr/bin/mkfifo [-m <i>mode</i>] <i>path</i>...</code>
描述	The <code>mkfifo</code> utility creates the FIFO special files named by its argument list. The arguments are taken sequentially, in the order specified; and each FIFO special file is either created completely or, in the case of an error or signal, not created at all. If errors are encountered in creating one of the special files, <code>mkfifo</code> writes a diagnostic message to the standard error and continues with the remaining arguments, if any. The <code>mkfifo</code> utility calls the library routine <code>mkfifo(3C)</code>, with the <i>path</i> argument is passed as the <i>path</i> argument from the command line, and <i>mode</i> is set to the equivalent of <code>a=rw</code>, modified by the current value of the file mode creation mask <code>umask(1)</code>.
选项	The following option is supported: <code>-m <i>mode</i></code> Set the file permission bits of the newly-created FIFO to the specified <i>mode</i> value. The <i>mode</i> option-argument will be the same as the <i>mode</i> operand defined for the <code>chmod(1)</code> command. In <code><symbolicmode></code> strings, the <i>op</i> characters <code>+</code> and <code>-</code> will be interpreted relative to an assumed initial mode of <code>a=rw</code>.
操作数	The following operand is supported: <i>file</i> A path name of the FIFO special file to be created.
用法	See <code>largefile(5)</code> for the description of the behavior of <code>mkfifo</code> when encountering files greater than or equal to 2 Gbyte (2 ³¹ bytes).
环境变量	See <code>environ(5)</code> for descriptions of the following environment variables that affect the execution of <code>mkfifo</code> : <code>LANG</code> , <code>LC_ALL</code> , <code>LC_CTYPE</code> , <code>LC_MESSAGES</code> , and <code>NLSPATH</code> .
退出状态	The following exit values are returned: <code>0</code> All the specified FIFO special files were created successfully. <code>>0</code> An error occurred.
属性	See <code>attributes(5)</code> for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed
Standard	See <code>standards(5)</code> .

另请参见

[chmod\(1\)](#), [umask\(1\)](#), [mkfifo\(3C\)](#), [attributes\(5\)](#), [environ\(5\)](#), [largefile\(5\)](#), [standards\(5\)](#)

引用名	mkfile – create a file
用法概要	mkfile [-nv] size [t g k b m] filename...
描述	mkfile creates one or more files that are suitable for use as NFS-mounted swap areas, or as local swap areas. When a root user executes <code>mkfile()</code> , the sticky bit is set and the file is padded with zeros by default. When non-root users execute <code>mkfile()</code> , they must manually set the sticky bit using <code>chmod(1)</code> . The default size is in bytes, but it can be flagged as terabytes, gigabytes, kilobytes, blocks, or megabytes, with the <code>t</code> , <code>g</code> , <code>k</code> , <code>b</code> , or <code>m</code> suffixes, respectively. Suffixes can be uppercase or lowercase.
选项	-n Create an empty <i>filename</i>. The size is noted, but disk blocks are not allocated until data is written to them. Files created with this option cannot be swapped over local UFS mounts. -v Verbose. Report the names and sizes of created files.
退出状态	The following exit values are returned: 0 Success. >0 An error occurred.
用法	See largefile(5) for the description of the behavior of <code>mkfile</code> when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [chmod\(1\)](#), [swap\(1M\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

引用名	mkfs – construct a file system								
用法概要	<pre>mkfs [-F FSType] [generic_options] [-o FSType-specific_options] raw_device_file [operands]</pre>								
描述	The <code>mkfs</code> utility constructs a file system on the <code>raw_device_file</code> by calling the specific <code>mkfs</code> module indicated by <code>-F FSType</code>. Note: ufs file systems are normally created with the newfs(1M) command. <i>generic_options</i> are independent of file system type. <i>FSType-specific_options</i> is a comma-separated list of <i>keyword=value</i> pairs (with no intervening spaces), which are <i>FSType</i>-specific. <i>raw_device_file</i> specifies the disk partition on which to write the file system. It is required and must be the first argument following the <i>specific_options</i> (if any). <i>operands</i> are <i>FSType</i>-specific. See the <i>FSType</i>-specific manual page of <code>mkfs</code> (for example, mkfs_ufs(1M)) for a detailed description.								
选项	The following are the generic options for <code>mkfs</code>: -F Specify the <i>FSType</i> to be constructed. If <code>-F</code> is not specified, the <i>FSType</i> is determined from <code>/etc/vfstab</code> by matching the <i>raw_device_file</i> with a <code>vfstab</code> entry, or by consulting the <code>/etc/default/fs</code> file. -V Echo the complete command line, but do not execute the command. The command line is generated by using the options and arguments provided and adding to them information derived from <code>/etc/vfstab</code> or <code>/etc/default/fs</code>. This option may be used to verify and validate the command line. -m Return the command line which was used to create the file system. The file system must already exist. This option provides a means of determining the command used in constructing the file system. -o Specify <i>FSType</i>-specific options. See the manual page for the <code>mkfs</code> module specific to the file system type.								
用法	See largefile(5) for the description of the behavior of <code>mkfs</code> when encountering files greater than or equal to 2 Gbyte (2 ³¹ bytes).								
文件	<table><tr><td><code>/etc/default/fs</code></td><td>Default file system type. Default values can be set for the following flags in <code>/etc/default/fs</code>. For example: <code>LOCAL=ufs</code></td></tr><tr><td></td><td><table><tr><td>LOCAL</td><td>The default partition for a command if no <i>FSType</i> is specified.</td></tr></table></td></tr><tr><td><code>/etc/vfstab</code></td><td>List of default parameters for each file system</td></tr></table>	<code>/etc/default/fs</code>	Default file system type. Default values can be set for the following flags in <code>/etc/default/fs</code> . For example: <code>LOCAL=ufs</code>		<table><tr><td>LOCAL</td><td>The default partition for a command if no <i>FSType</i> is specified.</td></tr></table>	LOCAL	The default partition for a command if no <i>FSType</i> is specified.	<code>/etc/vfstab</code>	List of default parameters for each file system
<code>/etc/default/fs</code>	Default file system type. Default values can be set for the following flags in <code>/etc/default/fs</code> . For example: <code>LOCAL=ufs</code>								
	<table><tr><td>LOCAL</td><td>The default partition for a command if no <i>FSType</i> is specified.</td></tr></table>	LOCAL	The default partition for a command if no <i>FSType</i> is specified.						
LOCAL	The default partition for a command if no <i>FSType</i> is specified.								
<code>/etc/vfstab</code>	List of default parameters for each file system								
属性	See attributes(5) for descriptions of the following attributes:								

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [mkfs_ufs\(1M\)](#), [newfs\(1M\)](#), [vfstab\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

Manual pages for the FSType-specific modules of `mkfs`.

附注 This command might not be supported for all *FSTypes*.

You can use `lofiadm` to create a file that appears to a `mkfs` command as a raw device. You can then use a `mkfs` command to create a file system on that device. See [lofiadm\(1M\)](#) for examples of creating a UFS and a PC (FAT) file system (using [mkfs_ufs\(1M\)](#) and [mkfs_pcfs\(1M\)](#)) on a device created by `lofiadm`.

引用名	mkfs_pcfs – construct a FAT file system
用法概要	<pre>mkfs -F pcfs [<i>generic_options</i>] [-o <i>FSType_specific_options</i>] <i>raw_device_file</i></pre>
描述	The pcfs-specific module of mkfs constructs a File Allocation Table (FAT) on removable media (JAZ disk, ZIP disk, PCMCIA card), a hard disk, or a file (see NOTES). FATs are the standard MS-DOS and Windows file system format. mkfs for pcfs determines an appropriate FAT size for the medium, then it installs an initial boot sector and an empty FAT. A sector size of 512 bytes is used. mkfs for pcfs can also install the initial file in the file system (see the pcfs-specific -o i option). This first file can optionally be marked as read-only, system, and/or hidden. If you want to construct a FAT with mkfs for pcfs on a medium that is not formatted, you must first perform a low-level format on the medium with format(1M). The media must also be partitioned with the fdisk(1M) utility. Note that all existing data on a disk partition, if any, is destroyed when a new FAT is constructed. <i>generic_options</i> are supported by the generic mkfs command. See mkfs(1M) for a description of these options. <i>raw_device_file</i> indicates the device on which to write unless the -o N option has been specified, or if the -V or -m generic options are passed from the generic mkfs module.
选项	See mkfs(1M) for the list of supported generic options. The following options are supported: -o <i>FSType_specific_options</i> Specify pcfs file system-specific options in a comma-separated list with no intervening spaces. If invalid options are specified, a warning message is printed and the invalid options are ignored. b=<i>label</i> Label the media with volume label. The volume label is restricted to 11 uppercase characters. B=<i>filename</i> Install <i>filename</i> as the boot loader in the file system's boot sector. If you don't specify a boot loader, an MS-DOS boot loader is installed. The MS-DOS boot loader requires specific MS-DOS system files to make a disk bootable. See NOTES for more information. fat=<i>n</i> The size of a FAT entry. Currently, 12, 16, and 32 are valid values. The default is 16. h Mark the first file installed as a hidden file. The -i option must also be specified.

hidden=*n*

Set the number of hidden sectors to *n*. This is the number of sectors on the physical disk preceding the start of the volume (which is the boot sector itself). This defaults to a computed value (based on the fdisk table) for disks. This option may be used only in conjunction with the `nofdisk` option.

i=*filename*

Install *filename* as the initial file in the new file system. The initial file's contents are guaranteed to occupy consecutive clusters at the start of the files area. When creating bootable media, a boot program should be specified as the initial file.

nofdisk

Do not attempt to find an fdisk table on the medium. Instead rely on the `size` option for determining the partition size. By default, the created FAT is 16 bits and begins at the first sector of the device. This origination sector can be modified with the `hidden` option (`-h`).

nsect=*n*

The number of sectors per track on the disk. If not specified, the value is determined by using a `dkio(7I)` ioctl to get the disk geometry.

ntrack=*n*

The number of tracks per cylinder on the disk. If not specified, the value is determined by using a `dkio(7I)` ioctl to get the disk geometry.

N

No execution mode. Print normal output, but do not actually write the file system to the medium. This is most useful when used in conjunction with the `verbose` option.

r

Mark the first file installed as read-only. The `-i` option must also be specified.

reserve=*n*

Set the number of reserved sectors to *n*. This is the number of sectors in the volume, preceding the start of the first FAT, including the boot sector. The value should always be at least 1, and the default value is exactly 1.

s

Mark the first file installed as a system file. The `-i` option must also be specified.

size=*n*

The number of sectors in the file system. If not specified, the value is determined from the size of the partition given in the fdisk table.

spc=*n*

The size of the allocation unit for space within the file system, expressed as a number of sectors. The default value depends on the FAT entry size and the size of the file system.

v

Verbose output. Describe, in detail, operations being performed.

文件

raw_device_file

The device on which to build the FAT.

mkfs_pcfs supports MBR (Master Boot Record) partitions and GPT (GUID Partition Table) partitions. GPT is part of the EFI (Extensible Firmware Interface) standard. For both x86 and SPARC, for a GPT-labeled disk, you can specify the partition using the logical device pathname with no suffix, for example, `/dev/rdisk/c0t0d0s0`. In GPT, this corresponds to the first partition on the disk.

On x86 for MBR partitions, you can specify the proper partition using the logical device pathname corresponding to the partition. For example, `/dev/rdisk/c0t0d0p1` corresponds to first partition in the MBR, or `/dev/rdisk/c0t0d0p5` corresponds to first logical partition in the extended partition. Alternatively, using a suffix is also acceptable. For example, in `/dev/rdisk/c0t0d0p0:c`, mkfs_pcfs recognizes `:c` as the first partition that can accept a FAT file system.

For removable media with MBR partitions on SPARC, you need to specify a disk device name with a suffix to indicate the proper partition. For example, in the name `/dev/rdisk/c0t0d0s2:c`, the `:c` suffix indicates that the partition can accept the new FAT.

For a file, *raw_device_file* is the block device name returned by `lofiadm(1M)`.

示例

The media in these examples must be formatted before running mkfs for pcfs. See DESCRIPTION for more details.

示例 1 Creating a FAT File System on a Disk

The following command creates a FAT file system on the second fdisk partition of a disk attached to an x86 based system:

```
mkfs -F pcfs /dev/rdisk/c0d0p0:d
```

示例 2 Creating a FAT File System on a ZIP Disk

The following command creates a FAT file system on a ZIP disk located on a SPARC based system:

```
mkfs -F pcfs /dev/rdisk/c0t4d0s2:c
```

示例 3 Creating a FAT File System on a JAZ Disk

The following command creates a FAT file system on a JAZ disk located on a SPARC based system and overrides the sectors/track and tracks/cylinder values obtained from the device's controller:

```
mkfs -F pcfs -o nsect=32,ntrack=64 /dev/rdisk/c0t3d0s2:c
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/pcfs
Interface Stability	Committed

另请参见

[fdisk\(1M\)](#), [format\(1M\)](#), [lofiadm\(1M\)](#), [mkfs\(1M\)](#), [attributes\(5\)](#), [dkio\(7I\)](#)

附注

You can use `lofiadm` to create a file that appears to a `mkfs` command (for example, `mkfs_pcfs` or `mkfs_ufs`) as a raw device. You can then use a `mkfs` command to create a file system on that device. See [lofiadm\(1M\)](#) for examples of creating a UFS and a PC (FAT) file system on a device created by `lofiadm`.

引用名 mkfs_udfs – construct a udfs file system

用法概要 mkfs -F udfs [*generic_options*] [-o *specific_options*] *raw_device_file* [*size*]

描述 This is the universal disk format file system (udfs) -specific module of the mkfs command. mkfs constructs a udfs file system with a root directory.

选项 See mkfs(1M) for the list of supported *generic_options*.

The following options are supported:

-o *specific_options* Specify a udfs-specific option. Specify udfs file system specific options in a comma-separated list with no intervening spaces. If invalid options are specified, a warning message is printed and the invalid options are ignored.

The following *specific_options* are available:

N Print the file system parameters without actually creating the file system.

label=*string* Specify the label to be written into the volume header structures. Specify *string* as the name of the label. If *string* is not specified, a default *string* is generated in the form of *NoLabel*.

操作数 The following operands are supported:

raw_device_file Specify the disk partition on which to write.

size Specify the number of 512-byte blocks in the file system.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/udfs

另请参见 [fsck\(1M\)](#), [mkfs\(1M\)](#), [attributes\(5\)](#)

诊断 not currently a valid file system

The specified device does not contain a valid udfs file system.

Invalid size: larger than the partition size

Number of blocks given as parameter to create the file system is larger than the size of the device specified.

is mounted can't mkfs

Device is in use, cannot create file system when the device is in use.

preposterous size

Negative size parameter provided is invalid.

sector size must be between 512, 8192 bytes

Sector size given is not in the valid range.

Volume integrity sequence descriptors too long

File set descriptor too long.

Not enough space to create volume integrity sequence or file set descriptor.

mkfs: argument out of range

One of the arguments is out of range.

mkfs: bad numeric arg

One of the arguments is potentially a bad numeric.

附注

You can use `lofiadm` to create a file that appears to a `mkfs` command (for example, `mkfs_pcfs` or `mkfs_ufs`) as a raw device. You can then use a `mkfs` command to create a file system on that device. See [lofiadm\(1M\)](#) for examples of creating a UFS and a PC (FAT) file system on a device created by `lofiadm`.

引用名	mkfs_ufs – construct a UFS file system												
用法概要	<pre>mkfs -F ufs [<i>generic_options</i>] [-o <i>FSType_specific_options</i>] <i>raw_device_file</i> [<i>size</i>]</pre>												
描述	The UFS-specific module of <code>mkfs</code> builds a UFS file system with a root directory and a <code>lost+found</code> directory (see fsck(1M)). The UFS-specific <code>mkfs</code> is rarely run directly. Use the newfs(1M) command instead. <i>raw_device_file</i> indicates the disk partition on which to create the new file system. If the <code>-o N</code>, <code>-V</code>, or <code>-m</code> options are specified, the <i>raw_device_file</i> is not actually modified. <i>size</i> specifies the number of disk sectors in the file system, where a disk sector is usually 512 bytes. This argument must follow the <i>raw_device_file</i> argument and is required (even with <code>-o N</code>), unless the <code>-V</code> or <code>-m</code> generic options are specified. <i>generic_options</i> are supported by the generic <code>mkfs</code> command. See mkfs(1M) for a description of these options.												
选项	The following generic options are supported: <code>-m</code> Print the command line that was used to create the existing file system. <code>-V</code> Print the current <code>mkfs</code> command line.												
选项	The following UFS-specific options are supported: <code>-o</code> Use one or more of the following values separated by commas (with no intervening spaces) to specify UFS-specific options:<table><tr><td><code>apc=<i>n</i></code></td><td>The number of alternate sectors per cylinder to reserve for bad block replacement for SCSI devices only. The default is 0.</td></tr><tr><td></td><td>This option is not applicable for disks with EFI labels and is ignored.</td></tr><tr><td><code>bsize=<i>n</i></code></td><td>The logical block size of the file system in bytes, either 4096 or 8192. The default is 8192. The sun4u architecture does not support the 4096 block size.</td></tr><tr><td><code>calcbinsb</code></td><td>Sends to stdout a binary (machine-readable) version of the superblock that would be used to create a file system with the specified configuration parameters.</td></tr><tr><td><code>calcsb</code></td><td>Sends to stdout a human-readable version of the superblock that would be used to create a file system with the specified configuration parameters.</td></tr><tr><td><code>cgsize=<i>n</i></code></td><td>The number of cylinders per cylinder group, ranging from 16 to 256. The default is calculated by dividing the number of sectors in the file system by the number of sectors in a gigabyte. Then, the result is multiplied by 32. The default value is always between 16 and 256.</td></tr></table>	<code>apc=<i>n</i></code>	The number of alternate sectors per cylinder to reserve for bad block replacement for SCSI devices only. The default is 0.		This option is not applicable for disks with EFI labels and is ignored.	<code>bsize=<i>n</i></code>	The logical block size of the file system in bytes, either 4096 or 8192. The default is 8192. The sun4u architecture does not support the 4096 block size.	<code>calcbinsb</code>	Sends to stdout a binary (machine-readable) version of the superblock that would be used to create a file system with the specified configuration parameters.	<code>calcsb</code>	Sends to stdout a human-readable version of the superblock that would be used to create a file system with the specified configuration parameters.	<code>cgsize=<i>n</i></code>	The number of cylinders per cylinder group, ranging from 16 to 256. The default is calculated by dividing the number of sectors in the file system by the number of sectors in a gigabyte. Then, the result is multiplied by 32. The default value is always between 16 and 256.
<code>apc=<i>n</i></code>	The number of alternate sectors per cylinder to reserve for bad block replacement for SCSI devices only. The default is 0.												
	This option is not applicable for disks with EFI labels and is ignored.												
<code>bsize=<i>n</i></code>	The logical block size of the file system in bytes, either 4096 or 8192. The default is 8192. The sun4u architecture does not support the 4096 block size.												
<code>calcbinsb</code>	Sends to stdout a binary (machine-readable) version of the superblock that would be used to create a file system with the specified configuration parameters.												
<code>calcsb</code>	Sends to stdout a human-readable version of the superblock that would be used to create a file system with the specified configuration parameters.												
<code>cgsize=<i>n</i></code>	The number of cylinders per cylinder group, ranging from 16 to 256. The default is calculated by dividing the number of sectors in the file system by the number of sectors in a gigabyte. Then, the result is multiplied by 32. The default value is always between 16 and 256.												

	The per-cylinder-group meta data must fit in a space no larger than what is available in one logical file system block. If too large a <i>cgsize</i> is requested, it is changed by the minimum amount necessary.
<code>fragsize=<i>n</i></code>	The smallest amount of disk space in bytes that can be allocated to a file. <i>fragsize</i> must be a power of 2 divisor of <i>bsize</i>, where: <i>bsize</i> / <i>fragsize</i> is 1, 2, 4, or 8. This means that if the logical block size is 4096, legal values for <i>fragsize</i> are 512, 1024, 2048, and 4096. When the logical block size is 8192, legal values are 1024, 2048, 4096, and 8192. The default value is 1024.
<code>free=<i>n</i></code>	For file systems greater than 1 terabyte or for file systems created with the <code>mtb=y</code> option, <i>fragsize</i> is forced to match block size (<i>bsize</i>). The minimum percentage of free space to maintain in the file system between 0% and 99%, inclusively. This space is off-limits to users. Once the file system is filled to this threshold, only the superuser can continue writing to the file system. The default is ((64 Mbytes/partition size) * 100), rounded down to the nearest integer and limited between 1% and 10%, inclusively. This parameter can be subsequently changed using the tunefs(1M) command.
<code>gap=<i>n</i></code>	Rotational delay. This option is obsolete in the Solaris 10 release. The value is always set to 0, regardless of the input value.
<code>maxcontig=<i>n</i></code>	The maximum number of logical blocks, belonging to one file, that are allocated contiguously. The default is calculated as follows: <i>maxcontig</i> = <i>disk drive maximum transfer size</i> / <i>disk block size</i> If the disk drive's maximum transfer size cannot be determined, the default value for <i>maxcontig</i> is calculated from kernel parameters as follows: If <i>maxphys</i> is less than <i>ufs_maxmaxphys</i>, which is typically 1 Mbyte, then <i>maxcontig</i> is set to <i>maxphys</i>. Otherwise, <i>maxcontig</i> is set to <i>ufs_maxmaxphys</i>. You can set <i>maxcontig</i> to any positive integer value. The actual value will be the lesser of what has been specified and what the hardware supports.

	You can subsequently change this parameter by using tunefs(1M) .
<code>mtb=y</code>	Set the parameters of the file system to allow eventual growth to over a terabyte in total file system size. This option sets <i>fragsize</i> to be the same as <i>bsize</i> , and sets <i>nbpi</i> to 1 Mbyte, unless the <code>-i</code> option is used to make it even larger. If you explicitly set the <i>fragsize</i> or <i>nbpi</i> parameters to values that are incompatible with this option, the user-supplied value of <i>fragsize</i> or <i>nbpi</i> is ignored.
<code>N</code>	Print out the file system parameters that would be used to create the file system without actually creating the file system.
<code>nbpi=n</code>	The number of bytes per inode, which specifies the density of inodes in the file system. The number is divided into the total size of the file system to determine the number of inodes to create. This value should reflect the expected average size of files in the file system. If fewer inodes are desired, a larger number should be used. To create more inodes, a smaller number should be given. The default is 2048. The number of inodes can increase if the file system is expanded with the <code>growfs</code> command.
<code>nrpos=n</code>	The number of different rotational positions in which to divide a cylinder group. The default is 8.
<code>nsect=n</code>	This option is not applicable for disks with EFI labels and is ignored. The number of sectors per track on the disk. The default is 32.
<code>ntrack=n</code>	The number of tracks per cylinder on the disk. The default is 16.
<code>opt=s t</code>	This option is not applicable for disks with EFI labels and is ignored. The file system can either be instructed to try to minimize the <i>time</i> spent allocating blocks, or to try to minimize the <i>space</i> fragmentation on the disk. The default is <i>time</i> . This parameter can be subsequently changed with the tunefs(1M) command.
<code>rps=n</code>	The rotational speed of the disk, in revolutions per second. The default is 60. Note that you specify <i>rps</i> for <code>mkfs</code> and <i>rpm</i> for <code>newfs</code> . This option is not applicable for disks with EFI labels and is ignored.

Alternatively, parameters can be entered as a list of space-separated values (without keywords) whose meaning is positional. In this case, the `-o` option is omitted and the list follows the size operand. This is the way `newfs` passes the parameters to `mkfs`.

操作数

The following operands are supported:

raw_device_file The disk partition on which to write.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[fsck\(1M\)](#), [mkfs\(1M\)](#), [newfs\(1M\)](#), [tunefs\(1M\)](#), [dir_ufs\(4\)](#), [attributes\(5\)](#), [ufs\(7FS\)](#)

诊断

The following error message typically occurs with very high density disks. On such disks, the file system structure cannot encode the proper disk layout information. However, such disks have enough onboard intelligence to make up for any layout deficiencies, so there is no actual impact on performance. The warning that performance might be impaired can be safely ignored.

Warning: insufficient space in super block for rotational layout tables with nsect *sblock.fs_nsect* and ntrak *sblock.fs_ntrak*. (File system performance may be impaired.)

The following error message occurs when the disk geometry results in a situation where the last truncated cylinder group cannot contain the correct number of data blocks. Some disk space is wasted.

Warning: inode blocks/cyl group (*grp*) >= data blocks (*num*) in last cylinder

If there is only one cylinder group and if the above condition holds true, `mkfs` fails with the following error:

File system creation failed. There is only one cylinder group and that is not even big enough to hold the inodes.

The following error message occurs when the best calculated file system layout is unable to include the last few sectors in the last cylinder group. This is due to the interaction between how much space is used for various pieces of meta data and the total blocks available in a cylinder group. Modifying `nbpi` and `cpg` might reduce this number, but it is rarely worth the effort.

Warning: *num* sector(s) in last cylinder group unallocated

附注

You can use `lofiadm` to create a file that appears to the `mkfs` command (for example, `mkfs_pcfs` or `mkfs_ufs`) as a raw device. You can then use the `mkfs` command to create a file system on that device. See [lofiadm\(1M\)](#) for examples of creating a UFS and a PC (FAT) file system on a device created by `lofiadm`.

Both the block and character devices, such as devices in `/dev/dsk` and `/dev/rdisk`, must be available prior to running the `mkfs` command.

引用名 mknod – make a special file

用法概要 mknod *name* b *major* *minor*
 mknod *name* c *major* *minor*
 mknod *name* p

描述 mknod makes a directory entry for a special file.

选项 The following options are supported:

b Create a block-type special file.
c Create a character-type special file.
p Create a FIFO (named pipe).

操作数 The following operands are supported:

major The *major* device number.

minor The *minor* device number; can be either decimal or octal. The assignment of major device numbers is specific to each system. You must be the super-user to use this form of the command.

name A special file to be created.

用法 See [largefile\(5\)](#) for the description of the behavior of `mknod` when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [ftp\(1\)](#), [mknod\(2\)](#), [symlink\(2\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

附注 If [mknod\(2\)](#) is used to create a device, the major and minor device numbers are always interpreted by the kernel running on that machine.

With the advent of physical device naming, it would be preferable to create a symbolic link to the physical name of the device (in the `/devices` subtree) rather than using `mknod`.

引用名	mkntfs – 创建 NTFS 文件系统
用法概要	<pre>mkntfs [options] device [number_of_sectors] mkntfs [-C] [-c cluster-size] [-F] [-f] [-H heads] [-h] [-I] [-L volume-label] [-l] [-n] [-p part-start-sect] [-Q] [-q] [-S sectors-per-track] [-s sector-size] [-T] [-V] [-v] [-z mft-zone-multiplier] [--debug] device [number-of-sectors]</pre>
描述	mkntfs 实用程序用于在设备上（通常是磁盘分区或文件）创建 NTFS 文件系统。<i>device</i> 操作数是对应于设备的特殊文件，例如 <code>/dev/dsk/c0d0p0</code>。<i>number-of-sectors</i> 操作数是设备上的块的数目。如果省略，则 mkntfs 将自动计算文件系统大小。
选项	下面列出了支持的选项。大多数选项都有单字母和全名两种形式。不接受参数的单字母选项可以多个组合到一起。例如，<code>-fv</code> 等效于 <code>-f -v</code>。全名选项可以缩写为其名称的唯一前缀。 选项划分为基本选项、高级选项、输出选项和帮助选项，如下所列。
基本选项	-C、--enable-compression 在卷上启用压缩。 -f、--fast 或 -q、--quick 执行快速格式化。此选项会跳过卷的零位调整和坏扇区检查。 -L、--label string 将文件系统的卷标签设置为 <i>string</i>。 -n、--no-action 使 mkntfs 不实际创建文件系统，但显示它创建文件系统时将执行的工作。将执行所有格式化步骤，只是不实际写入到磁盘。
高级选项	-c、--cluster-size bytes 以字节为单位指定簇的大小。有效的簇大小值是 2 的幂，最小值为 256 且最大值为 65536 字节/簇。如果省略，则 mkntfs 使用 4096 字节作为缺省的簇大小。 请注意，缺省簇大小设置为至少等于扇区大小，因为一个簇不能小于一个扇区。另请注意，大于 4096 的值有负面影响，即会在卷上禁用压缩。这是由 Windows 使用的 NTFS 压缩算法的局限性导致的。 -F、--force 强制 mkntfs 运行，即使指定的设备不是块特殊设备或者看起来是挂载的。 -H、--heads num 指定头的数量。最大值为 65535 (0xffff)。如果省略，mkntfs 会尝试自动确定头的数量。如果该尝试失败，则使用缺省值 0。请注意，要想使 Windows 能够从所创建的卷引导，必须指定 <i>num</i>。 -I、--no-indexing 在卷上禁用内容索引。此选项仅在 Windows 2000 和更高版本上有意义。Windows NT 4.0 和更早的版本会忽略此选项，因为它们没有实现内容索引。

- p、--partition-start *sector*
指定分区开始扇区。最大值为 4294967295 ($2^{32}-1$)。如果省略，mkntfs 会尝试自动确定 *sector*。如果该尝试失败，则使用缺省值 0。请注意，要想使 Windows 能够从所创建的卷引导，必须指定 *sector*。
- S、--sectors-per-track *num*
指定每个磁轨的扇区数。最大值为 65535 (0xffff)。如果省略，mkntfs 会尝试自动确定 sectors-per-track 的数目；如果该尝试失败，则使用缺省值 0。请注意，要想使 Windows 能够从所创建的卷引导，必须指定 sectors-per-track。
- s、--sector-size *bytes*
以字节为单位指定扇区的大小。有效的扇区大小值为 256、512、1024、2048 和 4096。如果省略，mkntfs 会尝试自动确定 sector-size。如果该尝试失败，则使用每扇区 512 字节这一缺省值。
- T、--zero-time
将时间伪造为 1970 年 1 月 1 日 00:00:00 UTC，而不使用当前的系统时间。这对于调试可能比较有用。
- z、--mft-zone-multiplier *num*
设置主文件表 (master file table, MFT) 区域乘数，它决定了要在卷上使用的 MFT 区域的大小。MFT 区域是在卷的开头为 MFT 保留的区域，它存储着盘上 inode (MFT 记录)。值得注意的是，小文件完全存储在 inode 内；因此，如果您期望将卷用于存储大量非常小的文件，则将区域倍数设置为一个较高的值比较有用。虽然 MFT 区域在 NTFS 驱动程序的运行过程中会根据需要动态调整大小，但是选择一个最佳值可以减少碎片。有效值为 1、2、3 和 4。这些值的含义如下所示：

MFT zone multiplier	MFT zone size (% of volume size)
1	12.5% (default)
2	25.0%
3	37.5%
4	50.0%

输出选项

- debug
包括来自 -v 选项的详细输出，以及对于调试 mkntfs 有用的其他输出。
- q、--quiet
以详细模式执行。错误将写入到 stderr，不会向 stdout 写入任何输出。如果在脚本中运行 mkntfs，则此选项比较有用。
- v、--verbose
以详细模式执行。

帮助选项

- h、--help
显示选项的列表及每个选项的简要说明。
- l、--license
显示 mkntfs 许可信息并退出。

-V、--version
显示 mkntfs 版本号并退出。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/ntfsprogs
接口稳定性	Uncommitted（未确定）

另请参见 [ntfsprogs\(1M\)](#)、[ntfsresize\(1M\)](#)、[ntfsundelete\(1M\)](#)、[attributes\(5\)](#)
<http://wiki.linux-ntfs.org>

作者 mkntfs 是由 Anton Altaparmakov、Richard Russon、Erik Sornes 和 Szabolcs Szakacsits 编写的。

引用名

mkpwdict – maintain password-strength checking database

用法概要

```
/usr/sbin/mkpwdict [-s dict1,...,dictN]
                  [-d destination-path]
```

描述

The `mkpwdict` command adds words to the dictionary-lookup database used by [pam_authtok_check\(5\)](#) and [passwd\(1\)](#).

Files containing words to be added to the database can be specified on the command-line using the `-s` flag. These source files should have a single word per line, much like `/usr/share/lib/dict/words`.

If `-s` is omitted, `mkpwdict` will use the value of `DICTIONLIST` specified in `/etc/default/passwd` (see [passwd\(1\)](#)).

The database is created in the directory specified by the `-d` option. If this option is omitted, `mkpwdict` uses the value of `DICTIONDBDIR` specified in `/etc/default/passwd` (see [passwd\(1\)](#)). The default location is `/var/passwd`.

选项

The following options are supported:

- `-s` Specifies a comma-separated list of files containing words to be added to the dictionary-lookup database.
- `-d` Specifies the target location of the dictionary-database.

文件

<code>/etc/default/passwd</code>	See passwd(1) .
<code>/var/passwd</code>	default destination directory

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [passwd\(1\)](#), [attributes\(5\)](#), [pam_authtok_check\(5\)](#)

引用名 modinfo – display information about loaded kernel modules

用法概要 /usr/sbin/modinfo [-c] [-w] [-i *module-id*]

描述 The modinfo utility displays information about the loaded modules. The format of the information is as follows:

Id Loadaddr Size Info Rev Module Name

where *Id* is the module ID, *Loadaddr* is the starting text address in hexadecimal, *Size* is the size of text, data, and bss in hexadecimal bytes, *Info* is module specific information, *Rev* is the revision of the loadable modules system, and *Module Name* is the filename and description of the module.

The module specific information is the block and character major numbers for drivers, the system call number for system calls, and unspecified for other module types.

选项 The following options are supported:

- c Display the number of instances of the module loaded and the module's current state.
- i *module-id* Display information about this module only.
- w Do not truncate module information at 80 characters.

示例 示例 1 Displaying the Status of a Module

The following example displays the status of module 2:

```
example% modinfo -i 2
Id Loadaddr Size Info Rev Module Name
2 ff08e000 1734 - 1 swapgeneric (root and swap configuration)
```

示例 2 Displaying the Status of Kernel Modules

The following example displays the status of some kernel modules:

```
example% modinfo
Id Loadaddr Size Info Rev Module Name
2 ff08e000 1734 - 1 swapgeneric
4 ff07a000 3bc0 - 1 specfs (filesystem for specfs)
6 ff07dbc0 2918 - 1 TS (time sharing sched class)
7 ff0804d8 49c - 1 TS_DPTBL (Time sharing dispatch table)
8 ff04a000 24a30 2 1 ufs (filesystem for ufs)
9 ff080978 c640 226 1 rpcmod (RPC syscall)
9 ff080978 c640 - 1 rpcmod (rpc interface str mod)
10 ff08cfb8 2031c - 1 ip (IP Streams module)
10 ff08cfb8 2031c 2 1 ip (IP Streams device)
```

示例 3 Using the -c Option

Using the `modinfo` command with the `-c` option displays the number of instances of the module loaded and the module's current state.

```
example% modinfo -c
Id      Loadcnt Module Name                               State
1        0 krtld                                     UNLOADED/UNINSTALLED
2        0 genunix                                    UNLOADED/UNINSTALLED
3        0 platmod                                       UNLOADED/UNINSTALLED
4        0 SUNW,UltraSPARC-IIi                         UNLOADED/UNINSTALLED
5        0 cl_bootstrap                                UNLOADED/UNINSTALLED
6        1 specfs                                       LOADED/INSTALLED
7        1 swapgeneric                                  UNLOADED/UNINSTALLED
8        1 TS                                           LOADED/INSTALLED
9        1 TS_DPTBL                                    LOADED/INSTALLED
10       1 ufs                                         LOADED/INSTALLED
11       1 fssnap_if                                    LOADED/INSTALLED
```

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [modload\(1M\)](#), [modunload\(1M\)](#), [attributes\(5\)](#)

引用名	modload – load a kernel module				
用法概要	modload [-p] [-e <i>exec_file</i>] <i>filename</i>				
描述	The <code>modload</code> command loads the loadable module <i>filename</i> into the running system. <i>filename</i> is an object file produced by <code>ld -r</code>. If <i>filename</i> is an absolute pathname then the file specified by that absolute path is loaded. If <i>filename</i> does not begin with a slash (/), then the path to load <i>filename</i> is relative to the current directory unless the <code>-p</code> option is specified. The kernel's <code>modpath</code> variable can be set using the <code>/etc/system</code> file. The default value of the kernel's <code>modpath</code> variable is set to the path where the operating system was loaded. Typically this is <code>/kernel /usr/kernel</code>. For example, the following command looks for <code>./drv/foo</code>: <pre>example# modload drv/foo</pre> The following command looks for <code>/kernel/drv/foo</code> and then <code>/usr/kernel/drv/foo</code>: <pre>example# modload -p drv/foo</pre>				
选项	The following options are supported: <code>-e <i>exec_file</i></code> Specify the name of a shell script or executable image file that is executed after the module is successfully loaded. The first argument passed is the module ID (in decimal). The other argument is module specific. The module specific information is: the block and character major numbers for drivers, the system call number for system calls, or, for other module types, the index into the appropriate kernel table. See modinfo(1M) <code>-p</code> Use the kernel's internal <code>modpath</code> variable as the search path for the module.				
属性	See attributes(5) for descriptions of the following attributes:				
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>system/core-os</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	system/core-os
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	system/core-os				
另请参见	ld(1) , add_drv(1M) , kernel(1M) , modinfo(1M) , modunload(1M) , system(4) , attributes(5) , modldrv(9S) , modlinkage(9S) , modlstrmod(9S) , module_info(9S)				
	《Writing Device Drivers》				
附注	Use add_drv(1M) to add device drivers, not <code>modload</code> . See 《Writing Device Drivers》 for procedures on adding device drivers.				

引用名 modunload – unload a module

用法概要 modunload -i *module_id* [-e *exec_file*]

描述 modunload unloads a loadable module from the running system. The *module_id* is the ID of the module as shown by [modinfo\(1M\)](#). If ID is 0, all modules that were autoloaded which are unloadable, are unloaded. Modules loaded by [modload\(1M\)](#) are not affected.

选项 The following options are supported:

-e *exec_file* Specify the name of a shell script or executable image file to be executed before the module is unloaded. The first argument passed is the module id (in decimal). There are two additional arguments that are module specific. For loadable drivers, the second argument is the driver major number. For loadable system calls, the second argument is the system call number. For loadable exec classes, the second argument is the index into the execsw table. For loadable filesystems, the second argument is the index into the vfsw table. For loadable streams modules, the second argument is the index into the fmodsw table. For loadable scheduling classes, the second argument is the index into the class array. Minus one is passed for an argument that does not apply.

-i *module_id* Specify the module to be unloaded.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [modinfo\(1M\)](#), [modload\(1M\)](#), [update_drv\(1M\)](#), [attributes\(5\)](#)

附注 The modunload command is often used on driver modules to force the system to reread the associated driver configuration file. While this works in the current Solaris release, it is not the supported way to reread the configuration file and is not guaranteed to work in future releases. The supported way for rereading driver configuration file is through the [update_drv\(1M\)](#) command.

引用名	mofcomp – compile MOF files into CIM classes
用法概要	<pre> /usr/sadm/bin/mofcomp [-c <i>cimom_hostname</i>] [-h] [-j <i>filename</i>] [-n <i>namespace</i>] [-o <i>dirname</i>] [-p <i>password</i>] [-CIQ] [-u <i>username</i>] [-v] [-version] [-x] <i>file</i> </pre>
描述	The <code>mofcomp</code> utility is executed during installation to compile managed object format (MOF) files that describe the Common Information Model (CIM) and Solaris Schemas into the CIM Object Manager Repository, a central storage area for management data. The CIM Schema is a collection of class definitions used to represent managed objects that occur in every management environment. The Solaris Schema is a collection of class definitions that extend the CIM Schema and represent managed objects in a typical Solaris operating environment. The <code>mofcomp</code> utility must be run as root or as a user with write access to the namespace in which you are compiling. MOF is a language for defining CIM classes and instances. MOF files are ASCII text files that use the MOF language to describe CIM objects. A CIM object is a computer representation or model of a managed resource, such as a printer, disk drive, or CPU. Many sites store information about managed resources in MOF files. Because MOF can be converted to Java, Java applications that can run on any system with a Java Virtual Machine can interpret and exchange this information. You can also use the <code>mofcomp</code> utility to compile MOF files at any time after installation.
选项	The following options are supported: -c <i>cimom_hostname</i> Specify a remote system running the CIM Object Manager. -C Run the compiler set with the class option, which updates a class if it exists, and returns an error if the class does not exist. If you do not specify this option, the compiler adds a CIM class to the connected namespace, and returns an error if the class already exists. -h List the arguments to the <code>mofcomp</code> utility. -I Run the compiler set with the instance option, which updates an instance if it exists, and returns an error if the instance does not exist. If you do not specify this option, the compiler adds a CIM instance to the connected namespace, and returns an error if the instance already exists. -j <i>filename</i> Generate Java Beans and Java Interfaces to manage the CIM instances related to the CIM classes in the MOF being compiled. The contents of <i>filename</i> are: <pre> PACKAGE=<i>Java package name</i> IMPORTS=<i>import1:...:importN</i> </pre>

<EXCEPTIONS=*exception1*:...:*exceptionN*

PACKAGE is a valid Java package name to include in all generated Java source. IMPORTS is an optional colon separated list of valid Java classes to be imported in all generated Java source. EXCEPTIONS is an optional colon separated list of valid Java exceptions to be thrown by the methods in all generated Java source.

- n *namespace* Requests that the compiler load the MOF file into the namespace specified as *namespace*. The default namespace (root\cimv2) is used unless this switch is used or a `#pragma namespace ("namespace")` statement appears in the MOF file. If both the -n *namespace* switch and the `#pragma namespace` construct are used, all namespaces are created, but the objects are created only in the `#pragma namespace`s.
- o *dirname* Run compiler in standalone mode, without the CIM Object Manager. Specify *dirname* as the directory in which the compiler output is to be stored. In this mode, the CIM Object Manager need not be running.
- p *password* Specify a password for connecting to the CIM Object Manager. Use this option for compilations that require privileged access to the CIM Object Manager. If you specify both -p and -u, you must type the password on the command line, which can pose a security risk. A more secure way to specify a password is to specify -u but not -p, so that the compiler will prompt for the password.
- Q Run the compiler set with the qualifier types option, which updates a qualifier type if it exists, and returns an error if the qualifier type does not exist. If you do not specify this option, the compiler adds a CIM qualifier type to the connected namespace, and returns an error if the qualifier type already exists.
- u *username* Specify user name for connecting to the CIM Object Manager. Use this option for compilations that require privileged access to the CIM Object Manager. If you specify both -p and -u, you must type the password on the command line, which can pose a security risk. A more secure way to specify a password is to specify -u but not -p, so that the compiler will prompt for the password.
- v Run the compiler in verbose mode, which displays compiler messages.
- version Display the version of the MOF compiler.
- x Generate XML documents for the CIM classes defined in the input MOF file.

- 操作数

The following operands are supported:
file The pathname of the file to be compiled.
- 退出状态

The mofcomp utility exits with 0 upon success and a positive integer upon failure.
- 文件

MOF files are installed in /usr/sadm/mof.
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWwbco

另请参见

[init.wbem\(1M\)](#), [mofreg\(1M\)](#), [wbemadmin\(1M\)](#), [wbemlogviewer\(1M\)](#), [attributes\(5\)](#), [wbem\(5\)](#)

引用名	mofreg – register MOF classes with WBEM services		
用法概要	<pre>/usr/sadm/bin/mofreg -r tag file /usr/sadm/bin/mofreg -s /usr/sadm/bin/mofreg -u tag [file]</pre>		
描述	The <code>mofreg</code> command is used by package and patch install scripts, or by any applications that wish to register managed object format (MOF) classes with Sun The Web-Based Enterprise Management (WBEM) services. The WBEM services daemon (Common Information Model or CIM object manager) processes at start up the files that are specified by <code>mofreg</code> commands. Files are processed in the order that the individual <code>mofreg</code> commands are executed. As an alternative to using the <code>mofreg</code> command, MOFs can be registered or unregistered by manipulating directories in <code>/var/sadm/wbem/logr</code>. Instead of running the <code>mofreg -r tag file</code> version of the command you can create a directory named <code>tag</code> under <code>/var/sadm/wbem/logr/preReg</code> and copy <code>file</code> to the <code>tag</code> directory. Similarly, instead of running the <code>mofreg -u tag [file]</code> command, you can create a directory named <code>tag</code> under <code>/var/sadm/wbem/logr/preUnreg</code> and copy the optional <code>file</code> to the <code>tag</code> directory. The entries are processed in increasing order of last modification time of the <code>tag</code> directories. If you issue <code>mofreg</code> commands in rapid succession, the timestamps might be the same. If you have a situation where the timestamp order is critical, you can place appropriate sleeps between the successive registration or unregistration operations. As with the <code>mofreg</code> command, processing is done at next restart or by using the <code>-s</code> option. This alternative mechanism is typically used in package install scripts which do not have access to <code>/usr</code>, and therefore do not have access to the <code>mofreg</code> command. This case arises when packages are installed for diskless clients.		
选项	The following options are supported: <table><tr><td><code>-r tag file</code></td><td> The <i>file</i> argument is the actual MOF registration file. Its form is identical to the MOF syntax as defined by the Distributed Management Task Force (DMTF). The only difference is the addition of the following 3 new pseudo-pragmas, which are variations of the namespace pragma. The name of file cannot end in <code>.unreg</code>. <pre>#pragma namespace("__create") #pragma namespace("__delete") #pragma namespace("__modify")</pre> These three pragmas are used specify if the elements following the pragmas should be created, deleted, or modified by the CIM object manager. The <code>__delete</code> pragma can currently only be applied for a <code>mofreg -u</code> command. </td></tr></table>	<code>-r tag file</code>	The <i>file</i> argument is the actual MOF registration file. Its form is identical to the MOF syntax as defined by the Distributed Management Task Force (DMTF). The only difference is the addition of the following 3 new pseudo-pragmas, which are variations of the namespace pragma. The name of file cannot end in <code>.unreg</code>. <pre>#pragma namespace("__create") #pragma namespace("__delete") #pragma namespace("__modify")</pre> These three pragmas are used specify if the elements following the pragmas should be created, deleted, or modified by the CIM object manager. The <code>__delete</code> pragma can currently only be applied for a <code>mofreg -u</code> command.
<code>-r tag file</code>	The <i>file</i> argument is the actual MOF registration file. Its form is identical to the MOF syntax as defined by the Distributed Management Task Force (DMTF). The only difference is the addition of the following 3 new pseudo-pragmas, which are variations of the namespace pragma. The name of file cannot end in <code>.unreg</code>. <pre>#pragma namespace("__create") #pragma namespace("__delete") #pragma namespace("__modify")</pre> These three pragmas are used specify if the elements following the pragmas should be created, deleted, or modified by the CIM object manager. The <code>__delete</code> pragma can currently only be applied for a <code>mofreg -u</code> command.		

The *tag* argument is a unique string that specifies the identity of the registry action. This tag can be set to the package name or the patch number if the mofreg script is being invoked through packages/patches, though any tag can be specified.

Errors and warnings that are encountered when the CIM object manager handles the mofreg script are logged. Processing of the mofreg script stops at the first error. Specific warnings include:

Element already defined - the element already exists and cannot be created.

Element not found - the element does not exist and cannot be modified.

The error conditions are:

Key modification - A class cannot be modified if its keys are being changed.

Other mod compilation errors.

- s Forces the CIM object manager to immediately process outstanding registry requests, instead of at the next restart. This currently requires Java.
- u *tag* [*file*] Undoes the operations performed during mof registry.

The *tag* argument must correspond to the value set during the original mofreg invocation. If no mofreg was done with the original *tag*, the command does not succeed.

If required, an *unreg* file can be specified. If no *unreg* file is specified, the CIM object manager automatically undoes the actions of the registry. Any class created by the registry process is removed and any classes modified by the registry revert to the old state.

The mofreg command does not take care of cases where packages and patches make conflicting changes to classes. This should be taken care of by the standard patch and package conflict resolution.

退出状态

The following exit values are returned:

- 0 Successful completion.
- 1 An error occurred. The reason for error is displayed.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWwbco

另请参见 [init.wbem\(1M\)](#), [mofcomp\(1M\)](#), [wbemadmin\(1M\)](#), [wbemlogviewer\(1M\)](#), [attributes\(5\)](#), [wbem\(5\)](#)

引用名	monitor – SPARC system PROM monitor
用法概要	STOP–A BREAK initial system power-on exit from a client program, e.g., the Operating System
描述	The CPU board of a workstation contains one or more EPROMs or EEPROMs. The program which executes from the PROMs is referred to as “the monitor”. Among other things, the monitor performs system initialization at power-on and provides a user interface.
Monitor Prompt	The monitor of earlier workstations was known as the SunMON monitor and displayed the > for its prompt. See the SunMON MONITOR USAGE section for further details. Existing workstations use a monitor which is known as the OpenBoot monitor. The OpenBoot monitor typically displays ok as its prompt, but it may also display the > prompt under certain circumstances. If the 'auto-boot?' NVRAM parameter is set to 'false' when the workstation is powered on, the system does not attempt to boot and the monitor issues its prompt. If 'auto-boot' is set to 'true', the system initiates the boot sequence. The boot sequence can be aborted by simultaneously pressing two keys on the system's keyboard: L1 and A (on older keyboards), or Stop and A (on newer keyboards). Either a lower case a or an upper case A works for the keyboard abort sequence. If a console has been attached by way of one of the system's serial ports then the abort sequence can be accomplished by sending a BREAK. See tip(1). When the NVRAM 'security-mode' parameter has been turned on, or when the value of the 'sunmon-compat?' parameter is true, then the OpenBoot monitor displays the message: Type b (boot), c (continue), or n (new command mode) and the > prompt appears.
Openboot Prom Usage	Some of the more useful commands that can be issued from OpenBoot's ok prompt are described here. Refer to the 《OpenBoot 2.x Command Reference Manual》 book for a complete list of commands.
Help	Help for various functional areas of the OpenBoot monitor can be obtained by typing help. The help listing provides a number of other key words which can then be used in the help command to provide further details.
NVRAM Parameters	Each workstation contains one or more NVRAM devices which contains unique system ID information, as well as a set of user-configurable parameters. The NVRAM parameters allow the user a certain level of flexibility in configuring the system to act in a given manner under a specific set of circumstances. See eeprom(1M) for a description of the parameters and information regarding setting the parameters from the OS level.

The following commands can be used at the OpenBoot monitor to access the NVRAM parameters.

<code>printenv</code>	Used to list the NVRAM parameters, along with their default values and current values.
<code>setenv <i>pn pv</i></code>	Used to set or modify a parameter. The <i>pn</i> represents the parameter name, and <i>pv</i> represents the parameter value.
<code>set-default <i>pn</i></code>	Used to set an individual parameter back to its default value.
<code>set-defaults</code>	Used to reset all parameters to their default values. (Note that 'set-defaults' only affects parameters that have assigned default values.)

Security Parameters Newer OpenBoot monitors contain user interfaces that support the storage and listing of keys for later use by client programs.

<code>list-security-keys</code>	Lists the names of keys currently stored on a machine.
<code>set-security-key <i>keyname</i> [<i>keydata</i>]</code>	Stores key data <i>keydata</i> in a key named <i>keyname</i> . Actual key data can be up to 32 bytes in length. The maximum length of <i>keyname</i> is 64 bytes, which allows for the hex-formatted ASCII used to present the key data. If <i>keydata</i> is not present, <i>keyname</i> and its corresponding data is deleted.

Hardware Checks and Diagnostics The following commands are available for testing or checking the system's hardware. If the 'diag-switch?' NVRAM parameter is set to true when the system is powered on, then a Power-On Self Test (POST) diagnostic is run, if present, sending its results messages to the system's serial port A. Not all of the commands shown are available on all workstations.

<code>test-all</code>	Run the diagnostic tests on each device which has provided a self-test.
<code>test /memory</code>	Run the main memory tests. If the NVRAM parameter 'diag-switch?' is set to true, then all of main memory is tested. If the parameter is false then only the amount of memory specified in the 'selftest-#megs' NVRAM parameter is tested.
<code>test net</code>	Test the network connection for the on-board network controller.
<code>watch-net</code>	Monitor the network attached to the on-board net controller.
<code>watch-net-all</code>	Monitor the network attached to the on-board net controller, as well as the network controllers installed in SBus slots.
<code>watch-clock</code>	Test the system's clock function.

System Information	The following commands are available for displaying information about the system. Not all commands are available on all workstations. <table> <tr> <td>banner</td><td>Display the power-on banner.</td></tr> <tr> <td>.enet-addr</td><td>Display the system's Ethernet address.</td></tr> <tr> <td>.idprom</td><td>Display the formatted contents of the IDPROM.</td></tr> <tr> <td>module-info</td><td>Display information about the system's processor(s).</td></tr> <tr> <td>probe-scsi</td><td>Identify the devices attached to the on-board SCSI controller.</td></tr> <tr> <td>probe-scsi-all</td><td>Identify the devices attached to the on-board SCSI controller as well as those devices which are attached to SBus SCSI controllers.</td></tr> <tr> <td>show-disks</td><td>Display a list of the device paths for installed SCSI disk controllers.</td></tr> <tr> <td>show-displays</td><td>Display a list of the device paths for installed display devices.</td></tr> <tr> <td>show-nets</td><td>Display a list of the device paths for installed Ethernet controllers.</td></tr> <tr> <td>show-sbus</td><td>Display list of installed SBus devices.</td></tr> <tr> <td>show-tapes</td><td>Display a list of the device paths for installed SCSI tape controllers.</td></tr> <tr> <td>show-ttys</td><td>Display a list of the device paths for tty devices.</td></tr> <tr> <td>.traps</td><td>Display a list of the SPARC trap types.</td></tr> <tr> <td>.version</td><td>Display the version and date of the OpenBoot PROM.</td></tr> </table>	banner	Display the power-on banner.	.enet-addr	Display the system's Ethernet address.	.idprom	Display the formatted contents of the IDPROM.	module-info	Display information about the system's processor(s).	probe-scsi	Identify the devices attached to the on-board SCSI controller.	probe-scsi-all	Identify the devices attached to the on-board SCSI controller as well as those devices which are attached to SBus SCSI controllers.	show-disks	Display a list of the device paths for installed SCSI disk controllers.	show-displays	Display a list of the device paths for installed display devices.	show-nets	Display a list of the device paths for installed Ethernet controllers.	show-sbus	Display list of installed SBus devices.	show-tapes	Display a list of the device paths for installed SCSI tape controllers.	show-ttys	Display a list of the device paths for tty devices.	.traps	Display a list of the SPARC trap types.	.version	Display the version and date of the OpenBoot PROM.
banner	Display the power-on banner.																												
.enet-addr	Display the system's Ethernet address.																												
.idprom	Display the formatted contents of the IDPROM.																												
module-info	Display information about the system's processor(s).																												
probe-scsi	Identify the devices attached to the on-board SCSI controller.																												
probe-scsi-all	Identify the devices attached to the on-board SCSI controller as well as those devices which are attached to SBus SCSI controllers.																												
show-disks	Display a list of the device paths for installed SCSI disk controllers.																												
show-displays	Display a list of the device paths for installed display devices.																												
show-nets	Display a list of the device paths for installed Ethernet controllers.																												
show-sbus	Display list of installed SBus devices.																												
show-tapes	Display a list of the device paths for installed SCSI tape controllers.																												
show-ttys	Display a list of the device paths for tty devices.																												
.traps	Display a list of the SPARC trap types.																												
.version	Display the version and date of the OpenBoot PROM.																												
Emergency Commands	These commands must be typed from the keyboard, they do not work from a console which is attached by way of the serial ports. With the exception of the Stop-A command, these commands are issued by pressing and holding down the indicated keys on the keyboard immediately after the system has been powered on. The keys must be held down until the monitor has checked their status. The Stop-A command can be issued at any time after the console display begins, and the keys do not need to be held down once they've been pressed. The Stop-D, Stop-F and Stop-N commands are not allowed when one of the security modes has been set. Not all commands are available on all workstations. <table> <tr> <td>Stop (L1)</td><td>Bypass the Power-On Self Test (POST). This is only effective if the system has been placed into the diagnostic mode.</td></tr> <tr> <td>Stop-A (L1-A)</td><td>Abort the current operation and return to the monitor's default prompt.</td></tr> <tr> <td>Stop-D (L1-D)</td><td>Set the system's 'diag-switch?' NVRAM parameter to 'true', which places the system in diagnostic mode. POST diagnostics, if present, are run, and the messages are displayed by way of the system's serial port A.</td></tr> <tr> <td>Stop-F (L1-F)</td><td>Enter the OpenBoot monitor before the monitor has probed the system for devices. Issue the 'fexit' command to continue with system initialization.</td></tr> </table>	Stop (L1)	Bypass the Power-On Self Test (POST). This is only effective if the system has been placed into the diagnostic mode.	Stop-A (L1-A)	Abort the current operation and return to the monitor's default prompt.	Stop-D (L1-D)	Set the system's 'diag-switch?' NVRAM parameter to 'true', which places the system in diagnostic mode. POST diagnostics, if present, are run, and the messages are displayed by way of the system's serial port A.	Stop-F (L1-F)	Enter the OpenBoot monitor before the monitor has probed the system for devices. Issue the 'fexit' command to continue with system initialization.																				
Stop (L1)	Bypass the Power-On Self Test (POST). This is only effective if the system has been placed into the diagnostic mode.																												
Stop-A (L1-A)	Abort the current operation and return to the monitor's default prompt.																												
Stop-D (L1-D)	Set the system's 'diag-switch?' NVRAM parameter to 'true', which places the system in diagnostic mode. POST diagnostics, if present, are run, and the messages are displayed by way of the system's serial port A.																												
Stop-F (L1-F)	Enter the OpenBoot monitor before the monitor has probed the system for devices. Issue the 'fexit' command to continue with system initialization.																												

Stop-N (L1-N) Causes the NVRAM parameters to be reset to their default values. Note that not all parameters have default values.

Line Editor Commands The following commands can be used while the monitor is displaying the ok prompt. Not all of these editing commands are available on all workstations.

CTRL-A	Place the cursor at the start of line.
CTRL-B	Move the cursor backward one character.
ESC-B	Move the cursor backward one word.
CTRL-D	Erase the character that the cursor is currently highlighting.
ESC-D	Erase the portion of word from the cursor's present position to the end of the word.
CTRL-E	Place the cursor at the end of line.
CTRL-F	Move the cursor forward one character.
ESC-F	Move the cursor forward one word.
CTRL-H	Erase the character preceding the cursor (also use Delete or Back Space)
ESC-H	Erase the portion of the word which precedes the cursor (use also CTRL-W)
CTRL-K	Erase from the cursor's present position to the end of the line.
CTRL-L	Show the command history list.
CTRL-N	Recall the next command from the command history list
CTRL-P	Recall a previous command from the command history list.
CTRL-Q	Quote the next character (used to type a control character).
CTRL-R	Retype the current line.
CTRL-U	Erase from the cursor's present position to the beginning of the line.
CTRL-Y	Insert the contents of the memory buffer into the line, in front (to the left) of the cursor.

nvrnrc The nvrnrc is an area of the system's NVRAM where users may store Forth programs. The programs which are stored in the nvrnrc are executed each time the system is reset, provided that the 'use-nvrnrc?' NVRAM parameter has been set to 'true'. Refer to the [《OpenBoot 2.x Command Reference Manual》](#) book for information on how to edit and use the nvrnrc.

Restricted Monitor The command 'old-mode' is used to move OpenBoot into a restricted monitor mode, causing the > prompt to be displayed. Only three commands are allowed while in the restricted monitor; the 'go' command (to resume a program which was interrupted with the

Stop-A command), the 'n' command (to return to the normal OpenBoot monitor), and boot commands. The restricted monitor's boot commands approximate the older SunMON monitor's boot command syntax. If a 'security-mode' has been turned on then the restricted monitor becomes the default monitor environment. The restricted monitor may also become the default environment if the 'sunmon-compat?' NVRAM parameter is set to true. Not all workstations have the 'sunmon-compat?' parameter.

SunMON Prom
Usage

The following commands are available systems with older SunMON-based PROM:

- +|-
Increment or decrement the current address and display the contents of the new location.
- ^C *source destination n*
(caret-C) Copy, byte-by-byte, a block of length *n* from the source address to the *destination* address.
- ^I *program*
(caret-I) Display the compilation date and location of *program*.
- ^T *virtual_address*
(caret-T) Display the physical address to which *virtual_address* is mapped.
- b [!] [*device* [(*c,u,p*)]] [*pathname*] [*arguments_list*]
b[?]
Reset appropriate parts of the system and bootstrap a program. A '!' (preceding the *device* argument) prevents the system reset from occurring. Programs can be loaded from various devices (such as a disk, tape, or Ethernet). 'b' with no arguments causes a default boot, either from a disk, or from an Ethernet controller. 'b?' displays all boot devices and their *devices*.

<i>device</i>	one of
le	Lance Ethernet
ie	Intel Ethernet
sd	SCSI disk, CDROM
st	SCSI 1/4- or 1/2-inch tape
fd	Diskette
id	IPI disk
mt	Tape Master 9-track 1/2-inch tape
xd	Xylogics 7053 disk
xt	Xylogics 1/2-inch tape
xy	Xylogics 440/450 disk

- c* A controller number (0 if only one controller),
- u* A unit number (0 if only one driver), and
- p* A partition.
- pathname* A pathname for a program such as /stand/diag.
- arguments_list* A list of up to seven arguments to pass to the program being booted.
- c [*virtual_address*]
Resume execution of a program. When given, *virtual_address* is the address at which execution resumes. The default is the current PC. Registers are restored to the values shown by the d, and r commands.
- d [*window_number*]
Display (dump) the state of the processor. The processor state is observable only after:
- An unexpected trap was encountered.
 - A user program dropped into the monitor (by calling *abortent*).
 - The user manually entered the monitor by typing L1-A or BREAK.
- The display consists of the following:
- The special registers: PSR, PC, nPC, TBR, WIM, and Y
 - Eight global registers
 - 24 window registers (8 *in*, 8 *local*, and 8 *out*), corresponding to one of the 7 available windows. If a Floating-Point Unit is on board, its status register along with 32 floating-point registers are also shown.
- window_number* Display the indicated *window_number*, which can be any value between 0 and 6, inclusive. If no window is specified and the PSR's current window pointer contains a valid window number, registers from the window that was active just prior to entry into the monitor are displayed. Otherwise, registers from window 0 are displayed.
- e [*virtual_address*] [*action*] . . .
Open the 16-bit word at *virtual_address* (default zero). The address is interpreted in the address space defined by the s command. See the a command for a description of *action*.
- f *virtual_address1* *virtual_address2* *pattern* [*size*]
Fill the bytes, words, or long words from *virtual_address1* (lower) to *virtual_address2* (higher) with the constant, *pattern*. The *size* argument can take one of the following values:
- b byte format (the default)
 - w word format
 - l long word format

For example, the following command fills the address block from 0x1000 to 0x2000 with the word pattern, 0xABCD:

```
f 1000 2000 ABCD W
```

g [*vector*] [*argument*]

g [*virtual_address*] [*argument*]

Goto (jump to) a predetermined or default routine (first form), or to a user-specified routine (second form). The value of *argument* is passed to the routine. If the *vector* or *virtual_address* argument is omitted, the value in the PC is used as the address to jump to.

To set up a predetermined routine to jump to, a user program must, prior to executing the monitor's **g** command, set the variable `*romp->v_vector_cmd` to be equal to the virtual address of the desired routine. Predetermined routines need not necessarily return control to the monitor.

The default routine, defined by the monitor, prints the user-supplied *vector* according to the format supplied in *argument*. This format can be one of:

%x hexadecimal

%d decimal

g0

Force a panic and produce a crash dump when the monitor is running as a result of the system being interrupted,

g4

(Sun-4 systems only) Force a kernel stack trace when the monitor is running as a result of the system being interrupted,

h

Display the help menu for monitor commands and their descriptions. To return to the monitor's basic command level, press ESCAPE or **q** before pressing RETURN.

i [*cache_data_offset*] [*action*] . . .

Modify cache data RAM command. Display and/or modify one or more of the cache data addresses. See the **a** command for a description of *action*.

j [*cache_tag_offset*] [*action*] . . .

Modify cache tag RAM command. Display and/or modify the contents of one or more of the cache tag addresses. See the **a** command for a description of *action*.

k [*reset_level*]

Reset the system, where *reset_level* is:

0 Reset VMEbus, interrupt registers, video monitor (Sun-4 systems). This is the default.

1 Software reset.

- 2 Power-on reset. Resets and clears the memory. Runs the EPROM-based diagnostic self test, which can take several minutes, depending upon how much memory is being tested.

kb

Display the system banner.

l [*virtual_address*] [*action*] . . .

Open the long word (32 bit) at memory address *virtual_address* (default zero). The address is interpreted in the address space defined by the s command (below). See the a command for a description of *action*.

m [*virtual_address*] [*action*] . . .

Open the segment map entry that maps *virtual_address* (default zero). The address is interpreted in the address space defined by the s command. See the a command for a description of *action*.

ne

ni

Disable, enable, or invalidate the cache, respectively.

o [*virtual_address*] [*action*] . . .

Open the byte location specified by *virtual_address* (default zero). The address is interpreted in the address space defined by the s command. See the a command for a description of *action*.

p [*virtual_address*] [*action*] . . .

Open the page map entry that maps *virtual_address* (default zero) in the address space defined by the s command. See the a command for a description of *action*.

q [*eprom_offset*] [*action*] . . .

Open the EEPROM *eprom_offset* (default zero) in the EEPROM address space. All addresses are referenced from the beginning or base of the EEPROM in physical address space, and a limit check is performed to insure that no address beyond the EEPROM physical space is accessed. This command is used to display or modify configuration parameters, such as: the amount of memory to test during self test, whether to display a standard or custom banner, if a serial port (A or B) is to be the system console, etc. See the a command for a description of *action*.

r [*register_number*]

r [*register_type*]

r [*w window_number*]

Display and/or modify one or more of the IU or FPU registers. A hexadecimal *register_number* can be one of:

0x00–0x0f window(0,i0)—window(0,i7), window(0,i0)—window(0,i7)

0x16–0x1f window(1,i0)—window(1,i7), window(1,i0)—window(1,i7)

0x20–0x2f	window(2,i0)–window(2,i7), window(2,i0)—window(2,i7)
0x30–0x3f	window(3,i0)–window(3,i7), window(3,i0)—window(3,i7)
0x40–0x4f	window(4,i0)–window(4,i7), window(4,i0)—window(4,i7)
0x50–0x5f	window(5,i0)–window(5,i7), window(5,i0)—window(5,i7)
0x60–0x6f	window(6,i0)–window(6,i7), window(6,i0)—window(6,i7)
0x70–0x77	g0, g1, g2, g3, g4, g5, g6, g7
0x78–0x7d	PSR, PC, nPC, WIM, TBR, Y.
0x7e–0x9e	FSR, f0–f31

Register numbers can only be displayed after an unexpected trap, a user program has entered the monitor using the *abortent* function, or the user has entered the monitor by manually typing L1–A or BREAK.

If a *register_type* is given, the first register of the indicated type is displayed. *register_type* can be one of:

f	floating-point
g	global
s	special

If *w* and a *window_number* (0–6) are given, the first *in*-register within the indicated window is displayed. If *window_number* is omitted, the window that was active just prior to entering the monitor is used. If the PSR's current window pointer is invalid, window 0 is used.

s [*asi*]

Set or display the Address Space Identifier. With no argument, s displays the current Address Space Identifier. The *asi* value can be one of:

0x2	control space
0x3	segment table
0x4	Page table
0x8	user instruction
0x9	supervisor instruction
0xa	user data
0xb	supervisor data
0xc	flush segment

0xd flush page

0xe flush context

0xf cache data

u [echo]

u [port] [options] [baud_rate]

u [u] [virtual_address]

With no arguments, display the current I/O device characteristics including: current input device, current output device, baud rates for serial ports A and B, an input-to-output echo indicator, and virtual addresses of mapped UART devices. With arguments, set or configure the current I/O device. With the u argument (uu. . .), set the I/O device to be the *virtual_address* of a UART device currently mapped.

echo Can be either e to enable input to be echoed to the output device, or ne, to indicate that input is not echoed.

port Assign the indicated *port* to be the current I/O device. *port* can be one of:

a serial port A

b serial port B

k the workstation keyboard

s the workstation screen

baud_rate Any legal baud rate.

options can be any combination of:

i input

o output

u UART

e echo input to output

ne do not echo input

r reset indicated serial port (a and b ports only)

If either a or b is supplied, and no *options* are given, the serial port is assigned for both input and output. If k is supplied with no options, it is assigned for input only. If s is supplied with no options, it is assigned for output only.

v *virtual_address1* *virtual_address2* [size]

Display the contents of *virtual_address1* (lower) *virtual_address2* (higher) in the format specified by size:

b byte format (the default)

- w word format
- l long word format

Enter return to pause for viewing; enter another return character to resume the display. To terminate the display at any time, press the space bar.

For example, the following command displays the contents of virtual address space from address 0x1000 to 0x2000 in word format:

```
v 1000 2000 W
```

- w [*virtual_address*] [*argument*]
Set the execution vector to a predetermined or default routine. Pass *virtual_address* and *argument* to that routine.

To set up a predetermined routine to jump to, a user program must, prior to executing the monitor's w command, set the variable *romp->v_vector_cmd to be equal to the virtual address of the desired routine. Predetermined routines need not necessarily return control to the monitor.

The default routine, defined by the monitor, prints the user-supplied *vector* according to the format supplied in *argument*. This format can be one of:

- %x hexadecimal
- %d decimal

- x
Display a menu of extended tests. These diagnostics permit additional testing of such things as the I/O port connectors, video memory, workstation memory and keyboard, and boot device paths.

- y c *context_number*
Display context number.
- y p | s *context_number virtual_address*
Flush the indicated context, context page, or context segment.
 - c flush context *context_number*
 - p flush the page beginning at *virtual_address* within context *context_number*
 - s flush the segment beginning at *virtual_address* within context *context_number*

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	SPARC

另请参见

[tip\(1\)](#), [boot\(1M\)](#), [eeprom\(1M\)](#), [attributes\(5\)](#)

《OpenBoot 2.x Command Reference Manual》

引用名

mount, umount – 挂载或卸载文件系统和远程资源

用法概要

```
mount [-p | -v]

mount [-F FSType] [generic_options] [-o specific_options]
      [-O special | mount_point]

mount [-F FSType] [generic_options] [-o specific_options]
      [-O] special mount_point

mount -a [-F FSType] [-V] [current_options]
      [-o specific_options] [mount_point]...

umount [-f] [-V] [-o specific_options] special | mount_point

umount -a [-f] [-V] [-o specific_options] [mount_point]...
```

描述

mount 将文件系统附加到文件系统分层结构中的 *mount_point*，这是一个目录的路径名。如果在执行 mount 操作之前，*mount_point* 中具有任何内容，则这些内容将被隐藏直至该文件系统被卸载。

umount 卸载当前已挂载的文件系统，该文件系统可以指定为 *mount_point* 或 *special*（文件系统所驻留的设备）。

可以通过检查已挂载文件系统信息文件来查看当前挂载的文件系统的表。这是由通常挂载在 */etc/mnttab* 上的一个文件系统提供的。[mnttab\(4\)](#) 中描述了已挂载的文件系统的信息。挂载一个文件系统会向挂载表添加一个条目，umount 会从该表删除一个条目。

当同时使用 *special* 和 *mount_point* 参数以及 -F 选项调用时，mount 会验证除 *special* 之外的所有参数并调用相应的特定于 *FSType* 的 mount 模块。在不使用参数调用时，mount 将列出挂载表 (*/etc/mnttab*) 中记录的所有已挂载的文件系统。当使用一个部分参数列表调用时（只指定 *special* 和 *mount_point* 中的一个，或同时指定 *special* 和 *mount_point* 但不指定 *FSType*），mount 将搜索 */etc/vfstab* 来查找将提供缺少的参数的条目。如果没有找到相应的条目，并且 *special* 参数以 / 开头，则将使用 */etc/default/fs* 中指定的缺省本地文件系统类型。否则，将使用缺省的远程文件系统类型。缺省的远程文件系统类型是由 */etc/dfs/fstypes* 文件中的第一个条目决定的。在填充缺少的参数后，mount 将调用特定于 *FSType* 的 mount 模块。

如果文件系统类型支持，可以直接将一个文件挂载为文件系统：通过将文件的全路径指定为 *special* 参数。在这种情况下，*nosuid* 选项是必需的。如果没有针对此类回送文件挂载提供特定文件系统支持，您仍然可以使用 [lofiadm\(1M\)](#) 来挂载文件系统映像。在此情况下，不需要特殊选项。

只有具有足够特权（至少为 *PRIV_SYS_MOUNT*）的用户可以使用 mount 和 umount 来挂载或卸载文件系统。不过，任何用户都可以使用 mount 来列出已挂载的文件系统和资源。

选项

-F *FSType*

用来指定要对其进行操作的 *FSType*。*FSType* 必须予以指定，或者可以从 `/etc/vfstab` 或通过查询 `/etc/default/fs` 或 `/etc/dfs/fstypes` 来确定。

-a [*mount_points...*]

尽可能以并行方式执行挂载或卸载操作。

如果未指定挂载点，`mount` 将挂载其 `/etc/vfstab` "mount at boot" 字段为 `yes` 的所有文件系统。如果指定了挂载点，则会忽略 `/etc/vfstab` "mount at boot" 字段。

如果指定了挂载点，则 `umount` 将仅**卸载**那些挂载点。如果未指定挂载点，则 `umount` 将尝试**卸载** `/etc/mnttab` 中的所有文件系统，但下面列出的这些系统必需文件系统除外：`/`、`/usr`、`/var`、`/var/adm`、`/var/run`、`/proc`、`/dev/fd` 和 `/tmp`。

-f

强制卸载某个文件系统。

未指定此选项时，如果文件系统上的某个文件正处于繁忙状态，则 `umount` 不允许卸载该文件系统。使用此选项可能会导致打开的文件的数据丢失；访问这些文件的程序在文件系统被卸载后将收到一个错误 (EIO)。

-p

以 `/etc/vfstab` 格式输出已挂载的文件系统的列表。必须是唯一指定的选项。请参见 BUG。

-v

以详细格式输出已挂载的文件系统的列表。必须是唯一指定的选项。

-V

回显完整命令行，但不执行命令。`umount` 使用用户提供的选项和参数并将由 `/etc/mnttab` 派生的信息添加到这些选项和参数，从而生成一个命令行。应当使用此选项来确认和验证命令行。

generic_options

特定于 *FSType* 的大多数命令模块通常会支持的选项。有下列选项可用：

-m

挂载文件系统但不在 `/etc/mnttab` 中创建条目。

-g

全局性地挂载文件系统。在群集系统上，这将在群集的所有节点上全局性地挂载文件系统。在非群集系统上，这不起作用。

-o

在以逗号分隔的（没有空格）子选项和关键字-属性对的列表中，指定将由该命令的特定于 *FSType* 的模块解释的特定于 *FSType* 的选项。（请参见 [mount_ufs\(1M\)](#)。）当您将 `-o` 用于在 `/etc/vfstab` 中有相应条目的文件系统时，`mount` 将忽略在 `/etc/vfstab` 中为该文件系统输入的选项。

所支持的选项如下：

devices | nodevices

允许或禁止打开设备特定文件。缺省值为 **devices**。

如果将 **nosuid** 与 **devices** 结合使用，则其行为等同于 **nosuid** 的行为。

exec | noexec

允许或禁止执行文件系统中的程序。允许或禁止在使用 **PROT_EXEC** 的情况下对文件系统内的文件执行 **mmap(2)**。缺省值为 **exec**。

loop

将被忽略以便兼容。

nbmand | nonbmand

在此文件系统上允许或禁止非阻塞强制性锁定语义。缺省情况下会禁止非阻塞强制性锁定。

如果文件系统是使用 **nbmand** 选项挂载的，则应用程序可以使用 **fcntl(2)** 接口在文件上放置非阻塞强制性锁定并且系统将实施那些语义。如果启用此选项，它会导致遵循标准的应用程序看到意外的错误。

为避免在系统文件上获取强制性锁定这一可能性，请不要对下列文件系统使用 **nbmand** 选项：

```
/
/usr
/etc
/var
/proc
/dev
/devices
/system/contract
/system/object
/etc/mnttab
/etc/dfs/sharetab
```

不要使用 **remount** 选项来更改文件系统的 **nbmand** 配置。**nbmand** 选项与全局选项是互斥的。请参见 **-g**。

ro | rw

指定只读或读写。缺省值为 **rw**。

setuid | nosetuid

允许或禁止 **setuid** 或 **setgid** 执行。缺省值为 **setuid**。

如果您同时指定了 **setuid** 与 **nosuid**，则其行为与 **nosuid** 的行为相同。

nosuid 等效于 **nosetuid** 和 **nodevices**。当 **suid** 或 **nosuid** 与 **setuid** 或 **nosetuid** 以及 **devices** 或 **nodevices** 结合使用时，这会实施最具限制性的选项。

如果文件系统是使用 `root=` 选项通过 NFS 共享的，则强烈建议使用此选项。如果不使用此选项，则 NFS 客户机将能够向服务器添加 `setuid` 程序或创建可能会带来安全漏洞的设备。

`suid | nosuid`

允许或禁止 `setuid` 或 `setgid` 执行。缺省值为 `suid`。此选项还可允许或禁止打开在文件系统内出现的任何设备特殊条目。

`nosuid` 等效于 `nosetuid` 和 `nodevices`。当 `suid` 或 `nosuid` 与 `setuid` 或 `nosetuid` 以及 `devices` 或 `nodevices` 结合使用时，这会实施最具限制性的选项。

如果文件系统是使用 `root=` 选项通过 NFS 共享的，则强烈建议使用此选项；这是因为，如果不使用此选项，则 NFS 客户机将能够向服务器添加 `setuid` 程序或者创建可能会带来安全漏洞的设备。

`rstchown | norstchown`

允许或禁止受限的 `chown`。如果文件系统是使用 `rstchown` 挂载的，则不允许文件的所有者更改文件的所有者 ID。如果文件系统是使用 `norstchown` 挂载的，则用户可以更改他们拥有的文件的所有权。只有超级用户或具有相应特权的用户可以任意更改所有者 ID。

`-O`

覆盖挂载。允许文件系统挂载到现有的挂载点之上，使得底层文件系统不可访问。如果在未设置此标志的情况下尝试在已存在的挂载点上进行挂载，则挂载将失败，并生成错误 "device busy"（设备繁忙）。

`-r`

将文件系统挂载为只读的。

示例

示例 1 直接挂载和卸载 DVD 映像

下面的命令挂载和卸载 DVD 映像。

```
# mount -F hsfs /images/solaris.iso /mnt/solaris-image
# umount /mnt/solaris-image
```

用法

有关遇到大于或等于 2 GB (2^{31} 字节) 的文件时 `mount` 和 `umount` 的行为，请参见 [largefile\(5\)](#) 中的描述。

文件

`/etc/mnttab`

其中列出了已挂载的文件系统的表。

`/etc/default/fs`

缺省的本地文件系统类型。可以在 `/etc/default/fs` 中为以下标志设置缺省值。例如：`LOCAL=ufs`

`LOCAL :`

未指定 `FSType` 时用于命令的缺省分区。

`/etc/vfstab`

每个文件系统的缺省参数的列表。

属性类型	属性值
可用性	system/core-os

引用名	mountall, umountall – mount, unmount multiple file systems
用法概要	<pre>mountall [-F FSType] [-l -r] [file_system_table] umountall [-k] [-s] [-F FSType] [-l -r] [-n] [-Z] umountall [-k] [-s] [-h host] [-n] [-Z]</pre>
描述	mountall is used to mount file systems specified in a file system table. The file system table must be in vfstab(4) format. If no <i>file_system_table</i> is specified, <i>/etc/vfstab</i> is used. If – is specified as <i>file_system_table</i>, mountall reads the file system table from the standard input. mountall mounts only those file systems with the mount at boot field set to yes in the <i>file_system_table</i>. For each file system in the file system table, the following logic is executed: if there exists a file <i>/usr/lib/fs/FSType/fsckall</i>, where <i>FSType</i> is the type of the file system, save that file system in a list to be passed later, and all at once, as arguments to the <i>/usr/lib/fs/FSType/fsckall</i> script. The <i>/usr/lib/fs/FSType/fsckall</i> script checks all of the file systems in its argument list to determine whether they can be safely mounted. If no <i>/usr/lib/fs/FSType/fsckall</i> script exists for the <i>FSType</i> of the file system, the file system is individually checked using fsck(1M). If the file system does not appear mountable, it is fixed using <i>fsck</i> before the mount is attempted. File systems with a – entry in the <i>fsckdev</i> field are mounted without first being checked. umountall causes all mounted file systems in the current zone except <i>root</i>, <i>/usr</i>, <i>/var</i>, <i>/var/adm</i>, <i>/var/run</i>, <i>/proc</i>, and <i>/dev/fd</i> to be unmounted. If the <i>FSType</i> is specified, mountall and umountall limit their actions to the <i>FSType</i> specified. There is no guarantee that umountall unmounts <i>busy</i> file systems, even if the <i>-k</i> option is specified.
选项	The following options are supported: -F Specify the <i>FSType</i> of the file system to be mounted or unmounted. -h <i>host</i> Unmount all file systems listed in <i>/etc/mnttab</i> that are remote-mounted from <i>host</i>. -k Use the <i>fuser -k mount-point</i> command. See the fuser(1M) for details. The <i>-k</i> option sends the SIGKILL signal to each process using the file. As this option spawns kills for each process, the kill messages might not show up immediately. There is no guarantee that umountall unmounts <i>busy</i> file systems, even if the <i>-k</i> option is specified. -l Limit the action to local file systems. -n List the actions that would be performed for the specified options, but do not actually execute these actions. Repeating the command without the <i>-n</i> option executes the listed actions, assuming that the <i>/etc/mnttab</i> file has not changed in the interval prior to repeating the command. -r Limit the action to remote file system types.

- s Do not perform the umount operation in parallel.
- Z Apply the action(s) only to the file systems mounted in non-global zones. By default, `umountall` unmounts only file systems mounted in the current zone. Option -Z is ignored if used in a non-global zone.

文件

`/etc/mnttab`
Mounted file system table

`/etc/vfstab`
Table of file system defaults

`/usr/lib/fs/FSType/fsckall`
Script called by `mountall` to perform the file system check of all file systems of type *FSType*

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	system/core-os
Interface Stability	Committed
Output Stability	Uncommitted

另请参见

[fsck\(1M\)](#), [fuser\(1M\)](#), [mount\(1M\)](#), [mnttab\(4\)](#), [vfstab\(4\)](#), [attributes\(5\)](#)

诊断

No messages are printed if the file systems are mountable and clean.

Error and warning messages come from [fsck\(1M\)](#) and [mount\(1M\)](#).

附注

At this time, NFS is the only remote file system supported by the -l, -r, and -h options.

引用名	mountd – server for NFS mount requests and NFS access checks
用法概要	<code>/usr/lib/nfs/mountd [-v] [-r]</code>
描述	mountd is an RPC server that answers requests for NFS access information and file system mount requests. It reads the file <code>/etc/dfs/sharetab</code> to determine which file systems are available for mounting by which remote machines. See sharetab(4). nfsd running on the local server will contact mountd the first time an NFS client tries to access the file system to determine whether the client should get read-write, read-only, or no access. This access can be dependent on the security mode used in the remotd procedure call from the client. See share_nfs(1M). The command also provides information as to what file systems are mounted by which clients. This information can be printed using the showmount(1M) command. The mountd daemon is automatically invoked by share(1M). Only super user can run the mountd daemon.
SMF Management	Since mountd must be running for nfsd to function properly, mountd is automatically started by the <code>svc:/network/nfs/server</code> service. Startup SMF parameters for mountd can be manipulated using the sharectl(1M) command. The currently supported parameters are: <pre>client_versmin=num client_versmax=num</pre> The NFS client uses only NFS versions in the range specified by these variables. Valid values or versions are: 2, 3, and 4. By default, these variables are unspecified (commented out) and the client's default minimum is Version 2. The default maximum is version 4. You can override this range on a per-mount basis by using the <code>-o vers=</code> option to mount_nfs(1M).
选项	The options shown below are supported for NVSv2/v3 clients. They are not supported for Solaris NFSv4 clients. - r Reject mount requests from clients. Clients that have file systems mounted will not be affected. - v Run the command in verbose mode. Each time mountd determines what access a client should get, it will log the result to the console, as well as how it got that result.
文件	<code>/etc/dfs/sharetab</code> shared file system table
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/file-system/nfs

另请参见

[mount_nfs\(1M\)](#), [nfsd\(1M\)](#), [share\(1M\)](#), [share_nfs\(1M\)](#), [sharectl\(1M\)](#), [showmount\(1M\)](#), [sharetab\(4\)](#), [attributes\(5\)](#)

附注

Some routines that compare hostnames use case-sensitive string comparisons; some do not. If an incoming request fails, verify that the case of the hostname in the file to be parsed matches the case of the hostname called for, and attempt the request again.

引用名 mount_hfs – mount hfs file systems

用法概要

```
mount -F hfs [generic_options]
      [-o FSType-specific_options] [-O ] special | mount_point

mount -F hfs [generic_options]
      [-o FSType-specific_options] [-O] special mount_point
```

描述

mount attaches an ISO 9660 filesystem (the High Sierra file system, hfs, is a draft predecessor to ISO 9660, so the name reflects the filesystem's history) to the file system hierarchy at the *mount_point*, which is the pathname of a directory. If *mount_point* has any contents prior to the mount operation, these are hidden until the file system is unmounted.

If mount is invoked with *special* or *mount_point* as the only arguments, mount will search /etc/vfstab to fill in the missing arguments, including the *FSType-specific_options*; see [mount\(1M\)](#) for more details.

The hfs file system supports direct mounting of files containing the file system as well as block devices. See [mount\(1M\)](#) and [lofiadm\(1M\)](#).

A file system conforming to ISO 9660 can contain extensions that allow it to overcome limitations of the original ISO 9660:1988 (version 1) standard. The following types of extensions are supported by hfs:

Rock Ridge extensions

This is the preferred type of extension as it allows file attributes, name length, and types equivalent to those on other UNIX-style filesystems. Example of supported features are device special files, POSIX permissions, symbolic links, and filenames of up to 255 bytes in length. Rock Ridge extensions also remove the ISO9660:1988 restriction on maximum nesting depth for directories (eight levels). hfs automatically detects the presence of Rock Ridge extensions and uses them, unless mount options are specified to disable the use of Rock Ridge or to use a different extension.

ISO9660:1999 (version 2) extensions

The first version of ISO9660, released in 1988, supported only uppercase ASCII filenames of no more than 31 characters in length. ISO9660 version 2, released in 1999, provides an extension that allows filenames of at least 207 bytes that can use UTF-8 characters and removes the limitation on the nesting depth for directories. Unlike Rock Ridge, it does not provide support for UNIX-style file types and file attributes. hfs automatically detects this extension and will use it for filename lookup if no Rock Ridge extensions are found on the media.

Joliet extensions

The Joliet extension was devised by Microsoft to allow Unicode (UCS-2) long filenames with CDROM-based media. It allows filename lengths of up to 110 Unicode characters and does not support UNIX-style file types and attributes. hfs falls back to using Joliet if such an extension is present and neither Rock Ridge nor ISO9660 version 2 extensions are found.

If filenames are longer than the 64 UCS-2 characters officially allowed by Microsoft (that is, 110 Unicode characters), they can translate to up to 330 UTF-8 octets. Filenames that translate to more than 255 UTF-8 octets will be truncated.

选项

generic_options

See [mount\(1M\)](#) for the list of supported options.

-o

Specify hfs file system specific options. If invalid options are specified, a warning message is printed and the invalid options are ignored. The following options are available:

global | *noglobal*

If *global* is specified and supported on the file system, and the system in question is part of a cluster, the file system will be globally visible on all nodes of the cluster. If *noglobal* is specified, the mount will not be globally visible. The default behavior is *noglobal*.

ro

Mount the file system read-only. This option is required.

rr | *nrr*

Enable (*rr*) or disable (*nrr*) the use of Rock Ridge. *rr* is the default and need not be specified. If you use *nrr* and Rock Ridge extensions are present in the file system, ignore them and search for other available extensions or fall back to plain ISO9660.

vers2 | *novers2*

Enable or disable the use of ISO9660 version 2 extensions. If *vers2* is specified and ISO9660 version 2 extensions are available, hfs will use ISO9660 version 2 even if the file system contains the preferred Rock Ridge extensions as well. If *novers2* is specified, it will fall back to using either Joliet extensions or plain ISO9660 even if ISO9660 version 2 extensions are available.

joliet | *nojoliet*

Enable or disable the use of Joliet extensions. If *joliet* is specified and Joliet extensions are available, hfs will use them even if the file system contains the preferred Rock Ridge and/or ISO9660 version 2 extensions. If *nojoliet* is specified, it will fall back to using plain ISO9660.

notraildot

File names on High Sierra file systems consist of a proper name and an extension separated by a '.' (dot) character. By default, the separating dot is always considered part of the file's name for all file access operations, even if there is no extension present. Specifying *notraildot* makes it optional to specify the trailing dot to access a file whose name lacks an extension.

Exceptions: This option is effective only on file systems for which Rock Ridge, ISO9660 version 2 or Joliet extensions are not active, either because they are not present on the CD-ROM, or they have been deliberately disabled via the *nrr*, *novers2* and *nojoliet* option. If either extension is active, hfs quietly ignores this option.

`nomaplower`
File names on High Sierra/ISO9660 CD-ROMs with no extensions present should be uppercase characters only. By default, `hfs` maps file names read from a non-Rock Ridge disk to all lowercase characters. `nomaplower` turns off this mapping. The exceptions for `notraildot` discussed above apply to `nomaplower`.

`-O`
Overlay mount. Allow the file system to be mounted over an existing mount point, making the underlying file system inaccessible. If a mount is attempted on a preexisting mount point without setting this flag, the mount will fail, producing the error: device busy.

示例
示例 1 Mounting and Unmounting a DVD Image Directly
The following commands mount and unmount a DVD image.

```
# mount -F hfs /images/solaris.iso /mnt/solaris-image
# umount /mnt/solaris-image
```

文件
`/etc/mnttab`
table of mounted file systems
`/etc/vfstab`
list of default parameters for each file system

属性
See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见
[lofiadm\(1M\)](#), [mount\(1M\)](#), [mountall\(1M\)](#), [mount\(2\)](#), [mnttab\(4\)](#), [vfstab\(4\)](#), [attributes\(5\)](#)

附注
If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on the directory to which the symbolic link refers, rather than on top of the symbolic link itself.

引用名

mount_nfs – mount remote NFS resources

用法概要

```
mount [-F nfs] [generic_options] [-o specific_options] [-O] resource
mount [-F nfs] [generic_options] [-o specific_options] [-O] mount_point
mount [-F nfs] [generic_options] [-o specific_options]
      [-O] resource mount_point
```

描述

The mount utility attaches a named *resource* to the file system hierarchy at the pathname location *mount_point*, which must already exist. If *mount_point* has any contents prior to the mount operation, the contents remain hidden until the *resource* is once again unmounted.

mount_nfs starts the [lockd\(1M\)](#) and [statd\(1M\)](#) daemons if they are not already running.

If the resource is listed in the /etc/vfstab file, the command line can specify either *resource* or *mount_point*, and mount consults /etc/vfstab for more information. If the -F option is omitted, mount takes the file system type from /etc/vfstab.

If the resource is not listed in the /etc/vfstab file, then the command line must specify both the *resource* and the *mount_point*.

host can be an IPv4 or IPv6 address string. As IPv6 addresses already contain colons, enclose *host* in a pair of square brackets when specifying an IPv6 address string. Otherwise the first occurrence of a colon can be interpreted as the separator between the host name and path, for example, [1080::8:800:200C:417A]:tmp/file. See [inet\(7P\)](#) and [inet6\(7P\)](#).

host:pathname

Where *host* is the name of the NFS server host, and *pathname* is the path name of the directory on the server being mounted. The path name is interpreted according to the server's path name parsing rules and is not necessarily slash-separated, though on most servers, this is the case.

nfs://host[:port]/pathname

This is an NFS URL and follows the standard convention for NFS URLs as described in NFS URL Scheme, RFC 2224. See the discussion of URL's and the public option under NFS FILE SYSTEMS for a more detailed discussion.

host:pathname nfs://host[:port]/pathname

host:pathname is a comma-separated list of *host:pathname*.

See the discussion of replicated file systems and failover under NFS FILE SYSTEMS for a more detailed discussion.

hostlist pathname

hostlist is a comma-separated list of hosts.

See the discussion of replicated file systems and failover under NFS FILE SYSTEMS for a more detailed discussion.

The `mount` command maintains a table of mounted file systems in `/etc/mnttab`, described in [mnttab\(4\)](#).

`mount_nfs` supports both NFSv3 and NFSv4 mounts. The default NFS version is NFSv4.

SMF Management

The NFS client service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/nfs/client:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

选项

See [mount\(1M\)](#) for the list of supported *generic_options*. See [share_nfs\(1M\)](#) for a description of server options.

`-o specific_options`

Set file system specific options according to a comma-separated list with no intervening spaces.

`acdirmax=n`

Hold cached attributes for no more than *n* seconds after directory update. The default value is 60.

`acdirmin=n`

Hold cached attributes for at least *n* seconds after directory update. The default value is 30.

`acregmax=n`

Hold cached attributes for no more than *n* seconds after file modification. The default value is 60.

`acregmin=n`

Hold cached attributes for at least *n* seconds after file modification. The default value is 3.

`actimeo=n`

Set *min* and *max* times for regular files and directories to *n* seconds. See “File Attributes,” below, for a description of the effect of setting this option to 0.

See “Specifying Values for Attribute Cache Duration Options,” below, for a description of how `acdirmax`, `acdirmin`, `acregmax`, `acregmin`, and `actimeo` are parsed on a `mount` command line.

`bg | fg`

If the first attempt fails, retry in the background, or, in the foreground. The default is `fg`.

forcedirectio | noforcedirectio

If **forcedirectio** is specified, then for the duration of the mount, forced direct I/O is used. If the filesystem is mounted using **forcedirectio**, data is transferred directly between client and server, with no buffering on the client. If the filesystem is mounted using **noforcedirectio**, data is buffered on the client. **forcedirectio** is a performance option that is of benefit only in large sequential data transfers. The default behavior is **noforcedirectio**.

grpuid

By default, the GID associated with a newly created file obeys the System V semantics; that is, the GID is set to the effective GID of the calling process. This behavior can be overridden on a per-directory basis by setting the set-GID bit of the parent directory; in this case, the GID of a newly created file is set to the GID of the parent directory (see [open\(2\)](#) and [mkdir\(2\)](#)). Files created on file systems that are mounted with the **grpuid** option obeys BSD semantics independent of whether the set-GID bit of the parent directory is set; that is, the GID is unconditionally inherited from that of the parent directory.

hard | soft

Continue to retry requests until the server responds (**hard**) or give up and return an error (**soft**). The default value is **hard**. Note that NFSv4 clients do not support **soft** mounts.

intr | nointr

Allow (do not allow) keyboard interrupts to kill a process that is hung while waiting for a response on a hard-mounted file system. The default is **intr**, which makes it possible for clients to interrupt applications that can be waiting for a remote mount.

llock

Use local locking (no lock manager). Note that this is a private interface.

noac

Suppress data and attribute caching. The data caching that is suppressed is the write-behind. The local page cache is still maintained, but data copied into it is immediately written to the server.

nocto

Do not perform the normal close-to-open consistency. When a file is closed, all modified data associated with the file is flushed to the server and not held on the client. When a file is opened the client sends a request to the server to validate the client's local caches. This behavior ensures a file's consistency across multiple NFS clients. When **nocto** is in effect, the client does not perform the flush on close and the request for validation, allowing the possibility of differences among copies of the same file as stored on multiple clients.

This option can be used where it can be guaranteed that accesses to a specified file system are made from only one client and only that client. Under such a condition, the effect of **nocto** can be a slight performance gain.

`port=n`

The server IP port number. The default is `NFS_PORT`. If the `port` option is specified, and if the resource includes one or more NFS URLs, and if any of the URLs include a port number, then the port number in the option and in the URL must be the same.

`posix`

Request POSIX.1 semantics for the file system. Requires a mount Version 2 [mountd\(1M\)](#) on the server. See [standards\(5\)](#) for information regarding POSIX.

`proto=netid | rdma`

By default, the transport protocol that the NFS mount uses is the first available RDMA transport supported both by the client and the server. If no RDMA transport is found, then it attempts to use a TCP transport or, failing that, a UDP transport, as ordered in the `/etc/netconfig` file. If it does not find a connection oriented transport, it uses the first available connectionless transport.

Use this option to override the default behavior.

`proto` is set to the value of *netid* or *rdma*. *netid* is the value of the `network_id` field entry in the `/etc/netconfig` file.

The UDP protocol is not supported for NFS Version 4. If you specify a UDP protocol with the `proto` option, NFS version 4 is not used.

`public`

The `public` option forces the use of the public file handle when connecting to the NFS server. The resource specified might not have an NFS URL. See the discussion of URLs and the `public` option under NFS FILE SYSTEMS for a more detailed discussion.

`quota | noquota`

Enable or prevent [quota\(1M\)](#) to check whether the user is over quota on this file system; if the file system has quotas enabled on the server, quotas are still checked for operations on this file system.

`remount`

Remounts a read-only file system as read-write (using the `rw` option). This option cannot be used with other `-o` options, and this option works only on currently mounted read-only file systems.

`retrans=n`

Set the number of NFS retransmissions to *n*. The default value is 5. For connection-oriented transports, this option has no effect because it is assumed that the transport performs retransmissions on behalf of NFS.

`retry=n`

The number of times to retry the mount operation. The default for the `mount` command is 10000.

The default for the automounter is 0, in other words, do not retry. You might find it useful to increase this value on heavily loaded servers, where automounter traffic is dropped, causing unnecessary server not responding errors.

`rsize=n`

Set the read buffer size to a maximum of *n* bytes. The default value is 1048576 when using connection-orientated transports with Version 3 or Version 4 of the NFS protocol, and 32768 when using connection-less transports. The default can be negotiated down if the server prefers a smaller transfer size. “Read” operations may not necessarily use the maximum buffer size. When using Version 2, the default value is 32768 for all transports.

`sec=mode`

Set the security *mode* for NFS transactions. If `sec=` is not specified, then the default action is to use AUTH_SYS over NFS Version 2 mounts, use a user-configured default auth over NFS version 3 mounts, or to negotiate a mode over Version 4 mounts.

The preferred mode for NFS Version 3 mounts is the default mode specified in `/etc/nfssec.conf` (see [nfssec.conf\(4\)](#)) on the client. If there is no default configured in this file or if the server does not export using the client's default mode, then the client picks the first mode that it supports in the array of modes returned by the server. These alternatives are limited to the security flavors listed in `/etc/nfssec.conf`.

NFS Version 4 mounts negotiate a security mode when the server returns an array of security modes. The client attempts the mount with each security mode, in order, until one is successful.

Only one mode can be specified with the `sec=` option. See [nfssec\(5\)](#) for the available *mode* options.

`secure`

This option has been deprecated in favor of the `sec=dh` option.

`timeo=n`

Set the NFS timeout to *n* tenths of a second. The default value is 11 tenths of a second for connectionless transports, and 600 tenths of a second for connection-oriented transports. This value is ignored for connectionless transports. Such transports might implement their own timeouts, which are outside the control of NFS.

`vers=NFS version number`

By default, the version of NFS protocol used between the client and the server is the highest one available on both systems. The default maximum for the client is Version 4. This can be changed by setting `client_versmax` to a valid version number (2, 3, or 4). Use the [sharectl\(1M\)](#) command to manipulate the `client_versmax` property. If the NFS server does not support the client's default maximum, the next lowest version attempted until a matching version is found.

`wsiz=n`

Set the write buffer size to a maximum of *n* bytes. The default value is 1048576 when using connection-orientated transports with Version 3 or Version 4 of the NFS protocol, and 32768 when using connection-less transports. The default can be negotiated down if the server prefers a smaller transfer size. “Write” operations may not necessarily use the maximum buffer size. When using Version 2, the default value is 32768 for all transports.

`xattr | noxattr`

Allow or disallow the creation and manipulation of extended attributes. The default is `xattr`. See `fsattr(5)` for a description of extended attributes.

`-O`

Overlay mount. Allow the file system to be mounted over an existing mount point, making the underlying file system inaccessible. If a mount is attempted on a pre-existing mount point without setting this flag, the mount fails, producing the error “device busy.”

Nfs File Systems

Background versus Foreground

File systems mounted with the `bg` option indicate that mount is to retry in the background if the server's mount daemon (`mountd(1M)`) does not respond. mount retries the request up to the count specified in the `retry=n` option. (Note that the default value for `retry` differs between mount and automount. See the description of `retry`, above.) Once the file system is mounted, each NFS request made in the kernel waits `timeo=n` tenths of a second for a response. If no response arrives, the time-out is multiplied by 2 and the request is retransmitted. When the number of retransmissions has reached the number specified in the `retrns=n` option, a file system mounted with the `soft` option returns an error on the request; one mounted with the `hard` option prints a warning message and continues to retry the request.

Hard versus Soft

File systems that are mounted read-write or that contain executable files should always be mounted with the `hard` option. Applications using `soft` mounted file systems can incur unexpected I/O errors, file corruption, and unexpected program core dumps. The `soft` option is not recommended.

Authenticated requests

The server can require authenticated NFS requests from the client. `sec=dh` authentication might be required. See `nfssec(5)`.

URLs and the public option

If the `public` option is specified, or if the *resource* includes an NFS URL, mount attempts to connect to the server using the public file handle lookup protocol. See WebNFS Client Specification, RFC 2054. If the server supports the public file handle, the attempt is successful; mount does not need to contact the server's `rpcbind(1M)` and the `mountd(1M)` daemons to get the port number of the mount server and the initial file handle of *pathname*, respectively. If the NFS client and server are separated by a firewall that allows all outbound connections through specific ports, such as `NFS_PORT`, then this enables NFS operations

through the firewall. The public option and the NFS URL can be specified independently or together. They interact as specified in the following matrix:

Resource Style		
	<i>host:pathname</i>	NFS URL
public option	Force public file handle and fail mount if not supported.	Force public file handle and fail mount if not supported.
	Use Native paths.	Use Canonical paths.
default	Use MOUNT protocol.	Try public file handle with Canonical paths. Fall back to MOUNT protocol if not supported.

A Native path is a path name that is interpreted according to conventions used on the native operating system of the NFS server. A Canonical path is a path name that is interpreted according to the URL rules. See Uniform Resource Locators (URL), RFC 1738. See “Examples,” below, for uses of Native and Canonical paths.

Replicated file systems and failover

resource can list multiple read-only file systems to be used to provide data. These file systems should contain equivalent directory structures and identical files. The file systems can be specified either with a comma-separated list of *host:/pathname* entries and/or NFS URL entries, or with a comma-separated list of hosts, if all file system names are the same. If multiple file systems are named and the first server in the list is down, failover uses the next alternate server to access files. If the read-only option is not chosen, replication is disabled. File access, for NFS Versions 2 and 3, is blocked on the original if NFS locks are active for that file.

File Attributes

To improve NFS read performance, files and file attributes are cached. File modification times get updated whenever a write occurs. However, file access times can be temporarily out-of-date until the cache gets refreshed.

The attribute cache retains file attributes on the client. Attributes for a file are assigned a time to be flushed. If the file is modified before the flush time, then the flush time is extended by the time since the last modification (under the assumption that files that changed recently are likely to change soon). There is a minimum and maximum flush time extension for regular files and for directories. Setting *actimeo=n* sets flush time to *n* seconds for both regular files and directories.

Setting *actimeo=0* disables attribute caching on the client. This means that every reference to attributes is satisfied directly from the server though file data is still cached. While this

guarantees that the client always has the latest file attributes from the server, it has an adverse effect on performance through additional latency, network load, and server load.

Setting the `noac` option also disables attribute caching, but has the further effect of disabling client write caching. While this guarantees that data written by an application is written directly to a server, where it can be viewed immediately by other clients, it has a significant adverse effect on client write performance. Data written into memory-mapped file pages (`mmap(2)`) are not written directly to this server.

Specifying Values for Attribute Cache Duration Options

The attribute cache duration options are `acdirmax`, `acdirmin`, `acregmax`, `acregmin`, and `actimeo`, as described under `OPTIONS`. A value specified for `actimeo` sets the values of all attribute cache duration options except for any of these options specified following `actimeo` on a mount command line. For example, consider the following command:

```
example# mount -o acdirmax=10,actimeo=1000 server:/path /localpath
```

Because `actimeo` is the last duration option in the command line, its value (1000) becomes the setting for all of the duration options, including `acdirmax`. Now consider:

```
example# mount -o actimeo=1000,acdirmax=10 server:/path /localpath
```

Because the `acdirmax` option follows `actimeo` on the command line, it is assigned the value specified (10). The remaining duration options are set to the value of `actimeo` (1000).

示例

示例 1 Mounting an NFS File System

To mount an NFS file system:

```
example# mount serv:/usr/src /usr/src
```

This is an example of the use of a native path.

示例 2 Mounting An NFS File System Read-Only With No suid Privileges

To mount an NFS file system read-only with no suid privileges:

```
example# mount -r -o nosuid serv:/usr/src /usr/src
```

示例 3 Mounting An NFS File System Over Version 2, with the UDP Transport

To mount an NFS file system over Version 2, with the UDP transport:

```
example# mount -o vers=2,proto=udp serv:/usr/src /usr/src
```

示例 4 Mounting an NFS File System Using An NFS URL

To mount an NFS file system using an NFS URL (a canonical path):

```
example# mount nfs://serv/usr/man /usr/man
```

示例 5 Mounting An NFS File System Forcing Use Of The Public File Handle

To mount an NFS file system and force the use of the public file handle and an NFS URL (a canonical path) that has a non 7-bit ASCII escape sequence:

```
example# mount -o public nfs://serv/usr/%A0abc /mnt/test
```

示例 6 Mounting an NFS File System Using a Native Path

To mount an NFS file system using a native path (where the server uses colons (":") as the component separator) and the public file handle:

```
example# mount -o public serv:C:doc:new /usr/doc
```

示例 7 Mounting a Replicated Set of NFS File Systems with the Same Pathnames

To mount a replicated set of NFS file systems with the same pathnames:

```
example# mount serv-a,serv-b,serv-c:/usr/man /usr/man
```

示例 8 Mounting a Replicated Set of NFS File Systems with Different Pathnames

To mount a replicated set of NFS file systems with different pathnames:

```
example# mount serv-x:/usr/man,serv-y:/var/man,nfs://serv-z/man /usr/man
```

文件

- /etc/mnttab
table of mounted file systems
- /etc/dfs/fstypes
default distributed file system type
- /etc/vfstab
table of automatically mounted resources

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/nfs

另请参见

[lockd\(1M\)](#), [mountall\(1M\)](#), [mountd\(1M\)](#), [nfsd\(1M\)](#), [quota\(1M\)](#), [sharectl\(1M\)](#), [statd\(1M\)](#), [mkdir\(2\)](#), [mmap\(2\)](#), [mount\(2\)](#), [open\(2\)](#), [umount\(2\)](#), [mnttab\(4\)](#), [nfssec.conf\(4\)](#), [attributes\(5\)](#), [fsattr\(5\)](#), [nfssec\(5\)](#), [standards\(5\)](#), [inet\(7P\)](#), [inet6\(7P\)](#), [lofs\(7FS\)](#)

Callaghan, Brent, WebNFS Client Specification, RFC 2054, October 1996.

Callaghan, Brent, NFS URL Scheme, RFC 2224, October 1997.

Berners-Lee, Masinter & McCahill , Uniform Resource Locators (URL), RFC 1738, December 1994.

附注

An NFS server should not attempt to mount its own file systems. See [lofs\(7FS\)](#).

If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on *the directory to which the symbolic link refers*, rather than being mounted on top of the symbolic link itself.

SunOS 4.x used the `biod` maintenance procedure to perform parallel read-ahead and write-behind on NFS clients. SunOS 5.x made `biod` obsolete with multi-threaded processing, which transparently performs parallel read-ahead and write-behind.

Since the root (/) file system is mounted read-only by the kernel during the boot process, only the `remount` option (and options that can be used in conjunction with `remount`) affect the root (/) entry in the `/etc/vfstab` file.

引用名	mount_pcfs – mount pcfs file systems
用法概要	<pre>mount -F pcfs [<i>generic_options</i>] [-o <i>FSType-specific_options</i>] <i>special</i> <i>mount_point</i> mount -F pcfs [<i>generic_options</i>] [-o <i>FSType-specific_options</i>] <i>special</i> <i>mount_point</i></pre>
描述	mount attaches an MS-DOS file system (pcfs) to the file system hierarchy at the <i>mount_point</i>, which is the pathname of a directory. If <i>mount_point</i> has any contents prior to the mount operation, these are hidden until the file system is unmounted. The pcfs file system supports direct mounting of files containing the file system as well as block devices. See mount(1M) and lofiadm(1M). If mount is invoked with <i>special</i> or <i>mount_point</i> as the only arguments, mount will search /etc/vfstab to fill in the missing arguments, including the <i>FSType-specific_options</i>; see mount(1M) for more details. The <i>special</i> argument is a special device file type, which is: A DOS logical drive on a hard disk expressed as <i>device-name:logical-drive</i>, where <i>device-name</i> specifies the special block device-file for the whole disk and <i>logical-drive</i> is either a drive letter (c through z) or a drive number (1 through 24). Examples are /dev/dsk/c0t0d0p0:c and /dev/dsk/c0t0d0p0:1. The <i>special</i> device file type must have a formatted MS-DOS file system with either a 12-bit, 16-bit, or 32-bit File Allocation Table.
选项	<i>generic_options</i> See mount(1M) for the list of supported options. -o Specify pcfs file system-specific options. The following options are supported: clamptime noclamptime File timestamps in pcfs cover a range between January 1st 1980 and December 31st 2127. This is not equal to the range of time_t on Unix for either 32-bit or 64-bit applications. In particular, 32-bit applications fail with EOVERFLOW errors on the stat(2) system call when timestamps beyond the range of 32-bit time_t are encountered. In order to prevent such spurious failures, pcfs by default clamps timestamps to the common subset of possible pcfs timestamps and the range available to 32-bit applications in Unix. The clamptime mount option therefore is active by default. If you want access to the full range of possible timestamps on pcfs, mount the file system with the noclamptime mount option. Note that if noclamptime is used, only 64-bit applications will have access to timestamps beyond January 19th 2038, 03:14:06 UTC; 32-bit applications will encounter EOVERFLOW errors.

`foldcase | nofoldcase`

Force uppercase characters in filenames to lowercase when reading them from the filesystem. This is for compatibility with the previous behavior of `pcfs`. The default is `nofoldcase`.

`hidden | nohidden`

Allow or disallow listing of files with hidden or system bits set. Option `hidden` is the default. When `nohidden` is effect, hidden and system files are neither visible nor accessible. Note that PCFS in previous releases of the Solaris operating system used the `nohidden` option as the default.

`atime | noatime`

Enable or disable write access timestamps on DOS-formatted media. Default for fixed disks is `atime`, while for removable media `noatime` is used. The latter default is so that writes to flash-based media (“memory sticks”) can be minimized, to prolong lifetime.

`timezone=timezone`

Timestamps on DOS-formatted media are recorded in the local time of the recording system. This can cause confusion when accessing removable media in which the recording and receiving system use different time zones. Use this option to force media timestamps to be interpreted for a specific time zone. The `mount_pcfs` command converts the given time zone name into a numerical offset that is passed to the `pcfs` kernel module, using the same rules as described in [environ\(5\)](#) for the `TZ` environment variable. By default, the `timezone` value is taken from the `TZ` environment variable.

`uid=user_id`

Specifies the owner for the contents of the file system. `user_id` is a decimal value. The default owner is the user ID of current process.

`gid=group_id`

Specifies the group ID for the contents of the file systems. `group_id` is a decimal value. The default group is the group ID of the current process.

`umask=user_mask`

Specifies the mask used for setting maximum access permissions on contents of the file system. The `user_mask` is an octal value, a bitmask, where each bit represents the permissions that should be disabled. The default `user_mask` is the file mode creation mask of the mounting process.

See [Intro\(2\)](#) for more information on masks.

文件

`/etc/mnttab`

table of mounted file systems

`/etc/vfstab`

list of default parameters for each file system

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/pcfs

另请参见 [lofiadm\(1M\)](#), [mount\(1M\)](#), [mountall\(1M\)](#), [Intro\(2\)](#), [mount\(2\)](#), [stat\(2\)](#), [time\(2\)](#), [mnttab\(4\)](#), [vfstab\(4\)](#), [attributes\(5\)](#), [environ\(5\)](#), [pcfs\(7FS\)](#)

附注 If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on the directory to which the symbolic link refers, rather than on top of the symbolic link itself.

引用名	mount_smbfs, umount_smbfs – 在 SMB 文件服务器上挂载和取消挂载共享资源
用法概要	<pre>/usr/sbin/mount [-F smbfs] [generic-options] [-o name=value] [-O] resource /usr/sbin/mount [-F smbfs] [generic-options] [-o name=value] [-O] mount-point /usr/sbin/mount [-F smbfs] [generic-options] [-o name=value] [-O] resource mount-point /usr/sbin/umount [-F smbfs] [generic-options] mount-point</pre>
描述	mount 实用程序将一个指定资源 <i>resource</i> 关联到路径名称位置 <i>mount-point</i> 处的文件系统结构，该位置必须已经存在。 如果 <i>mount-point</i> 在 mount 操作之前具有任何内容，这些内容将保持隐藏状态，直至取消资源的挂载。具有 SYS_MOUNT 特权的授权用户可执行 mount 操作。此外，用户可在自己拥有的目录上执行 SMBFS 挂载操作。 如果资源列在 /etc/vfstab 文件中，则可以指定 <i>resource</i> 或 <i>mount-point</i>，因为 mount 命令会向 /etc/vfstab 文件查询更多信息。如果省略 -F 选项，mount 会从 /etc/vfstab 文件所含的项中提取文件系统类型。 如果资源没有列在 /etc/vfstab 文件中，命令行必须同时指定 <i>resource</i> 和 <i>mount-point</i>。 umount 实用程序会从文件系统结构中分离挂载的文件系统。具有 SYS_MOUNT 特权的授权用户可执行 umount 操作。此外，用户可在自己拥有的目录上执行 SMBFS 取消挂载操作。 必须启用 network/smb/client 服务才能成功挂载 SMB 共享。缺省情况下将启用该服务。 要启用该服务，请输入以下 svcadm(1M) 命令： <pre># svcadm enable network/smb/client</pre> mount 命令支持下列操作数： <pre>resource //server/share</pre> 要挂载的资源名称。除了其名称以外，还可以指定以下有关资源的信息： ▪ <i>server</i> 是远程计算机的 DNS 或 NetBIOS 名称。 ▪ <i>share</i> 是远程服务器上的资源名称。 也可以指定用户帐户。请参见“选项”部分。 mount 命令可以从用户帐户的标准输入中读取口令。如果未提供口令，mount 将首先尝试使用 smbadm add-key 命令存储的口令（如果有）。如果该口令无法通过身份验证，mount_smbfs 命令将提示您输入一个口令（如果标准输入是 TTY）。
操作数	

mount-point

要挂载或取消挂载的文件系统所在位置的路径。`mount` 命令在 `/etc/mnttab` 文件中维护一个已挂载文件系统的表。请参见 [mnttab\(4\)](#) 手册页。

选项

有关支持的 *generic-options* 列表，请参见 [mount\(1M\)](#) 手册页。

-o name=value 或

-o name

设置文件系统特定的属性。可以将多个“名称-值”对指定为逗号分隔的对列表。列表中不允许有空格。这些属性如下所示：

acdirmax=*n*

目录更新后保留高速缓存的属性不超过 *n* 秒。缺省值为 60。

acdirmin=*n*

目录更新后将高速缓存的属性至少保留 *n* 秒。缺省值为 30。

acregmax=*n*

文件修改后保留高速缓存的属性不超过 *n* 秒。缺省值为 60。

acregmin=*n*

文件修改后将高速缓存的属性至少保留 *n* 秒。缺省值为 3。

actimeo=*n*

将正规文件与目录的最小和最大时间设置为 *n* 秒。有关将此选项设置为 0 的效果的说明，请参见下文的“文件属性”部分。

有关在 `mount` 命令行上如何解析 `acdirmax`、`acdirmin`、`acregmax`、`acregmin` 和 `actimeo` 的说明，请参见下面的“指定属性高速缓存持续时间选项的值”。

dirperms=*octaltriplet*

指定要分配给目录的权限。必须将值指定为三个一组的八进制数，例如 755。目录模式的缺省值来自 `fileperms` 设置，并且添加了执行权限，其中 `fileperms` 具有读取权限。

请注意，这些权限与 SMB 服务器授予的权限无关。

domain=*value*

指定用户名定义所在的工作组或 Windows 域的名称。如果未指定域名，将使用缺省系统的 SMB 域。

fileperms=*octaltriplet*

指定要分配给文件的权限。必须将值指定为三个一组的八进制数，例如 644。缺省值为 700。

请注意，这些权限与 SMB 服务器授予的权限无关。

gid=*groupid*

将指定的组 ID 分配给文件。缺省值是在其中挂载了卷的目录的组 ID。

intr|nointr

启用或禁用以下行为：在用户中断调用线程时（例如，当操作正在执行时按 Ctrl-C）取消 **smbfs(7FS)** I/O 操作。缺省值为 **intr**（启用中断），因此通常允许取消。

noac

抑制属性高速缓存。本地 **stat(2)** 调用始终会从 SMB 服务器请求属性。

noprompt

挂载共享时抑制口令提示。该属性可用于允许对共享进行匿名访问。匿名访问不需要口令。

如果需要口令，并且设置了 **noprompt** 属性，同时未通过 **smbadm add-key** 命令存储口令，**mount** 操作将失败。

retry_count=number

指定在将连接标记为中断之前 SMBFS 重新尝试的次数。缺省情况下将进行 4 次尝试。

timeout=seconds

指定 SMB 请求超时值。缺省情况下，超时值为 15 秒。

uid=userid

分配指定的用户 ID 文件。缺省值是在其中挂载了卷的目录的所有者 ID。

user=value

指定远程用户名。如果省略 **user**，则使用登录的用户 ID。

xattr|noxattr

启用（或禁用）此挂载点中的 Solaris 扩展属性。此选项缺省为 **xattr**（启用扩展属性），但请注意：如果 SMB 服务器不支持 SMB“命名流”，**smbfs(7FS)** 会将此选项强制设置为 **noxattr**。如果某个挂载设置了 **noxattr** 选项，使用 Solaris 扩展属性的尝试将失败，并生成 EINVAL 错误。

-O

覆盖挂载。允许文件系统挂载到现有的挂载点之上，使得底层文件系统不可访问。如果在未设置此标志的情况下在之前已存在的挂载点上尝试挂载，挂载将失败，并生成“设备忙”错误。

文件属性

为提高 **smbfs** 的性能，文件属性会进行高速缓存。只要发生任何本地修改，都会更新文件修改时间。但是，文件访问时间的高速缓存刷新之前可能会暂时过期。

属性高速缓存将文件属性保留在客户机上。系统会为文件的属性分配一个刷新时间。如果在刷新时间之前修改了文件，刷新时间会按照自上次修改以来的时间进行延长（假设最近更改的文件可能很快会发生更改）。正规文件和目录都有最小和最大刷新时间延长期。设置 **actimeo=n** 会将正规文件和目录的刷新时间设置为 *n* 秒。

设置 **actimeo=n** 会禁用客户机上的属性高速缓存。这意味着对属性的每个引用将直接由服务器提供。虽然这可确保客户机始终可从服务器获得最新的文件属性，但额外的等待时间、网络负载和服务器负载会对性能产生不利影响。

指定属性高速缓存持续时间选项的值

设置 `noac` 选项也会禁用属性高速缓存。当增强了 `smbfs` 以支持写高速缓存时，此选项会对禁用写高速缓存产生进一步影响。

属性高速缓存持续时间选项有 `acdirmax`、`acdirmin`、`acregmax`、`acregmin` 和 `actimeo`，如上面的“选项”部分所述。为 `actimeo` 指定的值会设置除了在 `mount` 命令行上 `actimeo` 之后指定的任何属性高速缓存持续时间选项以外的所有这些选项的值。例如，请考虑以下命令：

```
# mount -F smbfs -o acdirmax=10,actimeo=1000 \
//server/share /mountpoint
```

因为 `actimeo` 是命令行中最后一个持续时间选项，所以它的值 (`1000`) 将成为所有持续时间选项的设置，包括 `acdirmax`。现在请考虑：

```
# mount -F smbfs -o actimeo=1000,acdirmax=10 \
//server/share /mountpoint
```

因为 `acdirmax` 选项在命令行上位于 `actimeo` 之后，所以它被赋予指定值 (`10`)。其余的持续时间选项设置为 `actimeo` 的值 (`1000`)。

示例

示例1 挂载SMBFS共享

以下示例显示如何从本地 `/mnt` 挂载点上 `SALES` 工作组中的 `nano` 服务器挂载 `/tmp` 共享。您必须提供 `root` 用户的口令才能成功执行挂载操作。

```
# mount -F smbfs -o user=root,domain=SALES //nano.sfbay/tmp /mnt
Password:
```

示例2 验证SMBFS文件系统已挂载

以下示例显示如何从本地 `/mnt` 挂载点上的 `nano` 服务器挂载 `/tmp` 共享。您必须提供 `root` 用户的口令才能成功执行挂载操作。

```
# mount -F smbfs -o user=root //nano.sfbay/tmp /mnt
Password:
```

可通过下列方式验证共享是否已挂载：

- 查看 `/etc/mnttab` 文件中的文件系统项。

```
# grep mnt /etc/mnttab
//nano.sfbay/tmp /mnt smbfs dev=4900000 1177097833
```

- 查看 `mount` 命令的输出。

```
# mount | grep mnt
mnt on //nano.sfbay/tmp read/write/setuid/devices/dev=4900000 on
Tue Apr 20 13:37:13 2010
```

- 查看 `df /mnt` 命令的输出。

```
# df /mnt
/mnt                (//nano.sfbay/tmp): 3635872 blocks    -1 files
```

示例2 验证 SMBFS 文件系统已挂载 （续）

通过查看 `df -k /mnt` 命令的输出获取有关已挂载共享的信息。

```
# df -k /mnt
Filesystem            kbytes    used   avail capacity  Mounted on
//nano.sfbay/tmp
                    1882384   64448 1817936      4%    /mnt
```

示例3 取消挂载 SMB 共享

此示例假设 SMB 共享已挂载到 `/mnt` 挂载点。以下命令行将从挂载点取消共享的挂载。

```
# umount /mnt
/etc/mnttab
    其中列出了已挂载的文件系统的表。

/etc/dfs/fstypes
    缺省的分布式文件系统类型。

/etc/vfstab
    自动挂载资源的表。

/var/smb/smbfspasswd
    为 Solaris SMB 客户机存储每用户设置。
```

文件

属性

有关以下属性的说明，请参见 [attributes\(5\)](#) 手册页：

属性类型	属性值
可用性	system/file-system/ntfsprogs
接口稳定性	Committed（已确定）

另请参见

[mount\(1M\)](#)、[mountall\(1M\)](#)、[smbadm\(1M\)](#)、[svcadm\(1M\)](#)、[acl\(2\)](#)、[fcntl\(2\)](#)、[link\(2\)](#)、[mknod\(2\)](#)、[mount\(2\)](#)

附注

Solaris SMB 客户机会始终尝试使用 `gethostbyname()` 来解析主机名。如果不能解析主机名，SMB 客户机使用 NetBIOS 名称解析 (NetBIOS name resolution, NBNS)。Solaris SMB 客户机允许使用 NBNS，使 Windows 环境中的 Solaris SMB 客户机无需进行其他配置即可工作。

如果要在其上挂载文件系统的目录是一个符号链接，文件系统将挂载到符号链接所指的目录，而不会挂载到符号链接本身。

引用名	mount_tmpfs – mount tmpfs file systems
用法概要	mount [-F tmpfs] [-o <i>specific_options</i>] [-O] <i>special mount_point</i>
描述	tmpfs is a memory based file system which uses kernel resources relating to the VM system and page cache as a file system. mount attaches a tmpfs file system to the file system hierarchy at the pathname location <i>mount_point</i>, which must already exist. If <i>mount_point</i> has any contents prior to the mount operation, these remain hidden until the file system is once again unmounted. The attributes (mode, owner, and group) of the root of the tmpfs filesystem are inherited from the underlying <i>mount_point</i>, provided that those attributes are determinable. If not, the root's attributes are set to their default values. The <i>special</i> argument is usually specified as swap but is in fact disregarded and assumed to be the virtual memory resources within the system.
选项	-o <i>specific_options</i> Specify tmpfs file system specific options in a comma-separated list with no intervening spaces. If invalid options are specified, a warning message is printed and the invalid options are ignored. The following options are available: size=sz The sz argument controls the size of this particular tmpfs file system. If the argument is has a 'k' suffix, the number will be interpreted as a number of kilobytes. An 'm' suffix will be interpreted as a number of megabytes. No suffix is interpreted as bytes. In all cases, the actual size of the file system is the number of bytes specified, rounded up to the physical pagesize of the system. xattr noxattr Allow or disallow the creation and manipulation of extended attributes. The default is xattr. See fsattr(5) for a description of extended attributes. -O Overlay mount. Allow the file system to be mounted over an existing mount point, making the underlying file system inaccessible. If a mount is attempted on a pre-existing mount point without setting this flag, the mount will fail, producing the error device busy.
文件	/etc/mnttab Table of mounted file systems
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[mount\(1M\)](#), [mkdir\(2\)](#), [mount\(2\)](#), [open\(2\)](#), [umount\(2\)](#), [mnttab\(4\)](#), [attributes\(5\)](#), [fsattr\(5\)](#),
[tmpfs\(7FS\)](#)

附注

If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on the directory to which the symbolic link refers, rather than on top of the symbolic link itself.

引用名	mount_udfs – mount a udfs file system
用法概要	<pre>mount -F udfs [<i>generic_options</i>] [-o <i>specific_options</i>] [-O] <i>special mount_point</i> mount -F udfs [<i>generic_options</i>] [-o <i>specific_options</i>] [-O] <i>special</i> <i>mount_point</i></pre>
描述	The mount utility attaches a udfs file system to the file system hierarchy at the <i>mount_point</i>, which is the pathname of a directory. If <i>mount_point</i> has any contents prior to the mount operation, these are hidden until the file system is unmounted. If mount is invoked with either <i>special</i> or <i>mount_point</i> as the only arguments, mount searches <i>/etc/vfstab</i> to fill in the missing arguments, including the <i>specific_options</i>. See mount(1M). The udfs file system supports direct mounting of files containing the file system as well as block devices. See mount(1M) and lofiadm(1M). If <i>special</i> and <i>mount_point</i> are specified without any <i>specific_options</i>, the default is <i>rw</i>. If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on the directory to which the symbolic link refers, rather than on top of the symbolic link itself.
选项	See mount(1M) for the list of supported <i>generic_options</i>. The following options are supported: -o <i>specific_options</i> Specify udfs file system specific options in a comma-separated list with no intervening spaces. The following <i>specific_options</i> are available: m Mount the file system without making an entry in <i>/etc/mnttab</i>. remount Remount the file system as read-write. The option is used in conjunction with the <i>rw</i> option. A file system mounted read-only can be remounted as read-write. This option fails if the file system is not currently mounted. -O Overlay mount. Allow the file system to be mounted over an existing mount point, making the underlying file system inaccessible. If a mount is attempted on a pre-existing mount point without setting this flag, the mount fails, producing the error device busy.

文件

`/etc/mnttab`
Table of mounted file systems

`/etc/vfstab`
List of default parameters for each file system

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/udfs

另请参见

[fsck\(1M\)](#), [fsck_udfs\(1M\)](#), [lofiadm\(1M\)](#), [mount\(1M\)](#), [mountall\(1M\)](#), [mount\(2\)](#), [mnttab\(4\)](#), [vfstab\(4\)](#), [attributes\(5\)](#)

诊断

not super user
The command is run by a non-root user. Run as root.

no such device
The device name specified does not exist.

not a directory
The specified mount point is not a directory.

is not an udfs file system
The device specified does not contain a udf 1.50 file system or the udfs file system module is not available.

is already mounted
The specified device is already in use.

not a block device
The device specified is not a block device. Use block device to mount.

write-protected
The device is read-only.

is corrupted. needs checking
The file system is in an inconsistent state. Run `fsck`.

附注

Copy-protected files can be stored on DVD-ROM media using Universal Disk Format (UDF). Reading these copy-protected files is not possible as this involves an authentication process. Unless an authentication process between the host and the drive is completed, reading these copy-protected files after mounting and before the authentication process, returns an error.

引用名	mount_ufs – mount ufs file systems
用法概要	<pre>mount -F ufs [<i>generic_options</i>] [-o <i>specific_options</i>] [-O] <i>special</i> <i>mount_point</i> mount -F ufs [<i>generic_options</i>] [-o <i>specific_options</i>] [-O] <i>special</i> <i>mount_point</i></pre>
描述	The mount utility attaches a ufs file system to the file system hierarchy at the <i>mount_point</i>, which is the pathname of a directory. If <i>mount_point</i> has any contents prior to the mount operation, these are hidden until the file system is unmounted. The ufs file system supports direct mounting of files containing the file system as well as block devices. See mount(1M) and lofiadm(1M). If mount is invoked with <i>special</i> or <i>mount_point</i> as the only arguments, mount will search /etc/vfstab to fill in the missing arguments, including the <i>specific_options</i>. See mount(1M). If <i>special</i> and <i>mount_point</i> are specified without any <i>specific_options</i>, the default is rw. If the directory on which a file system is to be mounted is a symbolic link, the file system is mounted on the directory to which the symbolic link refers, rather than on top of the symbolic link itself.
选项	See mount(1M) for the list of supported <i>generic_options</i>. The following options are supported: -o <i>specific_options</i> Specify ufs file system specific options in a comma-separated list with no intervening spaces. If invalid options are specified, a warning message is printed and the invalid options are ignored. The following options are available: dfratime nodfratime By default, writing access time updates to the disk may be deferred (dfratime) for the file system until the disk is accessed for a reason other than updating access times. nodfratime disables this behavior. If power management is enabled on the system, do not set nodfratime unless noatime is also set. If you set nodfratime without setting noatime, the disk is spun up every time a file within a file system on the disk is accessed - even if the file is not modified. forcedirectio noforcedirectio If forcedirectio is specified and supported by the file system, then for the duration of the mount, forced direct I/O will be used. If the filesystem is mounted using forcedirectio, data is transferred directly between user address space and the disk. If the filesystem is mounted using noforcedirectio, data is buffered in kernel address space when data is transferred between user address space and the disk. forcedirectio is a performance option that is of benefit only in large sequential data transfers. The default behavior is noforcedirectio.

global | noglobal

If **global** is specified and supported on the file system, and the system in question is part of a cluster, the file system will be globally visible on all nodes of the cluster. If **noglobal** is specified, the mount will not be globally visible. The default behavior is **noglobal**.

intr | nointr

Allow (do not allow) keyboard interrupts to kill a process that is waiting for an operation on a locked file system. The default is **intr**.

largefiles | nolargefiles

If **nolargefiles** is specified and supported by the file system, then for the duration of the mount it is guaranteed that all regular files in the file system have a size that will fit in the smallest object of type **off_t** supported by the system performing the mount. The mount will fail if there are any files in the file system not meeting this criterion. If **largefiles** is specified, there is no such guarantee. The default behavior is **largefiles**.

If **nolargefiles** is specified, **mount** will fail for **ufs** if the file system to be mounted has contained a large file (a file whose size is greater than or equal to 2 Gbyte) since the last invocation of **fsck** on the file system. The large file need not be present in the file system at the time of the mount for the mount to fail; it could have been created previously and destroyed. Invoking **fsck** (see [fsck_ufs\(1M\)](#)) on the file system will reset the file system state if no large files are present. After invoking **fsck**, a successful mount of the file system with **nolargefiles** specified indicates the absence of large files in the file system; an unsuccessful mount attempt indicates the presence of at least one large file.

logging | nologging

If **logging** is specified, then logging is enabled for the duration of the mounted file system. Logging is the process of storing transactions (changes that make up a complete UFS operation) in a log before the transactions are applied to the file system. Once a transaction is stored, the transaction can be applied to the file system later. This prevents file systems from becoming inconsistent, therefore reducing the possibility that **fsck** might run. And, if **fsck** is bypassed, logging generally reduces the time required to reboot a system.

The default behavior is **logging** for all UFS file systems.

The log is allocated from free blocks in the file system, and is sized approximately 1 Mbyte per 1 Gbyte of file system, up to a maximum of 256 Mbytes. The log size may be larger (up to a maximum of 512 Mbytes) dependent upon the number of cylinder groups present in the file system.

Logging is enabled on any UFS file system, including root (/), except under the following conditions:

- When logging is specifically disabled.
- If there is insufficient file system space for the log. In this case, the following message is displayed and file system is still mounted:

```
# mount /dev/dsk/c0t4d0s0 /mnt
/mnt: No space left on device
Could not enable logging for /mnt on /dev/dsk/c0t4d0s0.
```

The log created by UFS logging is continually flushed as it fills up. The log is totally flushed when the file system is unmounted or as a result of the `lockfs -f` command.

m

Mount the file system without making an entry in `/etc/mnttab`.

noatime

By default, the file system is mounted with normal access time (`atime`) recording. If `noatime` is specified, the file system will ignore access time updates on files, except when they coincide with updates to the `ctime` or `mtime`. See [stat\(2\)](#). This option reduces disk activity on file systems where access times are unimportant (for example, a Usenet news spool).

`noatime` turns off access time recording regardless of `dfratime` or `nodfratime`.

The POSIX standard requires that access times be marked on files. `-noatime` ignores them unless the file is also modified.

nosec

By default, Access Control Lists (ACLs) are supported on a mounted UFS file system. Use this option to disallow the setting or any modification of an ACL on a file within a mounted UFS file system. See [getfacl\(1\)](#) for background on ACLs.

onerror = *action*

This option specifies the action that UFS should take to recover from an internal inconsistency on a file system. Specify *action* as `panic`, `lock`, or `umount`. These values cause a forced system shutdown, a file system lock to be applied to the file system, or the file system to be forcibly unmounted, respectively. The default is `panic`.

quota

Quotas are turned on for the file system.

remount

Remounts a file system with a new set of options. All options not explicitly set with `remount` revert to their default values.

rw

Read-write with quotas turned on. Equivalent to `rw`, `quota`.

xattr | noxattr

Allow or disallow the creation and manipulation of extended attributes. The default is `xattr`. See [fsattr\(5\)](#) for a description of extended attributes.

-O

Overlay mount. Allow the file system to be mounted over an existing mount point, making the underlying file system inaccessible. If a mount is attempted on a pre-existing mount

point without setting this flag, the mount will fail, producing the error “device busy”.

示例

示例 1 Turning Off (and On) Logging

The following command turns off logging on an already mounted file system. The subsequent command restores logging.

```
# mount -F ufs -o remount,nologging /export
# (absence of message indicates success)
# mount -F ufs -o remount,logging /export
```

In the preceding commands, the -F ufs option is not necessary.

文件

/etc/mnttab
table of mounted file systems

/etc/vfstab
list of default parameters for each file system

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[getfacl\(1\)](#), [fsck\(1M\)](#), [fsck_ufs\(1M\)](#), [lofiadm\(1M\)](#), [mount\(1M\)](#), [mountall\(1M\)](#), [fcntl\(2\)](#), [mount\(2\)](#), [stat\(2\)](#), [mnttab\(4\)](#), [vfstab\(4\)](#), [attributes\(5\)](#), [fsattr\(5\)](#), [largefile\(5\)](#)

附注

Since the root (/) file system is mounted read-only by the kernel during the boot process, only the remount option (and options that can be used in conjunction with remount) affect the root (/) entry in the /etc/vfstab file.

引用名	mpathadm – multipath discovery and administration														
用法概要	<code>mpathadm subcommand direct-object [options] [operand]</code>														
描述	The <code>mpathadm</code> command enables multipathing discovery and management. The <code>mpathadm</code> command is implemented as a set of subcommands, many with their own options, that are described in the section for that subcommand. Options not associated with a particular subcommand are described under OPTIONS. The <code>mpathadm</code> subcommands operate on a <i>direct-object</i>. These are described in this section for each subcommand. The <i>direct-objects</i>, <i>initiator-port</i>, <i>target-port</i>, and <i>logical-unit</i> in the subcommands are consistent with SCSI standard definitions. The <code>mpathadm</code> command supports the following subcommands, which are described in detail in subsections that follow. <table><tr><td><code>list</code></td><td>Display a list of discovered instances for a given object.</td></tr><tr><td><code>show</code></td><td>Display information about a given object instance.</td></tr><tr><td><code>modify</code></td><td>Modify properties of an object.</td></tr><tr><td><code>enable</code></td><td>Enable an object.</td></tr><tr><td><code>disable</code></td><td>Disable an object.</td></tr><tr><td><code>failover</code></td><td>Cause target port group failover for a logical-unit.</td></tr><tr><td><code>override</code></td><td>Set a path to be used over other paths on a logical-unit.</td></tr></table> The <code>mpathadm</code> subcommands operate on a <i>direct-object</i>. These are described in this section for each subcommand.	<code>list</code>	Display a list of discovered instances for a given object.	<code>show</code>	Display information about a given object instance.	<code>modify</code>	Modify properties of an object.	<code>enable</code>	Enable an object.	<code>disable</code>	Disable an object.	<code>failover</code>	Cause target port group failover for a logical-unit.	<code>override</code>	Set a path to be used over other paths on a logical-unit.
<code>list</code>	Display a list of discovered instances for a given object.														
<code>show</code>	Display information about a given object instance.														
<code>modify</code>	Modify properties of an object.														
<code>enable</code>	Enable an object.														
<code>disable</code>	Disable an object.														
<code>failover</code>	Cause target port group failover for a logical-unit.														
<code>override</code>	Set a path to be used over other paths on a logical-unit.														
list Subcommand	The syntax for the <code>list</code> subcommand is: <pre># mpathadm list direct-object [operands...]</pre> The <code>list</code> subcommand displays data for following direct-objects: <table><tr><td><code>mpath-support [mpath-support-name, ...]</code></td><td>List the multipathing support that can be administered by this CLI. This presents itself in the form of a library name registered through the MPAPI framework. If no <i>mpath-support name</i> <i>mpath-support-name</i> is specified, all registered multipathing support libraries will be displayed.</td></tr><tr><td><code>initiator-port [initiator-port-name, ...]</code></td><td>List the initiator ports that are discovered on this system. If no <i>initiator-port-name</i> is specified, all discovered initiator ports are displayed.</td></tr><tr><td><code>{logical-unit lu} [options] [logical-unit-name, ...]</code></td><td>List the information on multipath logical units. If no <i>logical-unit-name</i> is specified, all discovered logical-units will be displayed.</td></tr></table>	<code>mpath-support [mpath-support-name, ...]</code>	List the multipathing support that can be administered by this CLI. This presents itself in the form of a library name registered through the MPAPI framework. If no <i>mpath-support name</i> <i>mpath-support-name</i> is specified, all registered multipathing support libraries will be displayed.	<code>initiator-port [initiator-port-name, ...]</code>	List the initiator ports that are discovered on this system. If no <i>initiator-port-name</i> is specified, all discovered initiator ports are displayed.	<code>{logical-unit lu} [options] [logical-unit-name, ...]</code>	List the information on multipath logical units. If no <i>logical-unit-name</i> is specified, all discovered logical-units will be displayed.								
<code>mpath-support [mpath-support-name, ...]</code>	List the multipathing support that can be administered by this CLI. This presents itself in the form of a library name registered through the MPAPI framework. If no <i>mpath-support name</i> <i>mpath-support-name</i> is specified, all registered multipathing support libraries will be displayed.														
<code>initiator-port [initiator-port-name, ...]</code>	List the initiator ports that are discovered on this system. If no <i>initiator-port-name</i> is specified, all discovered initiator ports are displayed.														
<code>{logical-unit lu} [options] [logical-unit-name, ...]</code>	List the information on multipath logical units. If no <i>logical-unit-name</i> is specified, all discovered logical-units will be displayed.														

Options for `list logical-unit` are as follows:

- n, --name *name*

Return the logical unit name that is associated with the given name string. This name can be extracted from the output of the `mpathadm show lu` command.
- t, --target-port *target-port-name*

Return the list of logical units names that are associated with the given *target-port-name*.

show Subcommand The syntax for the show subcommand is:

```
# mpathadm show direct-object [operands...]
```

The show subcommand displays detailed information for following the direct-objects:

- `mpath-support` [*mpath-support-name, ...*]

Show the detailed information on the given *mpath-support-name* if the name exists. If the given *mpath-support-name* supports only a limited set of device products, the list of device products will be listed as part of the output.
- `initiator-port` *initiator-port-name*[,*initiator-port-name, ...*]

Show the detailed information for the given *initiator-port-name*.
- `{logical-unit | lu}` [*logical-unit-name, ...*]

Display the detailed information on multipath logical unit(s), including path and target port group information. Note that the name property in the logical unit information represents the identifier for this LUN, derived from the hardware, and used by this system. If the name is derived from SCSI Inquiry Vital Product Data (VPD) page 83h, the name type property represents an associated identifier type defined by the SCSI standards.

modify Subcommand The syntax for the modify subcommand is:

```
# mpathadm modify direct-object [options] [operands...]
```

The modify subcommand modifies characteristics of the following direct-objects:

- `mpath-support` [*options*] *mpath-support-name, ...*

Configuration management of an *mpath-support*. Options to modify mpath-support are as follows:

~~Set auto failback applyable~~
Set `auto failback applyable` only when mpath-support provides auto failback support.

~~Set auto probing applyable~~
Set `auto probing applyable` only when mpath-support provides auto probing support.

~~Change the default load balance type~~
The loadbalance type is one of the

		supported types listed in the show mpath-support output.
	<code>{logical-unit lu} [options] logical-unit-name, ...</code>	Configuration management of a logical unit. Options to modify logical-unit are as follows: <code>Set auto-failback</code> Applicable only when mpath-support provides auto failback support <code>Set auto-probing</code> Applicable only when mpath-support provides auto probing support. <code>Set load-balance-type</code> Set the load balance type of the logical unit. Applicable only when load balance configuration is supported at the logical unit level.
enable Subcommand	The syntax for the enable subcommand is: <pre># mpathadm enable [options]</pre> The enable subcommand supports the following direct-objects to be enabled: <pre>path -i initiator-port-name -t target-port-name -l logical-unit-name</pre>	The path that consists of the specified initiator port, target port, and logical unit will be enabled.
disable Subcommand	The syntax for the disable subcommand is: <pre># mpathadm disable [options]</pre> The disable subcommand supports the following direct-objects to be disabled: <pre>path -i initiator-port-name -t target-port-name -l logical-unit-name</pre>	The path that consists of the specified initiator port, target port, and logical unit will be disabled.
failover Subcommand	The syntax for the failover subcommand is: <pre># mpathadm failover direct-object [operand]</pre> The failover subcommand supports failover for the following direct-objects: <pre>{logical-unit lu} logical-unit-name</pre>	The target port group will failover when the given logical-unit is asymmetric and supports explicit

state change. The currently active target port group will be changed to the standby state and the standby target port group will be active.

override Subcommand The syntax for the override subcommand is:

```
# mpathadm override [options]
```

The override subcommand controls whether or not the following direct-objects override another:

```
path { -i initiator-port-name -t target-port-name | -c }  
-l logical-unit-name
```

Cause a path that consists of the specified initiator port, target port, and logical unit to override other paths on the logical unit. Once a path overrides other paths, the mpath-support uses only that path regardless of any other path selection configuration. The -c option cancels the setting. The path that consists of the specified initiator port, target port, and logical unit will be disabled.

Options for override path are as follows:

-i, --initiator-port <i>initiator-port-name</i>	Represent the initiator port element of a path. Options -t and -l must also be included.
-t, --target-port <i>target-port-name</i>	Represent the target port element of a path. Options -i and -l must also be included.
-l, --logical-unit <i>logical-unit</i>	Represent the logical unit element of a path. Options -i and -t must also be included.
-c, --cancel	Cancels overriding setting for the given logical unit. Option -l must also be included.

选项 The following options are supported:

-V, --version	Display the version information.
-, --help	Display context help. Can be used following an mpathadm command with no arguments, following a subcommand, or following a subcommand direct-object combination. Responds with help information appropriate for your entry. For example, if you enter:

```
# mpathadm add mpath-support-help  
  
... mpathadm responds with a display of the options available for that  
combination of subcommand and direct-object.
```

示例**示例 1** Obtaining a List of Multipathing Support

The following command uses the `list` subcommand to list all currently registered `mpath-support` libraries.

```
# mpathadm list mpath-support
mpath-support: libmpscsi_vhci.so
```

示例 2 Displaying the Properties of a Particular Multipathing Support

The following command uses the `show` subcommand to display the properties for a currently registered `mpath-support` library.

```
# mpathadm show mpath-support libmpscsi_vhci.so
mpath-support:    libmpscsi_vhci.so
                  Vendor:    Sun Microsystems
                  Driver Name: scsi_vhci
                  Default Load Balance: round-robin
                  Supported Load Balance Types:
                      round-robin
                      logical-block
Allows To Activate Target Port Group Access: yes
Allows Path Override: no
Supported Auto Failback Config:    1
Auto Failback:    on
Failback Polling Rate (current/max): 0/0
Supported Auto Probing Config:    0
Auto Probing:    NA
Probing Polling Rate (current/max): NA/NA
Supported Devices:
    Vendor:    SUN
    Product:    T300
    Revision:
    Supported Load Balance Types:
        round-robin
    Vendor:    SUN
    Product:    T4
    Revision:
    Supported Load Balance Types:
        round-robin
```

示例 3 Obtaining a List of Initiator Ports Discovered Through the `mpath-support` Libraries

The following command uses the `list initiator-port` subcommand to display a list of initiator ports discovered by the currently registered `mpath-support` libraries.

```
# mpathadm list initiator-port
Initiator-Port: iqn.1986-03.com.sun:01:080020b7ac2b.437a3b3e,4000002a0000
Initiator-Port: 2000000173018713
Initiator-Port: 2000000173818713
```

示例 4 Displaying the Properties of a Particular Initiator Port

The following command uses the `show initiator-port` subcommand to display the properties of a particular initiator port discovered using the `list initiator-port` subcommand in an example above.

```
# mpathadm show initiator-port 2000000173018713
initiator-port:      2000000173018713
      Transport Type:      Fibre Channel
      OS device File:      devices/pci@1f,4000/pci@2/SUNW,qlca@5/fp@0,0:fc
```

示例 5 Displaying the Properties of a Particular Logical Unit

The following command uses the `show logical-unit` subcommand to display the properties of the logical unit with the specified name.

```
# mpathadm show lu /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
Logical Unit: /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
      mpath-support libmpscsi_vhci.so
      Vendor: SUN
      Product: T4
      Revision: 0301
      Name Type: SCSI Inquiry VPD Page 83 type 3
      Name: 60003ba27d2120004204ac2b000dab00
      Asymmetric: yes
      Current Load Balance: round-robin
      Logical Unit Group ID: NA
      Aauto Failback: on
      Auto Probing: NA
```

Paths:

```
      Initiator Port Name: 2000000173818713
      Target Port Name: 20030003ba27d212
      Override Path: NA
      Path State: OK
      Disabled: no
```

```
      Initiator Port Name: 2000000173018713
      Target Port Name: 20030003ba27d095
      Override Path: NA
      Path State: OK
      Disabled: no
```

Target Port Group:

```
      ID: 2
      Explicit Failover: yes
      Access State: standby
      Target Ports:
          Name: 20030003ba27d212
```

示例 5 Displaying the Properties of a Particular Logical Unit (续)

```
Relative ID: 0

ID: 5
Explicit Failover: yes
Access State: active
Target Ports
    Name: 20030003ba27d095
    Relative ID: 0
```

示例 6 Enabling a Path

The following command uses the enable path subcommand to enable the path with the specified initiator port, target port, and logical unit.

```
# mpathadm enable path -i 2000000173018713 -t 20030003ba27d095 \
-l /dev/rdisk/c4t60003BA27D2120004204AC2B000DAB00d0s2
```

示例 7 Modifying mpath-support To Turn On autofailback

```
# mpathadm modify mpath-support -a on libmpscsi_vhci.so
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/storage/multipath-utilities
	system/library/storage/libmpapi (exec_attr(4) entry)
Interface Stability	Committed

另请参见

[stmsboot\(1M\)](#), [libMPAPI\(3LIB\)](#), [exec_attr\(4\)](#), [attributes\(5\)](#)

引用名

mpstat – report per-processor or per-processor-set statistics

用法概要

`/usr/bin/mpstat [-aq] [-p | -P set] [-T u | d] [interval [count]]`

描述

The `mpstat` command reports processor statistics in tabular form. Each row of the table represents the activity of one processor. The first table summarizes all activity since boot. Each subsequent table summarizes activity for the preceding interval. All values are rates listed as events per second unless otherwise noted.

During execution of the kernel status command, the state of the kernel can change. If relevant, a state change message is included in the `mpstat` output, in one of the following forms:

```
<<processor 3 moved from pset: -1 to: 1>>
<<pset destroyed: 1>>
<<pset created: 1>>
<<processors added: 1, 3>>
<<processors removed: 1, 3>>
```

The `mpstat` command reports the following information:

CPU or SET	Without the <code>-a</code> option, <code>mpstat</code> reports CPU statistics for a processor ID. With the <code>-a</code> option, <code>mpstat</code> reports SET statistics for a processor set ID.
<code>minf</code>	minor faults
<code>mjf</code>	major faults
<code>xcal</code>	inter-processor cross-calls
<code>intr</code>	interrupts
<code>ithr</code>	interrupts as threads (not counting clock interrupt)
<code>csw</code>	context switches
<code>icsw</code>	involuntary context switches
<code>migr</code>	thread migrations (to another processor)
<code>smtx</code>	spins on mutexes (lock not acquired on first try)
<code>srw</code>	spins on readers/writer locks (lock not acquired on first try)
<code>syscl</code>	system calls
<code>usr</code>	percent user time
<code>sys</code>	percent system time
<code>wt</code>	the I/O wait time is no longer calculated as a percentage of CPU time, and this statistic will always return zero.
<code>idl</code>	percent idle time

size number of processors in the requested processor set

set processor set membership of each CPU

选项

The following options are supported:

- a Aggregate output by processor set. Sort the output by set. The default output is sorted by CPU number.
 - p Report processor set membership of each CPU. Sort the output by set. The default output is sorted by CPU number.
 - P *set* Display only those processors in the specified *set*.
 - q Suppress messages related to state changes.
 - T *u* | *d* Specify *u* for a printed representation of the internal representation of time. See [time\(2\)](#). Specify *d* for standard date format. See [date\(1\)](#).
- interval* Report once each *interval* seconds.
- count* Only print *count* reports.

示例

示例 1 Using mpstat to Generate User and System Operation Statistics

The following command generates processor statistics over a five-second interval in two reports. The command shows the processor set membership of each CPU. The default output is sorted by CPU number, aggregated by *processor set*, for user (*usr*) and system (*sys*) operations.

example% mpstat -ap 5 2

```
SET minf mjf xcal intr ithr csw icsw migr smtx srw syscl usr sys wt idl size
 0    6    0  355  291  190  22    0    0    0    0   43    0    2    0  43    1
 1   24   17  534  207  200  70    1    0    2    0  600    4    1    0  84    2
 2   19    7  353  325  318  44    0    0    5    0  345    1    1    0  94    3
 3   36    2  149  237  236  14    0    0    4    0   97    0    0    0  98    2
SET minf mjf xcal intr ithr csw icsw migr smtx srw syscl usr sys wt idl size
 0    1    0  720  405  304  55    0    0   18    0   12    0   15    0  81    1
 1    0   69 1955  230  200 313   33    4   41    9 7086   34   10    0  19    2
 2    0   46  685  314  300 203   11    0   54    1 5287   36    6    0  28    3
 3    0    0   14  386  384    0    0    0    0    0    0    0    0 100    2
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	See below.

Invocation is evolving. Human readable output is unstable.

另请参见

[sar\(1\)](#), [date\(1\)](#), [iostat\(1M\)](#), [sar\(1M\)](#), [vmstat\(1M\)](#), [time\(2\)](#), [attributes\(5\)](#)

附注

The sum of CPU utilization might vary slightly from 100 due to rounding errors in the production of a percentage figure.

The total time used for CPU processing is the sum of `usr` and `sys` output values, reported for user and system operations. The `idl` value reports the time that the CPU is idle for any reason other than pending disk I/O operations.

Run the `iostat` command with the `-x` option to report I/O service times in `svc_t` output. The `iostat` utility also reports the same `wt`, user (`us`), and system (`sy`) statistics. See [iostat\(1M\)](#) for more information.

When executing in a zone and if the pools facility is active, `mpstat(1M)` will only provide information for those processors which are a member of the processor set of the pool to which the zone is bound.

引用名 msgid – generate message IDs

用法概要 /usr/sbin/msgid

描述 The msgid utility generates message IDs.

A message ID is a numeric identifier that uniquely identifies a message. Although the probability of two distinct messages having the same ID is high, this can be greatly reduced with the appropriate priority or facility.level designator (see [syslogd\(1M\)](#)). Specifically, the message ID is a hash signature on the message's unexpanded format string, generated by STRLOG_MAKE_MSGID() as defined in <sys/strlog.h>.

[syslogd\(1M\)](#) is a simple filter that takes strings as input and produces those same strings, preceded by their message IDs, as output. Every message logged by syslogd(1M) includes the message ID. The message ID is intended to serve as a small, language-independent identifier.

示例 示例 1 Using the msgid command to generate a message ID

The following example uses the msgid command to generate a message ID for the echo command.

```
example# echo hello | msgid205790 hello
```

示例 2 Using the msgid command to generate a message catalog

The following example uses the msgid command to enumerate all of the messages in the binary ufs, to generate a message catalog.

```
example# strings /kernel/fs/ufs | msgid
```

```
137713 free:
    freeing free frag, dev:0x%lx, blk:%ld, cg:%d, ino:%lu, fs:%s
567420 iallocg: block not in mapfs = %s
845546 alloc: %s: file system full
...
```

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [syslogd\(1M\)](#), [attributes\(5\)](#), [log\(7D\)](#)

引用名	mvdir – move a directory				
用法概要	<code>/usr/sbin/mvdir <i>dirname</i> <i>name</i></code>				
描述	mvdir moves directories within a file system. <i>dirname</i> must be a directory. If <i>name</i> does not exist, it will be created as a directory. If <i>name</i> does exist, and is a directory, <i>dirname</i> will be created as <i>name/dirname</i>. <i>dirname</i> and <i>name</i> may not be on the same path; that is, one may not be subordinate to the other. For example: example%<code>mvdir x/y x/z</code> is legal, but example%<code>mvdir x/y x/y/z</code> is not.				
操作数	<table><tr><td><i>dirname</i></td><td>The name of the directory that is to be moved to another directory in the filesystem.</td></tr><tr><td><i>name</i></td><td>The name of the directory into which <i>dirname</i> is to be moved. If <i>name</i> does not exist, it will be created. It may not be on the same path as <i>dirname</i>.</td></tr></table>	<i>dirname</i>	The name of the directory that is to be moved to another directory in the filesystem.	<i>name</i>	The name of the directory into which <i>dirname</i> is to be moved. If <i>name</i> does not exist, it will be created. It may not be on the same path as <i>dirname</i> .
<i>dirname</i>	The name of the directory that is to be moved to another directory in the filesystem.				
<i>name</i>	The name of the directory into which <i>dirname</i> is to be moved. If <i>name</i> does not exist, it will be created. It may not be on the same path as <i>dirname</i> .				
用法	See largefile(5) for the description of the behavior of mvdir when encountering files greater than or equal to 2 Gbyte (2 ³¹ bytes).				
退出状态	<table><tr><td>0</td><td>Successful operation.</td></tr><tr><td>>0</td><td>Operation failed.</td></tr></table>	0	Successful operation.	>0	Operation failed.
0	Successful operation.				
>0	Operation failed.				
属性	See attributes(5) for descriptions of the following attributes:				
另请参见	mkdir(1) , mv(1) , attributes(5) , largefile(5)				

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

引用名	named, in.named – Internet domain name server
用法概要	<pre>named [-fgsVv] [-c <i>config-file</i>] [-d <i>debug-level</i>] [-m <i>flag</i>] [-n <i>#cpus</i>] [-p <i>port</i>] [-S <i>#max-socks</i>] [-t <i>directory</i>] [-u <i>user</i>] [-x <i>cache-file</i>] [-4 -6]</pre>
描述	The named utility is a Domain Name System (DNS) server, part of the BIND 9 distribution from ISC. For more information on the DNS, see RFCs 1033, 1034, and 1035. When invoked without arguments, named reads the default configuration file <code>/etc/named.conf</code>, reads any initial data, and listens for queries. <code>in.named</code> is a link to <code>named</code>.
选项	The following options are supported: -4 Use only IPv4 transport. By default, both IPv4 and IPv6 transports can be used. Options -4 and -6 are mutually exclusive. -6 Use only IPv6 transport. By default, both IPv4 and IPv6 transports can be used. Options -4 and -6 are mutually exclusive. -c <i>config-file</i> Use <i>config-file</i> as the configuration file instead of the default <code>/etc/named.conf</code>. To ensure that reloading the configuration file continues to work after the server has changed its working directory due to a possible <i>directory</i> option in the configuration file, <i>config-file</i> should be an absolute pathname. -d <i>debug-level</i> Set the daemon's debug level to <i>debug-level</i>. Debugging traces from named become more verbose as the debug level increases. -f Run the server in the foreground (that is, do not run as a daemon). -g Run the server in the foreground and force all logging to <code>stderr</code>. -m <i>flag</i> Turn on memory usage debugging flags. Possible flags are <code>usage</code>, <code>trace</code>, and <code>record</code>, <code>size</code>, and <code>mctx</code>. These correspond to the <code>ISC_MEM_DEBUGXXXX</code> flags described in <code><isc/mem.h></code>. -n <i>#cpus</i> Create <i>#cpus</i> worker threads to take advantage of multiple CPUs. If not specified, named will try to determine the number of CPUs present and create one thread per CPU. If it is unable to determine the number of CPUs, a single worker thread will be created. -p <i>port</i> Listen for queries on port <i>port</i>. If not specified, the default is port 53.

-S #max-socks

Allow named to use up to *#max-socks* sockets.

This option should be unnecessary for the vast majority of users. The use of this option could even be harmful, because the specified value might exceed the limitation of the underlying system API. It therefore should be set only when the default configuration causes exhaustion of file descriptors and the operational environment is known to support the specified number of sockets. Note also that the actual maximum number is normally a little smaller than the specified value because named reserves some file descriptors for its internal use.

-s

Write memory usage statistics to *stdout* on exit.

This option is mainly of interest to BIND 9 developers and might be removed or changed in a future release.

-t directory

Change the root directory using [chroot\(2\)](#) to *directory* after processing the command line arguments, but before reading the configuration file.

This option should be used in conjunction with the **-u** option, as chrooting a process running as root does not enhance security on most systems; the way `chroot()` is defined allows a process with root privileges to escape a chroot jail.

-u user

Set the real user ID using [setuid\(2\)](#) to *user* after completing privileged operations, such as creating sockets that listen on privileged ports.

-V

Report the version number and build options, and exit.

-v

Report the version number and exit.

-x cache-file

Load data from *cache-file* into the cache of the default view.

Do not use this option. It is of interest only to BIND 9 developers and might be removed or changed in a future release.

扩展描述

This section describes additional attributes of named.

SMF Properties

When starting named from the service management facility, [smf\(5\)](#), named configuration is read from the service configuration repository. Use [svccprop\(1\)](#) to list the properties and [svccfg\(1M\)](#) to make changes.

The following application configuration properties are available to administrators:

options/server

Specifies the server executable to be used instead of the default server, `/usr/sbin/named`.

options/configuration_file

Specifies the configuration file to be used instead of the default, `/etc/named.conf`. A directory option might be specified in the configuration file. To ensure that reloading the configuration file continues to work in such a situation, *configuration_file* should be specified as an absolute pathname. This pathname should not include the *chroot_dir* pathname. This property is the equivalent of the `-c` option.

options/ip_interfaces

Specifies over which IP transport, IPv4 or IPv6, BIND will transmit. Possible values are IPv4 or IPv6. Any other setting assumes `all`, the default. This property is the equivalent of command line option `-4` or `-6`

options/listen_on_port

Specifies the default UDP and TCP port to be used for listening to DNS requests. This property is the equivalent of the command line option `-p port`.

options/debug_level

Specifies the default debug level. The default is 0, which means no debugging. The higher the number the more verbose debug information becomes. Equivalent of the command line option `-d debug_level`.

options/threads

Specifies the number of CPU worker threads to create. The default of 0 causes named to try and determine the number of CPUs present and create one thread per CPU. Equivalent of command line option `-n #cpus`.

options/chroot_dir

Specifies the directory to be used as the root directory after processing SMF properties and the command line arguments but before reading the configuration file. Use this property when using a [chroot\(2\)](#) environment. Synonymous to command line option `-t pathname`.

When using [chroot\(2\)](#), named is unable to disable itself when receiving [rndc\(1M\)](#) `stop` or `halt` commands. Instead, you must use the [svcadm\(1M\)](#) `disable` command.

In the event of a configuration error originating in one of the above SMF application options, named displays a message providing information about the error and the parameters that need correcting. The process then exits with exit code `SMF_EXIT_ERR_CONFIG`.

At startup, in the event of an error other than a configuration error, named exits with exit code `SMF_EXIT_ERR_FATAL`. Both of this code and `SMF_EXIT_ERR_CONFIG` cause the start method, [smf_method\(5\)](#), to place the service in the maintenance state, which can be observed with the [svcs\(1\)](#) command `svcs -x`.

In addition to the properties listed above, the following property can be used to invoke named as a user other than root:

start/user

Specifies the identity of the user that is invoking named. See [smf_method\(5\)](#) and [chroot\(2\)](#). Note that the user must have `solaris.smf.manage.bind` authorization. Without this role the named will be unable to manage its SMF FMRI and named will automatically be restarted by the SMF after an [rndc\(1M\)](#) stop or halt command. See EXAMPLES for a sequence of commands that establishes the correct authorization.

SIGNALS

In routine operation, signals should not be used to control the nameserver; [rndc\(1M\)](#) should be used instead.

SIGHUP

Force a reload of the server.

SIGINT, SIGTERM

Shut down the server.

The result of sending any other signals to the server is undefined.

Configuration

The named configuration file is too complex to describe in detail here. A list of configuration options is provided in the `named.conf` man page shipped with the BIND 9 distribution. A complete description is provided in the *BIND 9 Administrator Reference Manual*.

示例**示例 1** Configuring named to Transmit Only over IPv4 Networks

The following command sequence configures named such that it will transmit only over IPv4 networks.

```
# svccfg -s svc:network/dns/server:default setprop \  
> options/ip_interfaces=IPv4  
# svcadm refresh svc:network/dns/server:default  
#
```

示例 2 Listing Current Configuration File and Setting an Alternative File

The following sequence of commands lists the current named configuration file and sets an alternative file.

```
# svcprop -p options/configuration_file dns/server:default  
/etc/named.conf  
# svccfg -s dns/server:default setprop \  
> options/configuration_file=/var/named/named.conf  
# svcadm refresh dns/server:default  
# svcprop -p options/configuration_file dns/server:default  
/var/named/named.conf
```

示例 3 Establishing Appropriate Authorization for named

To have named start with the `solaris.smf.manage.bind` authorization, perform the steps shown below.

Add the user `dnsadmin` to the `solaris.smf.manage.bind` role:

示例 3 Establishing Appropriate Authorization for named (续)

```
# usermod -A solaris.smf.manage.bind dnsadmin
Observe effect of command:
# tail -1 /etc/user_attr
dnsadmin:::type=normal;auths=solaris.smf.manage.bind
```

Modify the service properties:

```
# svccfg
svc:> select svc:/network/dns/server:default
svc:/network/dns/server:default> setprop start/user = dnsadmin
svc:/network/dns/server:default> setprop start/group = dnsadmin
svc:/network/dns/server:default> exit
# svcadm refresh svc:/network/dns/server:default
# svcadm restart svc:/network/dns/server:default
```

Because only root has write access to create the default process-ID file, /var/run/named/named.pid, named must be configured to use an alternative path for the user dnsadmin. Here is an example of how to accomplish this:

```
# mkdir /var/named/tmp
# chown dnsadmin /var/named/tmp
```

Shown below is what you must add to named.conf to make use of the directory created above.

```
# head /etc/named.conf
options {
    directory "/var/named";
    pid-file "/var/named/tmp/named.pid";
};
```

/etc/named.conf
default configuration file

/var/run/named/named.pid
default process-ID file

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/dns/bind
Interface Stability	Volatile

[svcs\(1\)](#), [named-checkconf\(1M\)](#), [named-checkzone\(1M\)](#), [rndc\(1M\)](#), [rndc-confgen\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [svccprop\(1\)](#), [chroot\(2\)](#), [setuid\(2\)](#), [bind\(3SOCKET\)](#), [attributes\(5\)](#), [smf\(5\)](#), [smf_method\(5\)](#)

RFC 1033, RFC 1034, RFC 1035

文件

属性

另请参见

See the BIND 9 Administrator's Reference Manual. As of the date of publication of this man page, this document is available at <https://www.isc.org/software/bind/documentation>.

The `named.conf` man page shipped with the BIND 9 distribution

引用名	named-checkconf – named configuration file syntax checking tool
用法概要	named-checkconf [-h]vz] [-t <i>directory</i>] <i>filename</i>
描述	The named-checkconf utility checks the syntax, but not the semantics, of a specified configuration file.
选项	The following options are supported: -h Display the usage summary and exit. -j When loading a zonefile, read the journal if it exists. -t <i>directory</i> Change the root directory to <i>directory</i> so that include directives in the configuration file are processed as if run by a named configuration whose root directory has been similarly changed. -v Print the version of the named-checkconf program and exit. -z Perform a test load of the master zones found in <code>named.conf</code>.
操作数	The following operands are supported: <i>filename</i> The name of the configuration file to be checked. If not specified, it defaults to <code>/etc/named.conf</code>.
退出状态	0 No errors were detected. 1 An error was detected.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTE VALUE
Availability	service/network/dns/bind
Interface Stability	Volatile

另请参见	named(1M) , named-checkzone(1M) , attributes(5)
	See the BIND 9 Administrator's Reference Manual. As of the date of publication of this man page, this document is available at https://www.isc.org/software/bind/documentation .

引用名 named-checkzone, named-compilezone – zone file validity checking or converting tool

用法概要

```
named-checkzone [-Ddhjqv] [-c class] [-F format] [-f format]
                 [-i mode] [-k mode] [-M mode] [-m mode] [-n mode]
                 [-o filename] [-S mode] [-s style] [-t directory]
                 [-W mode] [-w directory] zonename filename

named-compilezone [-Ddhjqv] [-C mode] [-c class] [-F format]
                  [-f format] [-i mode] [-k mode] [-m mode] [-n mode]
                  [-o filename] [-s style] [-t directory]
                  [-W mode] [-w directory] zonename filename
```

描述

The named-checkzone utility checks the syntax and integrity of a zone file. It performs the same checks as [named\(1M\)](#) does when loading a zone. The named-checkzone utility is useful for checking zone files before configuring them into a name server.

named-compilezone is similar to named-checkzone, differing in that it always dumps the zone contents to a specified file in a specified format. Additionally, it applies stricter check levels by default, since the dump output will be used as an actual zone file loaded by [named\(1M\)](#). Unless manually specified otherwise, the check levels must be at least as strict as those specified in the named configuration file.

选项

For either or both utilities, the following options are supported:

- c *class*
Specify the class of the zone. If not specified, “IN” is assumed.
- D
Dump zone file in canonical format.
- d
Enable debugging.
- F *format*
Specify the format of the output file specified. Possible formats are text (default) and raw. For named-checkzone, this does not cause any effects unless it dumps the zone contents.
- f *format*
Specify the format of the zone file. Possible formats are text (default) and raw.
- h
Display usage message for named-checkzone.
- i *mode*
Perform post-load zone integrity checks. Possible modes are full (default), full-sibling, local, local-sibling, and none.

Mode full checks that MX records refer to the A or AAAA record (both in-zone and out-of-zone hostnames). Mode local checks only MX records that refer to in-zone hostnames.

Mode `full` checks that SRV records refer to the A or AAAA record (both in-zone and out-of-zone hostnames). Mode `local` checks only SRV records that refer to in-zone hostnames.

Mode `full` checks that delegation NS records refer to A or AAAA record (both in-zone and out-of-zone hostnames). It also checks that glue address records in the zone match those advertised by the child. Mode `local` checks only NS records that refer to in-zone hostnames or check that some required glue exists, that is, when the nameserver is in a child zone.

Mode `full-sibling` and `local-sibling` disable sibling glue checks, but are otherwise the same as `full` and `local`, respectively.

Mode `none` disables the checks.

`-k mode`

Perform “check-name” checks with the specified failure mode. Possible modes are `fail` (default for `named-compilezone`), `warn` (default for `named-checkzone`) and `ignore`.

`-j`

Read the journal, if it exists, when loading the zone file.

`-M mode`

Check if an MX record refers to a CNAME. Possible modes are `fail`, `warn` (default) and `ignore`.

`-m mode`

Specify whether MX records should be checked to see if they are addresses. Possible modes are `fail`, `warn` (default) and `ignore`.

`-n mode`

Specify whether NS records should be checked to see if they are addresses. Possible modes are `fail` (default for `named-compilezone`), `warn` (default for `named-checkzone`) and `ignore`.

`-o filename`

Write zone output to *filename*. If *filename* is `-` (a hyphen), then write to standard out. The hyphen mandatory for `named-compilezone`

`-q`

Run in quiet mode, reporting only the exit status.

`-S mode`

Check if a SRV record refers to a CNAME. Possible modes are `fail`, `warn` (default) and `ignore`.

`-s style`

Specify the style of the dumped zone file. Possible styles are `full` (default) and `relative`. The `full` format is most suitable for processing automatically by a separate script. The `relative` format is more human-readable and is thus suitable for editing by hand. For

named - checkzone this option does not cause any effects unless it dumps the zone contents. It also has no effect if the output format is not text.

- t *directory*
chroot to *directory* so that include directives in the configuration file are processed as if run by a similarly chrooted named.
- v
Print the version of the named - checkzone program and exit.
- W *mode*
Specify whether to check for non-terminal wildcards. Non-terminal wildcards are almost always the result of a failure to understand the wildcard matching algorithm (RFC 1034). Possible modes are warn (default) and ignore.
- w *directory*
chdir to *directory* so that relative filenames in master file \$INCLUDE directives work. This is similar to the directory clause in named.conf.

操作数 The following operands are supported:

- filename*
The name of the zone file.
- zonename*
The domain name of the zone being checked.

- 退出状态
- 0
No errors were detected.
 - 1
An error was detected.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/dns/bind
Interface Stability	Volatile

另请参见 [named\(1M\)](#), [named-checkconf\(1M\)](#), [attributes\(5\)](#)

RFC 1035

See the BIND 9 Administrator's Reference Manual. As of the date of publication of this man page, this document is available at <https://www.isc.org/software/bind/documentation>.

引用名	ncaconfd – Solaris Network Cache and Accelerator (NCA) configuration daemon
用法概要	<code>/usr/lib/inet/ncaconfd [-al] interface1 [interface2 ...]</code>
描述	Use the <code>ncaconfd</code> utility to set up NCA on a system. At boot time, the <code>ncakmod</code> initialization script reads in nca.if(4) to determine on which interface(s) NCA should run. <code>ncaconfd</code> then sets up the interface. <code>ncaconfd</code> also operates as a daemon if the <code>nca_active</code> key is set to enabled in ncakmod.conf(4) file. In this case, <code>ncaconfd</code> will continue as a daemon after all the NCA interfaces have been set up, listening for routing changes. The changes are then passed to NCA to control which interface NCA should use to make active outgoing TCP connections.
选项	The following options are supported: -a Enable active connections. -l Enable logging.
文件	<code>/etc/nca/ncakmod.conf</code>
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	system/network/http-cache-accelerator
Interface Stability	Committed

另请参见 [ncakmod\(1\)](#), [nca.if\(4\)](#), [ncakmod.conf\(4\)](#), [attributes\(5\)](#), [nca\(7d\)](#)

引用名	ncheck – generate a list of path names versus i-numbers
用法概要	<pre>ncheck [-F FSType] [-V] [generic_options] [-o FSType-specific_options] [special] . . .</pre>
描述	ncheck with no options generates a path-name versus i-number list of all files on <i>special</i> . If <i>special</i> is not specified on the command line the list is generated for all <i>specials</i> in <i>/etc/vfstab</i> which have a numeric <i>fsckpass</i> . <i>special</i> is the raw device on which the file system exists.
选项	-F Specify the <i>FSType</i> on which to operate. The <i>FSType</i> should either be specified here or be determinable from <i>/etc/vfstab</i> by finding an entry in the table that has a numeric <i>fsckpass</i> field and an <i>fsckdev</i> that matches <i>special</i>. -V Echo the complete command line, but do not execute the command. The command line is generated by using the options and arguments provided by the user and adding to them information derived from <i>/etc/vfstab</i>. This option may be used to verify and validate the command line. generic_options Options that are commonly supported by most <i>FSType</i>-specific command modules. The following options are available: -i i-list Limit the report to the files on the <i>i-list</i> that follows. The <i>i-list</i> must be separated by commas with no intervening spaces. -a Print the names “.” and “. .” which are ordinarily suppressed. -s Report only special files and files with set-user-ID mode. This option may be used to detect violations of security policy. -o Specify <i>FSType-specific_options</i> in a comma separated (without spaces) list of suboptions and keyword-attribute pairs for interpretation by the <i>FSType</i>-specific module of the command.
用法	See largefile(5) for the description of the behavior of ncheck when encountering files greater than or equal to 2 Gbyte (2 ³¹ bytes).
文件	<i>/etc/vfstab</i> list of default parameters for each file system
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [vfstab\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#) Manual pages for the *FSType*-specific modules of ncheck

附注

This command may not be supported for all *FSTypes*.

引用名

ncheck_ufs – generate pathnames versus i-numbers for ufs file systems

用法概要

ncheck -F ufs [*generic_options*] [-o m] [*special*]...

描述

ncheck -F ufs generates a pathname versus i-number list of files for the ufs file system residing on *special*. Names of directory files are followed by ‘/.’.

选项

See [ncheck\(1M\)](#) for the list of *generic_options* supported.

-o Specify ufs file system specific options. The available option is:

 m Print mode information.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[ff\(1M\)](#), [ncheck\(1M\)](#), [attributes\(5\)](#)

诊断

When the file system structure is improper, ‘??’ denotes the “parent” of a parentless file and a pathname beginning with ‘. . .’ denotes a loop.

引用名 ndd – get and set driver configuration parameters

用法概要 ndd [-set] *driver parameter* [*value*]

描述

ndd gets and sets selected configuration parameters in some kernel drivers. Currently, ndd only supports the drivers that implement the TCP/IP Internet protocol family. Each driver chooses which parameters to make visible using ndd. Since these parameters are usually tightly coupled to the implementation, they are likely to change from release to release. Some parameters may be read-only.

注 – It is strongly encouraged that you use [ipadm\(1M\)](#), rather than ndd, to modify or retrieve TCP/IP Internet protocols. The current ndd command will be made obsolete in a future release, replaced by [ipadm\(1M\)](#). Please see NOTES for more information.

If the ndd -set option is omitted, ndd queries the named *driver*, retrieves the value associated with the specified *parameter*, and prints it. If the -set option is given, ndd passes *value*, which must be specified, down to the named *driver* which assigns it to the named *parameter*.

By convention, drivers that support ndd also support a special read-only *parameter* named “?” which can be used to list the parameters supported by the driver.

示例

示例 1 Getting Parameters Supported By The TCP Driver

To see which parameters are supported by the TCP driver, use the following command:

```
example% ndd /dev/tcp \?
```

The parameter name “?” may need to be escaped with a backslash to prevent its being interpreted as a shell meta character.

The following command sets the value of the parameter *ip_forwarding* in the dual stack IP driver to zero. This disables IPv4 packet forwarding.

```
example% ndd -set /dev/ip ip_forwarding 0
```

Similarly, in order to disable IPv6 packet forwarding, the value of parameter *ip6_forwarding*

```
example% ndd -set /dev/ip ip6_forwarding 0
```

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [dladm\(1M\)](#), [ipadm\(1M\)](#), [ioctl\(2\)](#), [attributes\(5\)](#), [nca\(7d\)](#), [arp\(7P\)](#), [ip\(7P\)](#), [ip6\(7P\)](#), [tcp\(7P\)](#), [udp\(7P\)](#)

附注

The parameters supported by each driver may change from release to release. Like programs that read `/dev/kmem`, user programs or shell scripts that execute `nnd` should be prepared for parameter names to change.

The `ioctl()` command that `nnd` uses to communicate with drivers is likely to change in a future release. User programs should avoid making dependencies on it.

The use of `nnd` to administer Layer 2 (Data Link layer) drivers is strongly discouraged as this capability is to be obsoleted in a future release, replaced by [dladm\(1M\)](#). Please refer to the driver-specific man page in section 7D of the SunOS man pages.

The use of `nnd` to administer the drivers that implement the TCP/IP Internet protocol family (IP/TCP/SCTP/UDP/ICMP) is strongly discouraged as this capability is to be obsoleted in a future release, replaced by `ipadm`. Please see [ipadm\(1M\)](#) for instructions for modifying and retrieving supported protocol properties.

The meanings of many `nnd` parameters make sense only if you understand how the driver is implemented.

引用名	ndmpadm – 管理网络数据管理协议活动
用法概要	<code>/usr/sbin/ndmpadm [-?] subcommand [options] [direct-object]</code>
描述	可使用 <code>ndmpadm</code> 命令查询 ndmpd(1M) 守护进程，以获取活动会话的状态，终止会话，查询备份设备以及设置或获取当前 NDMP（Network Data Management Protocol，网络数据管理协议）服务变量和属性。<code>ndmpadm</code> 是作为一组子命令实现的（其中很多子命令都有自己的直接对象），下面将分别对这些子命令进行说明。某些子命令支持选项，这些选项将与子命令一起介绍。 <code>ndmpadm</code> 命令支持以下子命令： disable 禁用指定的验证口令处理。 enable 启用指定的验证口令处理。 get 获取 NDMP 配置属性值。 kill-sessions 终止活动会话。 set 设置 NDMP 配置属性的值。 show-devices 获取与服务器连接的磁带设备列表。 show-sessions 显示活动 NDMP 会话的详细信息。 选项 支持以下选项： -? 显示所有子命令和选项的列表。 子命令 <code>ndmpadm</code> 命令支持下述子命令。 disable 子命令 <code>disable</code> 子命令的语法如下： <pre># ndmpadm disable -a auth-type</pre> 该子命令为 NDMP 客户机的远程访问禁用 <i>auth-type</i> 指定的验证类型。<i>auth-type</i> 的有效值为 <code>cram-md5</code> 或 <code>cleartext</code>。 enable 子命令 <code>enable</code> 子命令的语法如下： <pre># ndmpadm enable -a auth-type -u username</pre> 该子命令提示输入两次用户口令以进行确认，并使用 NDMP 客户机访问的给定用户名和口令激活指定的验证类型。<i>auth-type</i> 的有效值为 <code>cram-md5</code> 或 <code>cleartext</code>。

get 子命令

get 子命令的语法如下：

```
# ndmpadm get [-p] [property] [[-p] property=value]...
```

属性名与用于 set 子命令的属性名相同，如下所述。如果未指定属性，get 子命令将返回所有配置属性。

kill-sessions 子命令

可以使用 kill-sessions 子命令终止会话编号 *ID*。

kill-sessions 子命令的语法如下：

```
# ndmpadm kill-sessions ID
```

set 子命令

set 子命令的语法如下：

```
# ndmpadm set [-p] property=value [[-p] property=value]...
```

[ndmp\(4\)](#) 手册页中介绍了可使用 set 子命令设置的属性。

show-devices 子命令

show-devices 子命令的语法如下：

```
# ndmpadm show-devices
```

该子命令可列出当前连接到系统的磁带驱动器和库的名称、供应商、序列号以及其他信息。

show-sessions 子命令

show-sessions 子命令显示会话的详细信息。show-sessions 子命令的语法如下：

```
# ndmpadm show-sessions [-i tape,scsi,data,mover] [ID]
```

show-sessions 子命令支持以下参数：

-i *tape,scsi,data,mover*

标识要获取其数据的接口的类型。如果未指定接口，show-sessions 将显示所有类型的接口的信息。

ID

标识要显示其数据的特定会话。如果未指定 *ID*，show-sessions 将显示所有会话的数据。

示例

示例1 获取所有 NDMP 连接的状态

以下命令将获取所有连接的状态。

```
# ndmpadm show-devices
```

示例2 获取特定类型的连接的状态

以下命令将获取磁带和 SCSI 接口的状态。

```
# ndmpadm show-sessions -i scsi,tape
```

示例3 限制协议版本

以下命令将限制 NDMP 协议使用版本 3。

```
# ndmpadm set -p version=3
```

示例4 获取当前版本号

以下命令将获取当前正在运行的 NDMP 的版本号。

```
# ndmpadm get -p version
```

示例5 断开特定会话

以下命令将断开会话 5。会话编号是之前通过 ndmpadm show-sessions 命令获取的。

```
# ndmpadm kill-session 5
```

示例6 获取所有 NDMP 属性的值

以下命令将获取所有 NDMP 属性的值。

```
# ndmpadm get
```

示例7 启用 CRAM-MD5 验证

以下命令将启用 CRAM-MD5 验证。

```
# ndmpadm enable -a cram-md5 -u admin
Enter new password:*****
Re-enter password:*****
```

示例8 禁用明文口令验证

以下命令将禁用明文口令验证。

```
# ndmpadm disable -a cleartext
```

退出状态

- 0 成功完成。
- 1 出现错误（如 ndmpd 守护进程未在运行），使得 ndmpadm 无法联系守护进程。
- 2 指定的命令行选项无效。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/storage/ndmp
接口稳定性	Committed（已确定）

另请参见

[dump\(1\)](#)、[tar\(1\)](#)、[ndmpd\(1M\)](#)、[ndmpstat\(1M\)](#)、[svccfg\(1M\)](#)、[syslogd\(1M\)](#)、[ndmp\(4\)](#)、[attributes\(5\)](#)

附注

[ndmpd\(1M\)](#) 守护进程由服务管理工具 ([smf\(5\)](#)) 管理，其服务标识符为：

```
svc:/system/ndmpd
```

还可使用 [svccfg\(1M\)](#) 对此服务执行管理操作，如设置和获取属性。例如，启用直接访问恢复 (Direct Access Recovery, DAR) 模式：

```
# svccfg -s svc:/system/ndmpd
```

```
svc:/system/ndmpd> setprop ndmpd/dar-support = yes
```

...以及获取属性列表：

```
# svccfg -s svc:/system/ndmpd
```

```
svc:/system/ndmpd> listprop
```

引用名 ndmpd – 网络数据管理协议守护进程

用法概要 /usr/lib/ndmp/ndmpd

描述 ndmpd 守护进程处理客户机网络数据管理协议 (Network Data Management Protocol, NDMP) 请求。NDMP 是一种基于网络的企业级开放数据管理协议，用于备份和恢复。通过 ndmpd 守护进程，用户可使用数据管理应用程序 (Data Management Application, DMA) 客户机管理数据备份和恢复。NDMP 协议用于协调 DMA 与 NDMP 服务器之间或两个 NDMP 服务器之间的数据移动和控制。

缺省情况下，ndmpd 处于禁用状态。

- 退出状态**
- 0 成功完成。
 - 1 出现使 ndmpd 守护进程无法进行初始化的错误，例如无法派生一个互斥锁初始化进程。
 - 2 指定的命令行选项无效。

文件

/usr/lib/ndmp/ndmpd
网络数据管理协议服务器二进制文件。

/var/ndmp/ndmp.log
网络数据管理协议日志消息文件。重新引导时将删除此文件。

/var/ndmp/dumpdates
存储有关转储备份的日期和级别信息的文本文件。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/storage/ndmp
接口稳定性	Committed（已确定）

另请参见 [svcs\(1\)](#)、[ndmpadm\(1M\)](#)、[svcadm\(1M\)](#)、[svccfg\(1M\)](#)、[syslogd\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

附注 ndmpd 守护进程由服务管理工具 ([smf\(5\)](#)) 管理，其服务标识符为：

svc:/system/ndmp:default

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

引用名	ndmpstat – 显示 NDMP 备份进度统计信息
用法概要	ndmpstat [tapes] [interval [count]]
描述	ndmpstat 实用程序报告网络数据管理协议 (Network Data Management Protocol, NDMP) 统计信息，其中包括 NDMP 工作线程、磁盘 IO、磁带 IO、文件操作、性能以及备份活动。 ndmpstat 报告所有磁带和磁盘的汇总统计信息。要获取特定磁带设备的统计信息，应将磁带设备名称作为参数传递给该实用程序。 由于调用了 NDMP 服务，ndmpstat 在调用时会首先显示一行 NDMP 守护进程活动摘要。
显示字段	ndmpstat 输出中的字段如下所述： wthr 报告以下四种状态中每个状态的工作线程数： - r 正在运行的工作线程数 - w 正在等待资源（如 I/O 和换页）的被阻塞的工作线程数 - b 当前正在运行的备份操作数 - r 当前正在运行的恢复操作数 file 报告文件系统的使用情况。 - rd 正在读取的文件数 - wr 正在写入的文件数 disk 报告每间隔的磁盘操作数。 - rd 正在读取的磁盘块数 - wr 正在写入的磁盘块数 tape 报告每间隔的磁带操作数。插槽中最多可放入四盘磁带，并标有一位数字编号。该编号表示 /dev/rmt 下设备的名称。

rd
正在读取的磁带块数

wr
正在写入的磁带块数

bytes
报告已传输的字节数。这是磁带和磁盘设备的总值。数值以千字节为单位。

rd
正在读取的千字节数

wr
正在写入的千字节数

perf
显示粗略估计的备份/恢复操作的性能（以兆字节/秒为单位）。

bk
备份性能

rs
恢复性能

prcnt
以百分比形式显示资源的比较使用情况。

dsk
磁盘 I/O 时间

tpe
磁带 I/O 时间

otr
其他时间（内存或空闲）

请参见“示例”部分。

操作数

支持下列操作数：

- count* 指定重复统计信息显示的次数。
- tape* 指定在输出中给定优先级的磁带。命令行限制为最多四个磁带设备。常用磁带名称为 */dev/rmt/n*，其中 *n* 为整数。
- interval* 指定 ndmpstat 汇总其间活动的秒数。终止命令之前，指定的间隔始终有效。

示例

示例1 使用 ndmpstat
以下命令每五秒钟显示一次 NDMP 守护进程活动的汇总信息。

示例1 使用 ndmpstat (续)

```
example% ndmpstat 5
wthr      file      disk      tape      bytes      perf      prcnt
r w b r rd  wr    rd  wr    rd  wr    rd  wr    bk  rs  dsk  tpe  otr
1 0 3 6 50   9  1250  0  32544 4455  42335 3234   5  4   20   40   40
1 0 0 1 1   0   128  0    0   128   64   64   1  0    0   80   20
1 0 0 1 2   0   128  0    0    0    64    0   1  0   80    0   20
1 0 0 1 1   0   128  0    0    0    64    0   1  0   80    0   20
1 0 0 1 3   0   128  0    0    0    64    0   0  0   80    0   20
1 0 0 1 1   0   128  0    0   128   64   64   1  0    0   80   20
^C
example%
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/storage/ndmp
接口稳定性	请参见下文

调用正在发展中。用户可读的输出不稳定。

另请参见 [iostat\(1M\)](#)、[mpstat\(1M\)](#)、[ndmpd\(1M\)](#)、[ndmpadm\(1M\)](#)、[attributes\(5\)](#)

附注 性能数字不精确，向上舍入为 MB/sec. 边界值。

引用名 netadm – 管理网络配置文件

用法概要

```
netadm enable [ -p profile-type ] [ -c ncu-class ] profile-name

netadm disable [ -p profile-type ] [ -c ncu-class ] profile-name

netadm list [ -x ] [ -p profile-type ] [ -c ncu-class ]
           [ profile-name ]

netadm show-events

netadm scan-wifi linkname

netadm select-wifi linkname

netadm help
```

描述

netadm 实用程序用于管理网络配置文件以及与 NWAM 守护进程交互。

有三种类型的网络配置文件：网络配置文件 (Network Configuration Profile, NCP)、位置和外部网络修饰器 (External Network Modifier, ENM)。

在任何给定的时间，系统上都有一个活动的 NCP 和一个活动的位置配置文件。（通过激活模式 `manual` 启用不同的 NCP 或位置配置文件将隐式禁用当前活动的 NCP 或位置配置文件。也可以禁用当前位置配置文件（如果其激活模式为 `manual`），但这样做会“关闭”系统联网功能的某些方面（如名称服务）。不允许显式禁用 NCP，因为这会切实关闭系统的基本网络连接。仅当启用不同的 NCP 时才会隐式禁用原来的 NCP。

与此相反，在任何给定的时间，可以不存在任何 ENM 或存在多个 ENM。因此，启用或禁用一个 ENM 对其他活动的 ENM 没有任何影响。

允许启用和禁用单个 NCU；指定的 NCU 必须是当前活动 NCP 的一部分，并且必须将其激活模式设置为 `manual`。如果未指定 NCU 类，将启用或禁用具有给定名称的所有 NCU（一个链接和/或一个接口）。

对象的启用和禁用以异步方式执行。因此，可能会出现启用或禁用请求成功，但操作本身失败的情况。此类失败将反映在对象状态中；`maintenance` 状态表示所执行的最后一项操作失败。

子命令

支持以下子命令：

```
enable [ -p profile-type ] [ -c ncu-class ] profile-name
```

启用指定的配置文件。如果配置文件名称不唯一，则必须指定配置文件类型以标识要启用的配置文件。如果配置文件类型为 NCU 并且名称不唯一（即，链接 NCU 和接口 NCU 同名），将启用这两个 NCU，除非使用 `-c` 选项指定 NCU 类。配置文件类型必须是 `ncp`、`ncu`、`loc` 或 `enm` 其中之一；NCU 类必须是 `phys` 或 `ip`。

```
disable [ -p profile-type ] [ -c ncu-class ] profile-name
```

禁用指定的配置文件。如果配置文件名称不唯一，则必须指定配置文件类型以标识要禁用的配置文件。如果配置文件类型为 NCU 并且名称不唯一（即，链接 NCU 和接口 NCU 同名），将禁用这两个 NCU，除非使用 `-c` 选项指定 NCU 类。配置文件类型必须是 `ncu`、`loc` 或 `enm` 其中之一；NCU 类必须是 `phys` 或 `ip`。

```
list [-x] [-p profile-type] [-c ncu-class] [profile-name]
```

列出所有可用配置文件及其当前状态。如果按名称指定某个特定配置文件，将仅列出该配置文件的当前状态。如果配置文件名称不唯一，将列出具有给定名称的所有配置文件；或者，也可以指定配置文件类型和/或 NCU 类来标识特定配置文件。如果仅提供一种类型，将列出该类型的所有配置文件。列出活动的 NCP 将包括构成该 NCP 的 NCU。

-x 选项可使 list 子命令在缺省显示的前三列标题（TYPE、PROFILE 和 STATE）之后，显示输出的第四列，该列以 AUXILIARY STATE 为标题。AUXILIARY STATE 列显示配置文件处于给定状态的原因。

可能的 STATE 值为：

disabled (禁用)

需要手动激活的配置文件尚未激活。

offline (脱机)

需要根据条件激活或由系统激活的配置文件尚未激活。该配置文件之所以未处于活动状态，可能是因为条件未满足；也可能是因为另一个配置文件具有更多已满足的特定条件而被激活（对于一次必须激活一个配置文件的配置文件类型可能会发生这种情况，如位置配置文件）。

online (联机)

需要根据条件激活或由系统激活的配置文件的条件已满足，并已成功激活；或者需要手动激活的配置文件在用户请求下已成功激活。

maintenance (维护)

试图激活配置文件，但失败。

initialized (已初始化)

配置文件表示一个尚未对其采取任何操作的有效配置对象。

uninitialized (未初始化)

配置文件表示一个系统中不存在的配置对象；例如，对应于已删除的物理链接的 NCU。

show-events

从 NWAM 守护进程侦听事件流并显示它们。

scan-wifi *linkname*

在链接 *linkname* 上启动无线扫描。

select-wifi *linkname*

从链接 *linkname* 上的扫描结果中选择一个要连接的无线网络。如有必要，可提示进行选择、提供 WiFi 密钥等。

help

显示其中包含每个子命令的简短描述的用法消息。

示例

示例1 启用用户指定的位置配置文件
以下命令将启用一个用户指定的位置。

```
# netadm enable -p loc office
Disabled loc 'home'.
Enabled loc 'office'
```

示例2 禁用 ENM
以下命令将禁用某个 ENM。

```
# netadm disable -p enm myvpn
Disabled enm 'myvpn'.
```

示例3 列出所有 NCP
以下命令将列出所有 NCP。

```
# netadm list -xp ncp
TYPE          PROFILE      STATE      AUXILIARY STATE
ncp           Automatic   disabled   disabled by administrator
ncp           User        online     active
ncu:phys      nge0        online     interface/link is up
ncu:ip        nge0        online     interface/link is up
ncu:phys      nge1        offline    interface/link is down
ncu:ip        nge1        offline    conditions for activation are
unmet
```

示例4 列出活动 NCP 中的 NCU
以下命令将列出活动 NCP 中的所有 ip NCU。

```
# netadm list -c ip
TYPE          PROFILE      STATE
ncu:ip        bge0        online
ncu:ip        bge1        disabled
```

示例5 强制执行扫描
以下命令将在无线接口 wpi0 上强制执行扫描。

```
# netadm scan-wifi wpi0
```

示例6 选择 WiFi 网络
以下命令将选择正在广播其 ESSID 的 WiFi 网络。

```
# netadm select-wifi wpi0
1: ESSID testing BSSID 0:40:96:29:e9:d8
2: ESSID sunwifi BSSID 0:b:e:9f:b5:80
3: ESSID sunwifi BSSID 0:b:e:85:26:c0
4: ESSID sunwifi BSSID 0:b:e:49:2f:80
```

示例6 选择 WiFi 网络 (续)

5: Other

Choose WLAN to connect to [1-5]: **2**

#

示例7 选择 WiFi 网络 (另一种情况)

以下命令将选择未在广播其 ESSID 的 WiFi 网络。

```
# netadm select-wifi wpi0
```

```
1: ESSID testing BSSID 0:40:96:29:e9:d8
```

```
2: ESSID sunwifi BSSID 0:b:e:85:26:c0
```

```
3: ESSID sunwifi BSSID 0:b:e:9f:b5:80
```

```
4: ESSID sunwifi BSSID 0:b:e:49:2f:80
```

```
5: ESSID sunwifi BSSID 0:b:e:49:62:c0
```

```
6: Other
```

Choose WLAN to connect to [1-6]: **6**

Enter WLAN name: **oraclewifi**

```
1: None
```

```
2: WEP
```

```
3: WPA
```

Enter security mode: **2**

Enter WLAN key for ESSID oraclewifi: **123456**

Enter key slot [1-4]: **1**

#

示例8 监视 nwamd

在切换位置时，以下命令将监视 [nwamd\(1M\)](#)。

```
# netadm show-events
```

```
EVENT          DESCRIPTION
OBJECT_ACTION  loc Automatic -> action refresh
OBJECT_STATE   loc Automatic -> state offline*, method/service executi
OBJECT_STATE   loc Automatic -> state online, active
OBJECT_ACTION  loc home -> action refresh
OBJECT_ACTION  loc NoNet -> action refresh
OBJECT_ACTION  loc User -> action refresh
OBJECT_ACTION  loc home -> action enable
OBJECT_STATE   loc home -> state offline*, method/service executing
OBJECT_STATE   loc Automatic -> state offline, conditions for activati
OBJECT_STATE   loc home -> state online, active
^C
#
```

属性 有关下列属性的描述，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）

另请参见 [netcfgd\(1M\)](#)、[netcfg\(1M\)](#)、[nwamd\(1M\)](#)、[attributes\(5\)](#)
另请参见 JDS/GNOME 手册页集合中的 [nwam-manager\(1M\)](#)。

引用名 netcfg – 创建和修改网络配置文件

用法概要

```
netcfg
netcfg subcommand [options...]
netcfg [-d] -f command-file
netcfg help [subcommand]
```

描述 netcfg 实用程序处理系统网络配置文件。可以通过单个子命令或通过指定一个包含一系列子命令的命令文件，以交互方式调用 netcfg。

netcfg 实用程序可对几种不同类型的配置文件进行操作：

- 网络配置文件 (NCP)
- 位置
- 外部网络修饰符 (ENM)
- 已知的 WLAN

有关这些配置文件类型的更多详细信息，请参阅“配置文件”部分。

netcfg 命令在作用域内执行。作用域有以下三种：全局、配置文件和 NCP。如果不带任何参数调用 netcfg，则会在全局作用域内开始编辑会话。在全局作用域内，NCP、位置和 ENM 配置文件以及已知的 WLAN 项均可用于操作。选择 NCP 将使编辑会话移到 NCP 作用域；在该作用域中，可以创建或选择单个 NCU 以移到配置文件作用域。此外，如果在全局作用域中选择或创建位置、ENM 或已知的 WLAN，将使编辑会话移到配置文件作用域。

在给定的配置文件作用域中，可查看和修改配置文件属性。

在交互模式下，调用提交之前不会将更改存储到持久性存储器。提交会在“结束”或“退出”时隐式调用，也可由用户显式调用。在调用提交时，将提交整个配置文件。为了维护持久性存储器的一致性，提交操作中加入了一个验证步骤；如果验证失败，则提交也会失败。如果隐式提交失败，则将为用户提供结束或退出的选项而不提交当前更改，或者保留在当前作用域内进行进一步更改。

配置文件 NWAM 服务将所需的属性值存储在配置文件中，以此管理网络配置。接着，该服务根据当前网络条件确定应在给定时间激活哪个配置文件，然后激活该配置文件。除了上一节中所讨论的网络配置文件 (NCP) 之外，nwamd 还管理位置和 ENM 配置文件。

网络配置文件 (NCP) NCP 指定本地网络组件（包括物理链路、IP 隧道链路和 IP 接口）的配置。IP 接口必须与任意类型的基础链路相关联。这些组件统称为网络配置单元 (Network Configuration Unit, NCU)。

可在系统上配置任意数量的 NCP。自动 NCP 由 nwamd 创建和管理，不能由用户修改。此 NCP 由一个链路 NCU 和一个接口 NCU 组成，它们均表示系统中的每个物理链路。当在系统中添加或删除链路时，也将在自动 NCP 中添加或删除它们的相应 NCU。通过此 NCP 实现的策略将选择有线链路而不是无线链路，并检查所有已连接有线链路上的 IP；如果未连接有线链路，该策略将选择一个无线链路。

在从早期版本的 NWAM 服务升级后，系统还会根据旧 `/etc/nwam/llp` 文件中找到的配置创建用户 NCP。在大多数情况下，如果 NWAM 服务在升级时处于活动状态，则最初是自动 NCP 处于活动状态（即使已创建用户 NCP）。但是，如果在旧配置中检测到显式更改（如静态 IP 地址更改），则最初是用户 NCP 处于活动状态。

最后，用户可以创建任意数量的附加 NCP。这些 NCP 完全由用户管理；必须显式添加或删除 NCU，可以添加未映射到当前已安装到系统中的任何链路的 NCU，也可以删除映射到系统中的链路的 NCU。用户可以确定适用于这些 NCP 的策略。

在进行初始 NCP 选择后，系统不会更改活动 NCP。用户可以随时使用 GUI 或 [netadm\(1M\)](#) 命令执行此操作。

位置

位置指定系统范围的网络配置，包括名称服务、域、IP 过滤器和 IPsec 配置。

外部网络修饰符 (ENM)

顾名思义，外部网络修饰符是 NWAM 服务的外部应用程序，可用于修改和/或创建网络配置。`nwamd` 可根据 ENM 配置文件中指定的条件来激活或取消激活 ENM。或者，用户可根据需要选择手动激活/取消激活 ENM。

当位置配置文件允许根据当前网络条件自动配置一组特定的网络相关服务的同时，也限制了该组服务。ENM 提供附加的灵活性，允许用户指定对 SMF 服务属性和/或状态或要在特定条件下应用的任何其他系统设置的更改。

属性

`netcfg` 支持以下类型的属性：

- NCU 属性
- 接口 NCU 的属性
- 所有链路 NCU 的通用属性
- 位置属性
- ENM 属性
- 已知的 WLAN 属性

以下小节中将介绍这些属性。

NCU 属性

以下属性是所有 NCU 的通用属性。

type：枚举值：`link` | `interface`

指定 NCU 类型（`link` 或 `interface`）。根据指定的类隐式确定值。

class：枚举值：`phys`（适用于链路 NCU）；`ip`（适用于接口 NCU）

指定 NCU 类。

parent：字符串：`name of parent NCP`

指定此 NCU 所属的 NCP。

type、**class** 和 **parent** 属性是在创建 NCU 时设置的，并且以后无法更改。

enabled：布尔值：`true` | `false`

如果 `activation-mode` 是 `manual`，则 `enabled` 属性将反映 NCU 的当前状态。此属性是只读属性；通过使用 [netadm\(1M\)](#) 启用或禁用 NCU，可间接更改此属性。

	缺省值为 <code>true</code> 。
接口 NCU 的属性	<code>ip-version</code>: 枚举值列表: <code>ipv4 ipv6</code> 在此 NCU 上应使用的 IP 版本。 缺省值为 <code>ipv4,ipv6</code>。 <code>ipv4-addrsrc</code>: 枚举值列表: <code>dhcp static</code> 标识分配给此 NCU 的 IPv4 地址的源; 可分配多个值。如果分配的值之一是 <code>static</code>, 则 <code>ipv4-addr</code> 属性必须包括至少一个要分配给 NCU 的 IPv4 地址。 缺省值为 <code>dhcp</code>。 <code>ipv4-addr</code>: IPv4 地址列表 一个或多个要分配给此 NCU 的 IPv4 地址。 <code>ipv4-default-route:IPv4 address</code> 缺省路由器的 IPv4 地址; 如果设置了此属性, 则在启动此接口时, IPv4 通信的缺省路由将与此接口关联。 <code>ipv6-addrsrc</code>: 枚举值列表: <code>dhcp autoconf static</code> 标识分配给此 NCU 的 IPv6 地址的源; 可分配多个值。如果分配的值之一是 <code>static</code>, 则 <code>ipv6-addr</code> 属性必须包括至少一个要分配给 NCU 的 IPv6 地址。 缺省值为 <code>dhcp,autoconf</code>。 <code>ipv6-addr</code>: IPv6 地址列表 一个或多个要分配给此 NCU 的 IPv6 地址。 <code>ipv6-default-route:IPv6 address</code> 缺省路由器的 IPv6 地址; 如果设置了此属性, 则在启动此接口时, IPv6 通信的缺省路由将与此接口关联。
所有链路 NCU 的通用属性	<code>activation-mode</code>: 枚举值: <code>manual prioritized</code> 用于自动激活此 NCU 的触发器类型。 缺省值为 <code>manual</code>。 <code>priority-group:uint64:group priority number</code> 如果 <code>activation-mode</code> 设置为 <code>prioritized</code>, 该属性将指定此 NCU 所属的优先级组。组由一个或多个 NCU 组成; 数字越小, 优先级越高。确定为可用的最高优先级组将通过 nwamd(1M) 激活, 并且一直保持到不再可用为止, 或保持到更高优先级组变为可用为止。 <code>priority-mode</code>: 枚举值: <code>exclusive shared all</code> 如果 <code>activation-mode</code> 设置为 <code>prioritized</code>, 此属性将指定用于确定优先级组的可用性和激活行为的模式。 <code>exclusive</code> 至少必须有一个 NCU 可用, 才能使组可用, 并且只会激活一个 NCU。如果有多个成员 NCU 可用, nwamd(1M) 将随机选择一个进行激活。

shared

至少必须有一个 NCU 可用，才能使组可用；将激活所有可用的 NCU。

all

所有组成员 NCU 均必须可用，才能使组可用；将激活所有成员 NCU。

link-mac-addr: 字符串: 48 位 MAC 地址

分配给此链路的 MAC 地址。缺省情况下，NWAM 将要求链路的出厂分配 MAC 地址或缺省 MAC 地址；在此设置一个不同的值可覆盖该选项。

link-autopush: 字符串列表: 要在链路上推动的模块

标识链路打开时应在其上自动推动的一系列模块。

link-mtu:uint64: 此链路的 MTU 大小

此属性将被自动设置为物理链路的缺省 MTU；将此属性设置为不同的值可覆盖该值。

请注意，允许的 MTU 值的范围取决于底层硬件。由于在提供附带有驱动程序的底层硬件之前可以创建 NCU，因此无法验证在编辑 NCU 时设置的值。如果由于 MTU 大小无效而未能激活 NCU，则会在系统日志中指明这种情况，并且 NCU 将被置于维护状态。

位置属性

activation-mode: 枚举值: `manual | system | conditional-all | conditional-any`

用于自动激活此位置的触发器类型。

缺省值为 `manual`。

enabled: 布尔值: `true | false`

如果 `activation-mode` 是 `manual`，则 `enabled` 属性将反映位置的当前状态。此属性是只读属性；通过使用 [netadm\(1M\)](#) 启用或禁用位置，可间接更改此属性。

缺省值为 `false`。

conditions: 字符串列表: *conditional expressions*

如果 `activation-mode` 设置为 `conditional-all` 或 `conditional-any`，则此属性会指定测试来确定是否应激活此位置。条件表达式由可赋予布尔值的一系列条件组成，例如 `"advertised-domain is sun.com"` 或 `"ip:bge0 is-not active"`。下面的“条件表达式”一节中定义了这些表达式的格式。如果指定了多个条件，则要满足激活要求，要么所有条件均为真（当 `activation-mode` 是 `conditional-all` 时），要么任何一个条件为真（当 `activation-mode` 是 `conditional-any` 时）。

注意 `advertised-domain` 与 `system-domain` 之间的区别。声明域可通过外部通信方式了解，例如 DHCP 服务器声明的 `DNSdomain` 或 `NISdomain`。此属性适用于位置的条件激活；例如，如果声明域是 `mycompany.com`，则激活“工作”位置。系统域是当前分配给系统的域，即由 [domainname\(1M\)](#) 命令返回的值。此属性适用于 ENM 的条件激活，因为只有在激活位置并且为该特定域配置系统后，条件才会变为真。

default-domain: 字符串: *system domain name*

应适用于系统的域名；此域名将由 NIS 和 LDAP 使用。

nameservices: 枚举值列表: *files|dns|nis|ldap*

指定应配置的名称服务, 例如 DNS、NIS 和 LDAP。

nameservices-config-file: 字符串: *path to nsswitch.conf file*

指定要使用的 *nsswitch.conf* 文件。此属性必须始终具有一个值。如果 **nameservices** 属性指定单个名称服务, 则此属性缺省情况下将包含 **nameservices** 属性中指定的名称服务的 */etc/nsswitch.nameservice*; 用户可更改该值。如果 **nameservices** 属性标识多个名称服务, 则用户必须为此属性另外指定一个值。

dns-nameservice-configsrc: 枚举值列表: *manual|dhcp*

指定应该用来获取 DNS 名称服务的配置信息的源。如果带有 *manual*, 则将使用其余的 **dns-nameservice-*** 属性。

dns-nameservice-domain: 字符串: *domain name*

dns-nameservice-servers: 字符串列表: *name server address(es)*

dns-nameservice-search: 字符串列表: *domain search string*

dns-nameservice-sortlist: 字符串列表: *IP address, netmask pair(s)*

dns-nameservice-options: 字符串列表: *resolver variable(s) to be modified*

如果 DNS 是已配置的名称服务之一, 并且 *manual* 是其配置源之一, 则这些属性可指定其配置参数。只有 **servers** 是必需属性; **search** 和 **domain** 是互斥的, 因此只能指定其中一项, 如果同时使用了 *dhcp*, 则会覆盖 DHCP 服务器提供的域。这些值的格式与 [resolv.conf\(4\)](#) 中的各个选项相同。

nis-nameservice-configsrc: 枚举值列表: *manual|dhcp*

指定应该用来获取 NIS 名称服务的配置信息的源。如果带有 *manual*, 则将使用其余的 **nis-nameservice-*** 属性。

nis-nameservice-servers: 字符串列表: *name server address(es)*

如果 NIS 是已配置的名称服务之一, 并且 *manual* 是其配置源之一, 则此属性可指定其服务器地址。如果未指定此属性, 则会以广播模式启动 NIS 客户机。

ldap-nameservice-configsrc: 枚举值列表: *manual*

指定应该用来获取 LDAP 名称服务的配置信息的源。*manual* 当前是 LDAP 的唯一选项; 因此必须提供 **ldap-nameservice-*** 属性才能完成 LDAP 配置。

ldap-nameservice-servers: 字符串列表: *name server address(es)*

如果 LDAP 是已配置的名称服务之一, 则此属性可指定其服务器地址。此属性为必需属性, 并且预计指定的服务器将具有用于完成客户机配置的客户机配置文件。

nfsv4-domain: 字符串: *domain name to be used for NFSv4*

指定要使用的 NFSv4 域。如果不指定此值, 则将使用名称服务域。

ipfilter-config-file: 字符串: *path to ipfilter IPv4 configuration file*

ipfilter-v6-config-file: 字符串: *path to ipfilter IPv6*

ipnat-config-file: 字符串: *path to ipnat configuration file*

ippool-config-file: 字符串: *path to ippool configuration file*

这些属性分别为 [ipfilter\(5\)](#) 配置的组件指定规则文件的路径。如果设置了给定的 **config-file** 属性, 则将启用对应的 IP 过滤器组件, 并从指定的文件读取配置。

ike-config-file: 字符串: *path to IKE configuration file*

指定 IKE 配置文件。如果指定了一个值, 则将启用 `svc:/network/ipsec/ike` 服务, 并从指定的文件读取配置。

ipsecpolicy-config-file: 字符串: *path to IPsec policy configuration file*

指定 IPsec 策略配置文件。如果指定了一个值, 则将启用 `svc:/network/ipsec/policy` 服务, 并从指定的文件读取配置。

ENM 属性

activation-mode: 枚举值: `manual | conditional-all | conditional-any`
用于自动激活此 ENM 的触发器类型。

缺省值为 `manual`。

enabled: 布尔值: `true | false`

如果 **activation-mode** 是 `manual`, 则 **enabled** 属性将反映 ENM 的当前状态。此属性是只读属性; 通过使用 [netadm\(1M\)](#) 启用或禁用 ENM, 可间接更改此属性。

缺省值为 `false`。

conditions: 字符串列表: *conditional expressions*

如果 **activation-mode** 设置为 `conditional-all` 或 `conditional-any`, 则此属性会指定测试来确定是否应激活此 ENM。条件表达式由可赋予布尔值的一系列条件组成, 例如 `"system-domain is sun.com"` 或 `"ip:bge0 is-not active"`。下面的“条件表达式”一节中定义了这些表达式的格式。如果指定了多个条件, 则要满足激活要求, 要么所有条件均为真 (当 **activation-mode** 是 `conditional-all` 时), 要么任何一个条件为真 (当 **activation-mode** 是 `conditional-any` 时)。

注意 **advertised-domain** 与 **system-domain** 之间的区别。声明域可通过外部通信方式了解, 例如 DHCP 服务器声明的 `DNSdomain` 或 `NISdomain`。此属性适用于位置的条件激活; 例如, 如果声明域是 `mycompany.com`, 则激活“工作”位置。系统域是当前分配给系统的域, 即由 `|domainname|(1M)` 命令返回的值。此属性适用于 ENM 的条件激活, 因为只有当在激活位置并且为该特定域配置系统后, 条件才会变为真。

fmri: 字符串: *service FMRI*

如果此 ENM 已作为 SMF 服务执行, 则此属性会标识该服务。如果指定了此属性, 则将通过启用服务来激活 ENM, 通过禁用服务来取消激活 ENM。

start: 字符串: *start command*

如果此 ENM 未作为 SMF 服务执行, 则此属性会标识应执行以启动或激活 ENM 的命令。如果已设置 FMRI 属性, 则将忽略此属性。

stop: 字符串: *stop command*

如果此 ENM 未作为 SMF 服务执行, 则此属性会标识应执行以停止/取消激活 ENM 的命令。如果已设置 FMRI 属性, 则将忽略此属性。

已知的 WLAN 属性

bssids: 字符串列表: WLAN BSSID (AP MAC 地址)

如果一个特定访问点的优先级高于具有相同名称/ESSID 的其他访问点, 则此属性允许指定 AP 的 BSSID。

priority:uint64: 一个数字优先级值
此 WLAN 的相对优先级；数字越小表示优先级越高。

请注意，可以更改该数字，前提是对已知 WLAN 对象的设置的后续更改需要进行此类更改。以配置有两个已知的 WLAN 对象 wlanA 和 wlanB 的系统为例。wlanA 优先级为 1，wlanB 优先级为 2。创建一个新的已知 WLAN wlanC，并且分配的优先级为 2。在此情况下，整个列表将被更新，使 wlanA 的优先级为 1，wlanC 的优先级为 2，而 wlanB 的优先级为 3。任何两个 WLAN 都不得具有相同的优先级值，因此增加优先级为 2 的 wlanC 会强制 wlanB 的优先级下降。如果存在优先级低于 wlanB 的任何其他 WLAN，则它们的优先级也会相应地改变。

keyslot:uint64: 要用于此 WLAN 的密钥插槽
如果 WLAN 使用支持多个密钥插槽的加密模式，则可在此属性中指定放置密钥的插槽。如果未指定，缺省情况下将使用插槽 1。

keyname: 字符串列表：安全对象名称
允许用户将使用 [dladm\(1M\)](#) 创建的安全对象与已知的 WLAN 进行关联。

条件表达式

可根据用户指定的一组条件激活位置和 ENM。下表概述了这些条件表达式的语法。

对象类型	对象	条件
ncp ncu enm loc	name	is/is-not active
Object Type	Condition	Object
-----	-----	-----
essid	is/is-not/contains/does-not-contain	name string
bssid	is/is-not	bssid string
ip-address	is/is-not	IPv4 or IPv6 address
ip-address	is-in-range/is-not-in-range	IPv4 or IPv6 address plus netmask/prefixlen
advertised-domain	is/is-not/contains/does-not-contain	name string
system-domain	is/is-not/contains/does-not-contain	name string

选项

以下选项与任何特定的 netcfg 子命令无关联：

- d
在从命令文件（参见以下选项）读取子命令之前删除所有配置。
- f *command_file*
读取并执行 *command_file* 中的 netcfg 子命令。

子命令

所支持的子命令如下。

cancel
结束当前配置文件指定，而不向持久性存储器提交当前所做的更改，然后进入下一更高级别的作用域。

此子命令在 NCP 和配置文件作用域中有效。

clear *prop-name*

清除指定属性的值。

此子命令在配置文件作用域中有效。

commit

将当前配置文件提交至持久性存储器。由于待提交的配置必须正确，因此该操作还会在对象上自动执行验证。在离开当前作用域（使用 **end** 或 **exit** 子命令）时，会自动尝试执行提交操作。

此子命令在配置文件作用域中有效。

请注意，在非交互模式中不需要进行提交，这是由于更改值的任何子命令均采用隐式提交。

create [-t *template*] *object-type* [*class*] *object-name*

创建指定类型和名称的内存中配置文件。-t *template* 选项指定新对象应与 *template* 完全相同，其中 *template* 是相同类型的现有对象的名称。如果不使用 -t 选项，则用缺省值创建新对象。对于 NCP，只可创建一个“用户”NCP。

此子命令在全局和 NCP 作用域中有效。

destroy { -a | *object-type* [*class*] *object-name* }

从内存和持久性存储器中删除所有配置文件或指定的配置文件。此操作是即时操作，不需要提交。已销毁的对象无法进行恢复。

如果指定了特定对象，此子命令将在全局和 NCP 作用域中有效；-a 选项仅在全局作用域中有效。

end

结束当前配置文件指定，然后进入下一更高级别的作用域。在结束前将验证并提交当前对象；如果验证或提交失败，则会发出一条相应的错误消息，并为用户提供结束机会，而不提交当前所做的更改，或者保留在当前作用域中继续进行编辑。

此子命令在任何作用域中均有效。

exit

退出 netcfg 会话。在结束前将验证并提交当前配置文件；如果验证或提交失败，则会发出一条相应的错误消息，并为用户提供退出机会，而不提交当前所做的更改，或者保留在当前作用域中继续进行编辑。

此子命令在任何作用域中均有效。

export [-d] [-f *output-file*] [*object-type* [*class*] *object-name*]

将当前作用域或指定作用域中的当前配置输出到标准输出，或输出到使用 -f 选项指定的文件。-d 选项生成 **destroy -a** 作为首行输出。此子命令以适用于命令文件的格式生成输出。无法导出系统创建的对象（包括 Automatic NCP 和 Automatic、NoNet 以及 User 位置）。

此子命令在任何作用域中均有效。

get [-v] *prop-name*

获取指定属性的当前（内存中）值。缺省情况下，将会同时输出属性名称和属性值；如果指定了 -v 选项，则只会显示属性值。

此子命令在配置文件作用域中有效。

help [*subcommand*]

显示一般帮助或有关特定子命令的帮助。

此子命令在任何作用域中均有效。

list [-a] [*object-type* [*class*] *object-name*]

列出当前作用域或指定作用域中的所有配置文件、属性值对和资源。在列出对象的属性时，缺省行为是仅列出适用于所指定配置的属性。也就是说，如果列出 `ipv4-addrsrc` 是 `dhcp` 的 IP NCU，则不会列出 `ipv4-addr` 属性。包含 -a 选项将导致列出所有属性，无论这些属性是否适用于当前设置。

此子命令在任何作用域中均有效。

revert

删除对当前配置文件所做的任何当前更改，并从持久性存储器中恢复值。

此子命令在配置文件作用域中有效。

select *object-type* [*class*] *object-name*

选择当前作用域级别中可用的配置文件之一，向下跳转到该对象的作用域中。所选对象将从持久性存储器装入到内存中。

此子命令在全局和 NCP 作用域中有效。

set *prop-name=**value1*[,*value2*...]

设置指定属性的当前（内存中）值。如果在非交互模式下执行，则会同时将所做的更改提交到持久性存储器。

用于分隔多值属性的值的分隔符是 ","（逗号）。如果此类属性中的任何单个值包含逗号，则必须对其进行转义（即表示为 \, ）。如果属性只有一个值，其中的逗号不会被解释为分隔符，因此无需进行转义。

此子命令在配置文件作用域中有效。

verify

验证当前的内存中对象的配置是否有效。

此子命令在配置文件作用域中有效。

walkprop [-a]

遍历与当前配置文件关联的每个属性。对于每个属性，均会显示名称和当前值，并给出提示以使用户更改当前值。

用于分隔多值属性的值的分隔符是","（逗号）。如果此类属性中的任何单个值包含逗号，则必须对其进行转义（即表示为\,）。如果属性只有一个值，其中的逗号不会被解释为分隔符，因此无需进行转义。

缺省情况下，只会遍历对于已设置的属性而言所需要的属性；也就是说，如果将 `ipv4-addrsrc` 设置为 `dhcp`，则不会遍历 `ipv4-addr`。包含 `-a` 选项将导致遍历所有可用的属性。

此子命令在配置文件作用域中有效。

此子命令仅在交互模式下有意义。

示例

示例1 设置NCU属性

以下命令从命令行（即在非交互模式下）设置NCU属性。

```
# netcfg "select ncp User; select ncu phys net1; set link-mtu=1492"
```

示例2 列出顶层配置文件

以下命令从命令行列出所有顶层配置文件。

```
# netcfg list
```

NCPs:

```
Automatic
User
```

Locations:

```
Automatic
home
NoNet
office
```

ENMs:

```
enmtest
myenm
```

WLANS:

```
sunwifi
coffeeshop
linksys
```

示例3 销毁位置配置文件

以下命令从命令行销毁位置配置文件。

```
# netcfg destroy loc home
```

```
Destroyed loc 'home'
```

示例4 创建NCU配置文件

以下命令序列以交互方式创建NCU配置文件。

```
# netcfg
```

```
netcfg> select ncp user
```

示例4 创建NCU配置文件 (续)

```
netcfg:ncp:User> create ncu ip net1
Created ncu 'net1'. Walking properties ...
ip-version (ipv4,ipv6) [ipv4|ipv6]> ipv4
ipv4-addrsrc (dhcp) [dhcp|static]> static
ipv4-addr> 168.1.2.3
netcfg:ncp:User:ncu:net1> list
ncu:net1
    type:                interface
    class:                ip
    parent:               "User"
    enabled               true
    ip version:           ipv4
    ipv4-addrsrc:         static
    ipv4-addr:            "168.1.2.3"
    ipv6-addrsrc          dhcp,autoconf
netcfg:ncp:User:ncu:net1> commit
Committed changes
netcfg:ncp:User:ncu:net1> end
netcfg:ncp:User> exit
```

示例5 处理 ENM

以下命令序列选择现有的 ENM、显示其内容并更改属性值。

```
# netcfg
netcfg>select enm myenm
netcfg:enm:myenm>list
ENM:myenm
    activation-mode manual
    enabled             true
    start               "/usr/local/bin/myenm start"
    stop                "/usr/local/bin/myenm stop"
netcfg:enm:myenm>set stop="/bin/alt_stop"
netcfg:enm:myenm>list
ENM:myenm
    activation-mode manual
    enabled             true
    start               "/usr/local/bin/myenm start"
    stop                "/bin/alt_stop"
netcfg:enm:myenm>exit
Committed changes
```

示例6 配置静态地址

以下命令序列使用 Home NCP 在接口 bge0 上配置静态地址 192.168.2.12/24。在激活 Home NCP 时，缺省情况下会启用此接口。

示例6 配置静态地址 (续)

```
# netcfg
netcfg> create ncp Home
```

首先配置 phys NCU :

```
netcfg:ncp:Home> create ncu phys bge0
Created ncu 'bge0'. Walking properties ...
activation-mode (manual) [manual|prioritized]>
link-mac-addr>
link-autopush>
link-mtu>
netcfg:ncp:Home:ncu:bge0> end
Committed changes
netcfg:ncp:Home>
```

然后配置 IP NCU :

```
netcfg:ncp:Home> create ncu ip bge0
Created ncu 'bge0'. Walking properties ...
ip-version (ipv4,ipv6) [ipv4|ipv6]>
ipv4-addrsrc (dhcp) [dhcp|static]> static
ipv4-addr> 192.168.2.10/24
ipv4-default-route>
ipv6-addrsrc (dhcp,autoconf) [dhcp|autoconf|static]>
ipv6-default-route>
netcfg:ncp:Home:ncu:bge0> list
ncu:bge0
      type                interface
      class               ip
      parent              "Home"
      enabled              true
      ip-version           ipv4,ipv6
      ipv4-addrsrc         static
      ipv4-addr            "192.168.2.10/24"
      ipv6-addrsrc         dhcp,autoconf
netcfg:ncp:Home:ncu:bge0> exit
Committed changes
#
```

使用 [netadm\(1M\)](#) 切换到 Home NCP。

示例7 根据条件配置位置

以下命令序列根据条件配置位置。只要 IP 地址位于 10.0.8.0/24 子网中，就会激活此位置。此外，在此位置也会配置 NIS。

```
netcfg> select loc office
netcfg:loc:office> list
```

示例7 根据条件配置位置 (续)

```
loc:office
    activation-mode          conditional-any
    conditions               "ip-address is 10.0.8.0/24"
    enabled                 false
    nameservices             dns,nis
    nameservices-config-file "/etc/nsswitch.nis"
    dns-nameservice-configsrc dhcp
    nis-nameservice-configsrc manual
    nis-nameservice-servers  "10.2.18.24"
    default-domain          "labs.east.sun.com"
netcfg:loc:office>
```

示例8 创建已知的 WLAN

以下命令序列使用 WEP 密钥建立一个名为 coffeeshop 的已知 WLAN。

```
netcfg> select wlan coffeeshop
netcfg:wlan:coffeeshop> list
known wlan:coffeeshop
    priority                2
    keyname                 "foo"
    security-mode           wep
netcfg:wlan:coffeeshop> set priority=1
netcfg:wlan:coffeeshop> end
Committed changes
```

示例9 将当前配置导出到文件

以下命令将当前配置导出到文件。

```
netcfg> export -f /tmp/nwam.config
或者，从 Unix 命令行执行同样的任务：
# netcfg export -f /tmp/nwam.config
```

示例10 从文件中导入当前配置

以下命令从文件中导入当前配置。

```
# netcfg -f /tmp/nwam.config
```

属性

有关下列属性的描述，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）

另请参见

[dladm\(1M\)](#)、[domainname\(1M\)](#)、[netadm\(1M\)](#)、[netcfgd\(1M\)](#)、[nwamd\(1M\)](#)、[resolv.conf\(4\)](#)、[att](#)

另请参见 JDS/GNOME 手册页集合中的 [nwam-manager\(1M\)](#)。

引用名

netsservices – 启用或禁用网络服务

用法概要

netsservices open
netsservices limited

描述

netsservices 命令使用 Solaris 服务管理工具 [smf\(5\)](#) 控制通过网络从远程客户机接受的服务。

使用 limited 命令行参数调用 netsservices 时，将禁用除安全 shell 守护进程 sshd(1M) 之外的所有网络服务，或将其限制为仅响应本地请求。

使用 open 命令行参数调用 netsservices 将启用一大组网络服务，在以前的 Solaris 发行版中就是如此。

要自定义 netsservices 启用的配置集，请使用 [svcadm\(1M\)](#) 启用或禁用各个服务。使用 [svccfg\(1M\)](#) 设置确定服务是否接受来自远程客户机的输入的属性。有关服务实例的名称及其属性，请参见各个服务的手册页。

请注意，netsservices 命令的接口稳定性为 Obsolete（已过时）。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Obsolete（已过时）

另请参见

[svcadm\(1M\)](#)、[svccfg\(1M\)](#)、[sshd\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

引用名

netstat – 显示网络状态

用法概要

```
netstat [-anvR] [-f address_family] [-P protocol]
netstat -g [-nv] [-f address_family]
netstat -p [-n] [-f address_family]
netstat -s [-f address_family] [-P protocol]
    [-T u | d ] [interval [count]]
netstat -m [-T u | d ] [-v] [interval [count]]
netstat -i [-I interface] [-an] [-f address_family]
    [-T u | d ] [interval [count]]
netstat -r [-anvR] [-f address_family | filter]
netstat -M [-ns] [-f address_family]
netstat -D [-I interface] [-f address_family]
netstat -d [-f address_family]
```

描述

netstat 命令可通过各种格式显示与网络有关的特定数据结构的内容（具体格式取决于选定的选项）。

netstat 命令具有以上“用法概要”一节中显示的多种形式，如下所列：

- 此命令的第一种形式（不带必需参数）显示每项协议的活动套接字列表。
- 第二、第三和第四种形式（-g、-p 和 -s 选项）显示来自各种网络数据结构的信息。
- 第五种形式（-m 选项）显示 STREAMS 内存统计信息。
- 第六种形式（-i 选项）显示接口的状态。
- 第七种形式（-r 选项）显示路由表。
- 第八种形式（-M 选项）显示多播路由表。
- 第九种形式（-D 选项）显示一个接口或所有接口的 DHCP 状态。
- 第十种形式（-d 选项）显示目标高速缓存项的表。

下面将更详细地介绍这些形式。

如果不带任何参数（第一种形式），则 netstat 会显示 PF_INET、PF_INET6 和 PF_UNIX 的连接套接字，除非 -f 选项另外修改。

选项

-a

显示所有套接字、所有路由表项或所有接口（物理接口和逻辑接口）的状态。通常不会显示服务器进程使用的侦听器套接字。在大多数情况下，只会显示接口、主机、网络和缺省路由，并且只显示物理接口的状态。

-d

显示目标高速缓存项表。请参见下面的“显示”部分。

-f *address_family*

限制对指定的 *address_family* 的所有显示。*address_family* 的值可以为下列值之一：

- inet 适用于显示 IPv4 信息的 AF_INET 地址族。
- inet6 适用于显示 IPv6 信息的 AF_INET6 地址族。
- unix 适用于 AF_UNIX 地址族。
- sdp 适用于套接字描述协议 (Socket Description Protocol, SDP) 协议和地址族。下面列出了针对 SDP 套接字显示的地址状态。netstat 显示的标志后跟其含义。

- LST Listen
- EST Established
- PL Path Lookup
- HS Hello Request Sent
- HR Hello Request Received
- HAR Hello Ack Recvd
- HAS Hello Ack sent
- DR Fin received
- DS Fin sent
- DSA Fin Ack recvd
- DRC Simultaneous Disconnect
- DSC Disconnect sent (peer already closed)
- TW1 Time Wait 1
- TW2 Time Wait 2
- CLD Closed
- ERR Error
- INV Invalid
- UNK Unknown

对于 SDP 协议和地址族，netstat 显示以下列标题：

- Local Address 本地 IP 地址
- Remote Address 远程 IP 地址
- State 套接字的当前状态
- RxBPending 未读字节
- TxBQueued 排队等待 Tx 的字节（包括 TxBPosted）
- TxBPosted 发送给 HW 进行传输的字节
- LAdvtsz 本地通告的缓冲区大小
- RAdvtsz 远程通告的缓冲区大小
- LAdvtsBuff 本地通告的 Rx 缓冲区的数量
- RAdvtsBuff 远程通告的 Rx 缓冲区的数量

LPostBuff

当前发布的 Rx 缓冲区的数量

-f *filter*

如果只带有 **-r** 选项，显示的路由将限于匹配指定过滤器的路由。过滤器规则由 *keyword:value* 对组成。已知关键字和值的语法如下：

af: {inet inet6 unix sdp <i>number</i> }	选择地址族。这与 -f <i>address_family</i> 相同，两种语法均受支持。
outif: { <i>name</i> <i>ifIndex</i> any none}	选择输出接口。可以按名称（如 hme0 ）或按 ifIndex 编号（如 2）指定接口。如果使用 any ，过滤器会匹配具有指定接口（除 null 接口以外的任何接口）的所有路由。如果使用 none ，过滤器会匹配具有 null 接口的所有路由。请注意，可以使用 ifconfig(1M) 的 -a 选项查看接口的索引号 (<i>ifIndex</i>)。
dst: { <i>ip-address</i> [/ <i>mask</i>] any none}	选择目标 IP 地址。如果指定有掩码长度，则会选择具有匹配掩码或更长（更多特定）掩码的任何路由。如果使用 any ，则会选择除 0 以外的所有地址。如果使用 none ，则会选择地址 0。
flags: [+ -]?[ABDGHLSU]+	选择标有指定标志的路由。缺省情况下，必须设置指定的标志才能进行匹配。如果带有前导 + ，则必须设置指定的标志，而忽略其他标志。如果带有前导 - ，则不得设置指定的标志，而允许设置其他标志。

您可以指定多个 **-f** 实例来指定多个过滤器。例如：

```
% netstat -nr -f outif:hme0 -f outif:hme1 -f dst:10.0.0.0/8
```

上述显示 10.0.0.0/8 网络中掩码长度为 8 或更长、输出接口为 **hme0** 或 **hme1** 的路由，并排除所有其他路由。

-g

显示所有接口的多播组成员。如果包含 **-v** 选项，则还会显示特定于源的成员信息。请参见下文的“显示”部分。

-i

显示用于 IP 通信的接口的状态。通常，此选项会显示物理接口的统计信息。与 **-a** 选项结合使用时，此选项还会报告逻辑接口的信息。请参见 **ifconfig(1M)**。

-m

显示 STREAMS 内存统计信息。

-n

以数字形式显示网络地址。**netstat** 通常以符号形式显示地址。此选项可与任何显示格式配合使用。

- p**
显示网络-介质表。请参见下文的“显示”部分。
- r**
显示路由表。通常，只会显示接口、主机、网络和缺省路由；但是，当此选项与 **-a** 选项结合使用时，将显示所有路由，包括高速缓存。如果尚未设置多播路由，则 **-ra** 可能不会显示任何多播路由项，但内核将根据需要派生此类项。
- s**
显示每个协议的统计信息。与 **-M** 选项配合使用时，改为显示多播路由统计信息。与 **-a** 选项配合使用时，将会显示每个接口的统计信息，以及显示系统的全局统计信息（如果可用）。请参见下文的“显示”部分。
- T u | d**
显示时间戳。

指定 **u** 表示时间的内部表示形式的印刷表示形式。请参见 [time\(2\)](#)。指定 **d** 表示标准日期格式。请参见 [date\(1\)](#)。
- v**
详细模式。显示套接字的其他信息、STREAMS 内存统计信息、路由表和多播组成员。
- I interface**
显示特定接口的状态。*interface* 可以是任何有效的接口，例如 **hme0** 或 **eri0**。通常，会显示物理接口的状态和统计信息。当此选项与 **-a** 选项结合使用时，还会报告逻辑接口的信息。
- M**
显示多播路由表。与 **-s** 选项配合使用时，改为显示多播路由统计信息。
- P protocol**
将统计信息的显示或所有套接字的状态限制为适用于 *protocol* 的项目。*protocol* 可以是 **ip**、**ipv6**、**icmp**、**icmpv6**、**icmp**、**icmpv6**、**igmp**、**udp**、**tcp** 和 **rawip** 之一。**rawip** 也可指定为 **raw**。此命令只接受全小写形式的协议选项。
- D**
显示 DHCP 配置的接口的状态。
- R**
此修饰符显示套接字和路由表项的扩展安全属性。**-R** 修饰符仅在系统配置有 Solaris Trusted Extensions 功能时可用。

如果只带有 **-r** 选项，此选项将显示路由项的网关安全属性。有关安全属性的更多信息，请参见 [route\(1M\)](#)。

当使用此命令的第一种形式显示套接字信息时，此选项会显示多级别端口 (Multi-Level Port, MLP) 套接字的其他信息。其中包括：
- 对等方的标签（如果已连接套接字）。

- 以下标志可附加到套接字的“状态”输出：
 - P 套接字是区域专用 IP 地址上的 MLP。
 - S 套接字是区域之间共享的 IP 地址上的 MLP。

操作数

interval 每 *interval* 秒显示一次自上次显示以来累积的统计信息，并不断重复直到指定了 *count*。用 *interval* 进行调用时，第一行 netstat 输出显示自上次重新引导以来累积的统计信息。

以下选项支持 *interval*：-i、-m、-s 和 -Ms。有些值是配置参数，仅在每次间隔时重新显示。

count 按 *count* 指定的次数以及 *interval* 指定的间隔显示接口统计信息。

显示

活动套接字（第一种形式） 每个活动套接字的显示内容均显示本地地址和远程地址、发送队列和接收队列的大小（字节）、发送窗口和接收窗口（字节）以及协议的内部状态。

通常用于显示套接字地址的符号格式为：

hostname.port

前提是指定了主机的名称，或

network.port

前提是套接字地址指定网络，而不指定特定主机。

与套接字关联的数字主机地址或网络号用于在 *hosts* 或 *networks* 数据库中查找对应的符号主机名或网络名称。

如果地址的网络或主机名未知，或者指定了 -n 选项，则会显示数字网络地址。未指定的地址和端口或“通配符”地址和端口以星号 (*) 显示。有关 Internet 命名约定的更多信息，请参见 [inet\(7P\)](#) 和 [inet6\(7P\)](#)。

对于 SCTP 套接字，由于端点可由多个地址表示，因此详细选项 (-v) 会列显所有本地地址和远程地址。

TCP Sockets TCP 套接字的可能状态值如下：

- BOUND 已绑定，准备连接或侦听。
- CLOSED 已关闭。套接字没有被使用。
- CLOSING 已关闭，然后远程关机；正在等待确认。
- CLOSE_WAIT 远程关机；正在等待套接字关闭。
- ESTABLISHED 已建立连接。

FIN_WAIT_1	套接字已关闭；正在关闭连接。
FIN_WAIT_2	套接字已关闭；正在等待远程关机。
IDLE	空闲，已打开但是未绑定。
LAST_ACK	远程关机，然后关闭；正在等待确认。
LISTEN	正在侦听传入连接。
SYN_RECEIVED	正在对连接进行初始同步。
SYN_SENT	正在积极尝试建立连接。
TIME_WAIT	关闭后等待远程关机重新传输。

SCTP Sockets

SCTP 套接字的可能状态值如下：

CLOSED	已关闭。套接字没有被使用。
LISTEN	正在侦听传入关联。
ESTABLISHED	已建立关联。
COOKIE_WAIT	INIT 已被发送到对等方，正在等待确认。
COOKIE_ECHOED	INIT-ACK 的状态 cookie 已被发送到对等方，正在等待确认。
SHUTDOWN_PENDING	已接收到来自上一层的 SHUTDOWN，正在等待确认对等方的所有未处理 DATA。
SHUTDOWN_SENT	所有未处理数据已在 SHUTDOWN_SENT 状态中确认。SHUTDOWN 已被发送到对等方，正在等待确认。
SHUTDOWN_RECEIVED	已接收到来自对等方的 SHUTDOWN，正在等待确认所有未处理 DATA。
SHUTDOWN_ACK_SENT	所有未处理数据已在 SHUTDOWN_RECEIVED 状态中确认。SHUTDOWN_ACK 已被发送到对等方。

网络数据结构（第二至第五种形式）

显示形式取决于选定的 -g、-m、-p 或 -s 选项。

- g 显示多播组成员的列表。
- m 显示内存使用率，例如 STREAMS mblk。
- p 显示网络-介质映射表。对于 IPv4，显示地址解析表。请参见 [arp\(1M\)](#)。对于 Pv6，显示邻居高速缓存。
- s 显示各种协议层的统计信息。

统计信息使用 MIB 指定的变量。为 ipForwarding 定义的值是：

forwarding(1) 充当网关。

not-forwarding(2) 不充当网关。

IPv6 和 ICMPv6 协议层保留每个接口的统计信息。如果指定 -a 选项并附带 -s 选项，则会显示每个接口的统计信息及其总和。否则，只会显示统计信息的总和。

对于此命令的第二、第三和第四种形式，必须至少指定 -g、-p 或 -s。您可以指定这些选项的任意组合。此外，也可以指定 -m（第五种形式）以及 -g、-p 和 -s 选项的任意组合形式。如果指定这些选项中的多个选项，则 netstat 会显示每个选项的信息。

接口状态（第六种形式）

接口状态显示列出所有当前接口的信息，每行一个接口。如果使用 -I 选项指定接口，则只会显示指定接口的信息。

列表包含接口名称、mtu（最大传输单元，或最大包大小）（参见 ifconfig(1M)）、接口连接到的网络、每个接口的地址以及与接口关联的计数器。计数器分别显示输入包、输入错误、输出包、输出错误和冲突的数量。对于点对点接口，"Net/Dest"（网络/目标）字段为链路另一侧的名称或地址。

如果指定 -a 选项并附带 -i 选项或 -I 选项，则输出包括物理接口的名称、每个逻辑接口的输入包和输出包的数量以及其他信息。

如果指定了 -n 选项，则列表会显示 IP 地址而不是接口名称。

如果指定了可选的 interval，则会在 interval 秒内连续显示输出，直到用户中断或达到 count。请参见“操作数”部分。

使用 -I 选项可指定物理接口。与 interval 操作数配合使用时，-I 选项的输出具有以下格式：

input					output				
input	eri0				input	(Total)			
packets	errs	packets	errs	colls	packets	errs	packets	errs	colls
227681	0	659471	1	502	261331	0	99597	1	502
10	0	0	0	0	10	0	0	0	0
8	0	0	0	0	8	0	0	0	0
10	0	2	0	0	10	0	2	0	0

如果不指定输入接口，则会显示地址族 inet 或 inet6 的第一个接口。

路由表（第七种形式）

路由表显示列出可用的路由以及每个路由的状态。每个路由都由目标主机或网络以及用于转发包的网关组成。flags 列显示路由的状态。这些标志如下：

- U 指示路由正在运行。
- G 路由指向网关。
- H 路由指向主机而不是网络。
- M 用 -multirt 选项建立的冗余路由。
- S 路由是使用 -setsrc 选项建立的。

- D 已通过重定向动态创建路由。
- B 包将被无提示放弃（设置了 RTF_BLACKHOLE）。
- R 包将被放弃，并发送 ICMP 错误（设置了 RTF_REJECT）。
- I 用 -indirect 选项建立的间接路由（非直接可达的网关）。
- Z （仅限于非全局专用 IP 区域）路由是根据在全局区域中使用 zonecfg(1M) 配置的路由信息在引导时静态添加的。

如果指定了 -a 选项，则会产生带有以下标志的路由项：

- b 多播地址。
- C 克隆链路目标上的接口主机路由项。
- L 主机的本地地址。

接口路由是为连接到本地主机的每个接口创建的；此类路由项的网关字段会显示传出接口的地址。

use 列显示使用相关路由发送或转发的包数量。

interface 项指示路由所使用的网络接口。

多播路由表（第八种形式）

多播路由表由虚拟接口表和实际路由表组成。

DHCP 接口信息（第九种形式）

DHCP 接口信息包含接口名称、接口当前状态、租用信息、包数量和标志列表。

这些状态与 RFC 2131 中设定的规范相关。

租用信息包括：

- 开始租用的时间；
- 开始续租的时间；以及
- 租用到期的时间。

当前定义的标志包括：

- BOOTP 接口含有通过 BOOTP（仅限于 IPv4）获得的租用。
- BUSY 接口正忙于 DHCP 事务。
- PRIMARY 接口为主接口。请参见 dhcpinfo(1) 和 ifconfig(1M)。
- FAILED 接口处于故障状态，并且必须手动重新启动。

系统会为已发送包的数量、已接收包的数量以及 DHCP 客户端拒绝提供的租用数量保留数据包计数。所有三个计数器均初始化为零，然后在获得租用后递增。在开始接口

的续租期限时，这些计数器将被重置。因此，这些计数器表示获得当前租用时已发送、已接收和拒绝的数据包数量，或表示尝试获得未来租用时已发送、已接收和拒绝的数据包数量。

目标高速缓存项表（第十种形式）

目标高速缓存项显示列出记录的路径 MTU、该项的存在时间（秒）和标志。p 标志指示已记录路径 MTU。s 标志指示路径 MTU 小于 IP 将允许的最小值。u 标志指示在目标高速缓存项中缓存了一些传输度量（往返时间等）。

文件

/etc/default/inet_type DEFAULT_IP 设置

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见

[arp\(1M\)](#)、[dhcpcinfo\(1\)](#)、[dhcpcagent\(1M\)](#)、[ifconfig\(1M\)](#)、[iostat\(1M\)](#)、[kstat\(1M\)](#)、[mibiisa\(1M\)](#)、

Droms, R. 编著的《RFC 2131, Dynamic Host Configuration Protocol》，Network Working Group，1997 年 3 月出版。

Droms, R. 编著的《RFC 3315, Dynamic Host Configuration Protocol for IPv6 (DHCPv6)》，Cisco Systems。2003 年 7 月出版。

附注

显示接口信息时，netstat 接受 /etc/default/inet_type 中的 DEFAULT_IP 设置。如果将其设置为 IP_VERSION4，则 netstat 将省略与 IPv6 接口、统计信息、连接、路由等有关的信息。

但是，可以在命令行上覆盖 /etc/default/inet_type 中的 DEFAULT_IP。例如，如果已使用命令行并通过 inet6 地址族或 IPv6 协议之一以显式方式请求 IPv6 信息，则会覆盖 DEFAULT_IP 设置。

如果需要检查内核崩溃之后的网络状态信息，请对 [savecore\(1M\)](#) 输出使用 [mdb\(1\)](#) 实用程序。

netstat 实用程序通过打开 /dev/tcp 并发出查询来获取系统中的 TCP 统计信息。因此，netstat 在报告连接状态时可能会显示处于 IDLE 状态的额外未用连接。

以前版本的 netstat 采用未记录的方法报告使用 [kstat\(7D\)](#) 工具发布的内核统计信息。此功能已被删除。改为使用 [kstat\(1M\)](#)。

netstat 将其输出限制为与 netstat 运行区域相关的信息。（这对共享 IP 区域和专用 IP 区域均适用。）

引用名	netstrategy – return network configuration information
用法概要	/usr/sbin/netstrategy
描述	The netstrategy command determines the network configuration strategy in use on a system and returns information in a form that is easily consumable by a script. The command returns three tokens: <i><root filesystem type> <primary interface> <network config strategy></i> These tokens are described as follows: <i><root filesystem type></i> Type of filesystem that contains the bootable kernel, as would be specified in the <i>fstype</i> column of the mnttab(4). <i><primary interface></i> Name of the primary network interface. For a diskless machine, this is the interface used to load the kernel. <i><network config strategy></i> The means by which a system obtains its IP address for booting. This can be one of rarp, dhcp, or none. The netstrategy command is not intended for use on a command line.
选项	The netstrategy command has no options.
退出状态	0 Success. !=0 An error occurred.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [ifconfig\(1M\)](#), [mnttab\(4\)](#), [attributes\(5\)](#)

引用名	newaliases – rebuild the data base for the mail aliases file																
用法概要	newaliases																
描述	newaliases rebuilds the random access data base for the mail aliases file <code>/etc/mail/aliases</code>. newaliases accepts all the flags that sendmail(1M) accepts. However, most of these flags have no effect, except for the <code>-C</code> option and three of the Processing Options that can be set from a configuration file with the <code>-o</code> option: <table><tr><td><code>-C /path/to/alt/config/file</code></td><td>Use alternate configuration file.</td></tr><tr><td><code>-oAfile</code></td><td>Specify possible alias files.</td></tr><tr><td><code>-oLn</code></td><td>Set the default log level to <i>n</i>. Defaults to 9.</td></tr><tr><td><code>-on</code></td><td>Validate the RHS of aliases when rebuilding the aliases(4) database.</td></tr></table> newaliases runs in verbose mode (<code>-v</code> option) automatically.	<code>-C /path/to/alt/config/file</code>	Use alternate configuration file.	<code>-oAfile</code>	Specify possible alias files.	<code>-oLn</code>	Set the default log level to <i>n</i> . Defaults to 9.	<code>-on</code>	Validate the RHS of aliases when rebuilding the aliases(4) database.								
<code>-C /path/to/alt/config/file</code>	Use alternate configuration file.																
<code>-oAfile</code>	Specify possible alias files.																
<code>-oLn</code>	Set the default log level to <i>n</i> . Defaults to 9.																
<code>-on</code>	Validate the RHS of aliases when rebuilding the aliases(4) database.																
示例	示例 1 Running the newaliases Command The following command runs newaliases on an alias file different from the <code>/etc/mail/aliases</code> default in sendmail(1M): <pre>example% newaliases -oA/path/to/alternate/alias/file</pre>																
退出状态	newaliases returns an exit status describing what it did. The codes are defined in <code>/usr/include/sysexits.h</code>. <table><tr><td>EX_OK</td><td>Successful completion on all addresses.</td></tr><tr><td>EX_NOUSER</td><td>User name not recognized.</td></tr><tr><td>EX_UNAVAILABLE</td><td>Catchall. Necessary resources were not available.</td></tr><tr><td>EX_SYNTAX</td><td>Syntax error in address.</td></tr><tr><td>EX_SOFTWARE</td><td>Internal software error, including bad arguments.</td></tr><tr><td>EX_OSERR</td><td>Temporary operating system error, such as “cannot fork”.</td></tr><tr><td>EX_NOHOST</td><td>Host name not recognized.</td></tr><tr><td>EX_TEMPFAIL</td><td>Message could not be sent immediately, but was queued.</td></tr></table>	EX_OK	Successful completion on all addresses.	EX_NOUSER	User name not recognized.	EX_UNAVAILABLE	Catchall. Necessary resources were not available.	EX_SYNTAX	Syntax error in address.	EX_SOFTWARE	Internal software error, including bad arguments.	EX_OSERR	Temporary operating system error, such as “cannot fork”.	EX_NOHOST	Host name not recognized.	EX_TEMPFAIL	Message could not be sent immediately, but was queued.
EX_OK	Successful completion on all addresses.																
EX_NOUSER	User name not recognized.																
EX_UNAVAILABLE	Catchall. Necessary resources were not available.																
EX_SYNTAX	Syntax error in address.																
EX_SOFTWARE	Internal software error, including bad arguments.																
EX_OSERR	Temporary operating system error, such as “cannot fork”.																
EX_NOHOST	Host name not recognized.																
EX_TEMPFAIL	Message could not be sent immediately, but was queued.																
文件	<table><tr><td><code>/etc/aliases</code></td><td>Symbolic link to <code>/etc/mail/aliases</code></td></tr><tr><td><code>/etc/mail/aliases.pag</code></td><td></td></tr><tr><td><code>/etc/mail/aliases.dir</code></td><td>ndbm files maintained by newaliases</td></tr><tr><td><code>/etc/mail/aliases.db</code></td><td>Berkeley DataBase file maintained by newaliases</td></tr></table>	<code>/etc/aliases</code>	Symbolic link to <code>/etc/mail/aliases</code>	<code>/etc/mail/aliases.pag</code>		<code>/etc/mail/aliases.dir</code>	ndbm files maintained by newaliases	<code>/etc/mail/aliases.db</code>	Berkeley DataBase file maintained by newaliases								
<code>/etc/aliases</code>	Symbolic link to <code>/etc/mail/aliases</code>																
<code>/etc/mail/aliases.pag</code>																	
<code>/etc/mail/aliases.dir</code>	ndbm files maintained by newaliases																
<code>/etc/mail/aliases.db</code>	Berkeley DataBase file maintained by newaliases																

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/smtp/sendmail

另请参见

[sendmail\(1M\)](#), [aliases\(4\)](#), [attributes\(5\)](#)

引用名

newfs – construct a UFS file system

用法概要

newfs [-NSBTv] [*mkfs-options*] *raw-device*

描述

newfs is a friendly front-end to the [mkfs\(1M\)](#) program for making UFS file systems on disk partitions. newfs calculates the appropriate parameters to use and calls mkfs.

If run interactively (that is, standard input is a tty), newfs prompts for confirmation before making the file system.

If the -N option is not specified and the inodes of the device are not randomized, newfs calls [fsirand\(1M\)](#).

You must be super-user or have appropriate write privileges to use this command.

**Creating a
Multiterabyte UFS File
System**

Keep the following limitations in mind when creating a multiterabyte UFS file system:

- *nbpi* is set to 1 Mbyte unless you specifically set it higher. You cannot set *nbpi* lower than 1 Mbyte on a multiterabyte UFS file system.
- *fragsize* is set equal to *bsize*.

选项

The following options are supported:

-N

Print out the file system parameters that would be used to create the file system without actually creating the file system. [fsirand\(1M\)](#) is not called here.

-S

Sends to stdout a human-readable version of the superblock that would be used to create a filesystem with the specified configuration parameters.

-B

Sends to stdout a binary (machine-readable) version of the superblock that would be used to create a filesystem with the specified configuration parameters.

-T

Set the parameters of the file system to allow eventual growth to over a terabyte in total file system size. This option sets *fragsize* to be the same as *bsize*, and sets *nbpi* to 1 Mbyte, unless the -i option is used to make it even larger. If you use the -f or -i options to specify a *fragsize* or *nbpi* that is incompatible with this option, the user-supplied value of *fragsize* or *nbpi* is ignored.

-v

Verbose. newfs prints out its actions, including the parameters passed to mkfs.

mkfs-options

Options that override the default parameters are:

-a *apc*

The number of alternate sectors per cylinder to reserve for bad block replacement for SCSI devices only. The default is 0.

This option is not applicable for disks with EFI labels and is ignored.

-b *bsize*

The logical block size of the file system in bytes, either 4096 or 8192. The default is 8192. The sun4u architecture does not support the 4096 block size.

-c *cgsize*

The number of cylinders per cylinder group, ranging from 16 to 256. The default is calculated by dividing the number of sectors in the file system by the number of sectors in a gigabyte. Then, the result is multiplied by 32. The default value is always between 16 and 256.

mkfs can override this value. See [mkfs_ufs\(1M\)](#) for details.

This option is not applicable for disks with EFI labels and is ignored.

-C *maxcontig*

The maximum number of logical blocks, belonging to one file, that are allocated contiguously. The default is calculated as follows:

$\text{maxcontig} = \text{disk drive maximum transfer size} / \text{disk block size}$

If the disk drive's maximum transfer size cannot be determined, the default value for *maxcontig* is calculated from kernel parameters as follows:

If *maxphys* is less than *ufs_maxmaxphys*, which is typically 1 Mbyte, then *maxcontig* is set to *maxphys*. Otherwise, *maxcontig* is set to *ufs_maxmaxphys*.

You can set *maxcontig* to any positive integer value.

The actual value will be the lesser of what has been specified and what the hardware supports.

You can subsequently change this parameter by using [tunefs\(1M\)](#).

-d *gap*

Rotational delay. This option is obsolete in the Solaris 10 release. The value is always set to 0, regardless of the input value.

-f *fragsize*

The smallest amount of disk space in bytes that can be allocated to a file. *fragsize* must be a power of 2 divisor of *bsize*, where:

$\text{bsize} / \text{fragsize}$ is 1, 2, 4, or 8.

This means that if the logical block size is 4096, legal values for *fragsize* are 512, 1024, 2048, and 4096. When the logical block size is 8192, legal values are 1024, 2048, 4096, and 8192. The default value is 1024.

For file systems greater than 1 terabyte or for file systems created with the -T option, *fragsize* is forced to match block size (*bsize*).

-i *nbpi*

The number of bytes per inode, which specifies the density of inodes in the file system. The number is divided into the total size of the file system to determine the number of inodes to create.

This value should reflect the expected average size of files in the file system. If fewer inodes are desired, a larger number should be used. To create more inodes, a smaller number should be given. The default for *nbpi* is as follows:

Disk size	Density
Less than 1GB	2048
Less than 2GB	4096
Less than 3GB	6144
3GB to 1 Tbyte	8192
Greater than 1 Tbyte or created with -T	1048576

The number of inodes can increase if the file system is expanded with the `growfs` command.

-m *free*

The minimum percentage of free space to maintain in the file system, between 0% and 99%, inclusively. This space is off-limits to users. Once the file system is filled to this threshold, only the super-user can continue writing to the file system.

The default is $((64 \text{ Mbytes/partition size}) * 100)$, rounded down to the nearest integer and limited between 1% and 10%, inclusively.

This parameter can be subsequently changed using the [tuneufs\(1M\)](#) command.

-n *nrpos*

The number of different rotational positions in which to divide a cylinder group. The default is 8.

This option is not applicable for disks with EFI labels and is ignored.

-o *space* | *time*

The file system can either be instructed to try to minimize the *time* spent allocating blocks, or to try to minimize the *space* fragmentation on the disk. The default is *time*.

This parameter can subsequently be changed with the [tuneufs\(1M\)](#) command.

-r *rpm*

The rotational speed of the disk in revolutions per minute. The default is driver- or device-specific.

Note that you specify *rpm* for `newfs` and *rps* for `mkfs`.

This option is not applicable for disks with EFI labels and is ignored.

-s size

The size of the file system in sectors. The default is to use the entire partition.

-t ntrack

The number of tracks per cylinder on the disk. The default is taken from the disk label.

This option is not applicable for disks with EFI labels and is ignored.

操作数

The following operands are supported:

raw-device

The name of a raw special device residing in the /dev directory (for example, /dev/rdisk/c0t0d0s6) on which to create the file system.

用法

See [largefile\(5\)](#) for the description of the behavior of newfs when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).

示例

示例 1 Displaying the Parameters for the Raw Special Device

The following example verbosely displays the parameters for the raw special device, c0t0d0s6. It does not actually create a new file system:

```
example# newfs -Nv /dev/rdisk/c0t0d0s6
mkfs -F ufs -o N /dev/rdisk/c0t0d0s6 1112940 54 15 8192 1024 16 10 60
2048 t 0 -1 8 /dev/rdisk/c0t0d0s6: 1112940 sectors in
1374 cylinders of 15 tracks, 54 sectors 569.8MB in 86 cyl
groups (16 c/g, 6.64MB/g, 3072 i/g) super-block backups
(for fsck -b #) at:
32, 13056, 26080, 39104, 52128, 65152, 78176, 91200, 104224, . . .
```

示例 2 Creating a UFS File System That Will Eventually Be Grown to a Multiterabyte UFS File System

The following example creates a UFS file system that will eventually be grown to a multiterabyte UFS file system.

This command creates a 800-Gbyte file system on the volume, /dev/md/rdisk/d99.

```
# newfs -T /dev/md/rdisk/d99
newfs: construct a new file system /dev/md/rdisk/d99: (y/n)? y
/dev/md/rdisk/d99: 1677754368 sectors in 45512 cylinders of
144 tracks, 256 sectors
819216.0MB in 1821 cyl groups (25 c/g, 450.00MB/g, 448 i/g) . . .
```

Then, if you increase the volume size for this file system, you can use the growfs command to expand the file system. The file system is grown to 1.2 terabytes in this example:

```
# growfs -v /dev/md/rdisk/d99
/usr/lib/fs/ufs/mkfs -G /dev/md/rdisk/d99 2516631552 /dev/md/rdisk/d99:
2516631552 sectors in 68268 cylinders of 144 tracks, 256 sectors
1228824.0MB in 2731 cyl groups (25 c/g, 450.00MB/g, 448 i/g) . . .
```

退出状态

The following exit values are returned:

- 0
The operation was successful.
 - 1, 10
Usage error or internal error. A message is output to STDERR explaining the error.
- Other exit values may be returned by [mkfs\(1M\)](#), which is called by newfs.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[fsck\(1M\)](#), [fsck_ufs\(1M\)](#), [fsirand\(1M\)](#), [mkfs\(1M\)](#), [mkfs_ufs\(1M\)](#), [tunefs\(1M\)](#), [attributes\(5\)](#), [largefile\(5\)](#), [ufs\(7FS\)](#)

诊断

- newfs: No such file or directory
The device specified does not exist, or a disk partition was not specified.
- special: cannot open
You must write access to the device to use this command.

引用名	newkey – create a new Diffie-Hellman key pair in the publickey database										
用法概要	<pre>newkey -h <i>hostname</i> [-s nis files ldap] newkey -u <i>username</i> [-s nis files ldap]</pre>										
描述	newkey establishes new public keys for users and machines on the network. These keys are needed when using secure RPC or secure NFS service. newkey prompts for a password for the given <i>username</i> or <i>hostname</i> and then creates a new public/secret Diffie-Hellman 192 bit key pair for the user or host. The secret key is encrypted with the given password. The key pair can be stored in the <code>/etc/publickey</code> file or the NIS <code>publickey</code> map. newkey consults the <code>publickey</code> entry in the name service switch configuration file (see nsswitch.conf(4)) to determine which naming service is used to store the secure RPC keys. If the <code>publickey</code> entry specifies a unique name service, newkey will add the key in the specified name service. However, if there are multiple name services listed, newkey cannot decide which source to update and will display an error message. The user is required to specify the source explicitly with the <code>-s</code> option. In the case of NIS, newkey should be run by the superuser on the master NIS server for that domain. In the case of LDAP, newkey should be run by the superuser on a machine that also recognizes the directory manager's bind distinguished name (DN) and password to perform an LDAP update for the host.										
选项	<table><tr><td>-h <i>hostname</i></td><td>Create a new public/secret key pair for the privileged user at the given <i>hostname</i>. Prompts for a password for the given <i>hostname</i>.</td></tr><tr><td>-u <i>username</i></td><td>Create a new public/secret key pair for the given <i>username</i>. Prompts for a password for the given <i>username</i>.</td></tr><tr><td>-s nis</td><td></td></tr><tr><td>-s files</td><td></td></tr><tr><td>-s ldap</td><td>Update the database in the specified source: <i>nis</i> (for NIS), <i>files</i>, or <i>ldap</i> (LDAP). Other sources may be available in the future.</td></tr></table>	-h <i>hostname</i>	Create a new public/secret key pair for the privileged user at the given <i>hostname</i> . Prompts for a password for the given <i>hostname</i> .	-u <i>username</i>	Create a new public/secret key pair for the given <i>username</i> . Prompts for a password for the given <i>username</i> .	-s nis		-s files		-s ldap	Update the database in the specified source: <i>nis</i> (for NIS), <i>files</i> , or <i>ldap</i> (LDAP). Other sources may be available in the future.
-h <i>hostname</i>	Create a new public/secret key pair for the privileged user at the given <i>hostname</i> . Prompts for a password for the given <i>hostname</i> .										
-u <i>username</i>	Create a new public/secret key pair for the given <i>username</i> . Prompts for a password for the given <i>username</i> .										
-s nis											
-s files											
-s ldap	Update the database in the specified source: <i>nis</i> (for NIS), <i>files</i> , or <i>ldap</i> (LDAP). Other sources may be available in the future.										
属性	See attributes(5) for descriptions of the following attributes:										

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [chkey\(1\)](#), [keylogin\(1\)](#), [nsswitch.conf\(4\)](#), [publickey\(4\)](#), [attributes\(5\)](#)

引用名 nfs4cbd – NFS 版本 4 回调守护进程

用法概要 /usr/lib/nfs/nfs4cbd

描述 nfs4cbd 守护进程管理 NFS 版本 4 协议回调程序的通信端点。nfs4cbd 在 NFS 版本 4 客户机上运行，并为可通过其发送回调的每个传输创建侦听器端口。

nfs4cbd 守护进程仅适用于 NFS 版本 4 客户机。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/ntfsprogs

另请参见 [svcs\(1\)](#)、[mount_nfs\(1M\)](#)、[svcadm\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

附注 nfs4cbd 服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

svc:/network/nfs/cbd

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

如果此服务处于禁用状态，它将在第一次进行 NFSv4 挂载时由 [mount_nfs\(1M\)](#) 和 [automountd\(1M\)](#) 启用，除非其 application/auto_enable 属性设置为 false。

未来的 Solaris 发行版中可能不存在此守护进程。

引用名	nfsd – NFS 守护进程
用法概要	<pre>/usr/lib/nfs/nfsd [-a] [-c #_conn] [-l listen_backlog] [-p protocol] [-t device] [nservers]</pre>
描述	nfsd 是处理客户机文件系统请求的守护进程。只有具有 {PRIV_SYS_NFS} 且具有足够的特权向 /var/run 进行写入的用户可以运行此守护进程。 当使用带 -a 选项的 share(1M) 时，会自动调用 nfsd 守护进程。 缺省情况下，对于版本 2 和 3，nfsd 会在 TCP 和 UDP 传输上启动。缺省情况下，对于版本 4，它会在 TCP 传输上启动。您可以使用 -p 选项更改此行为。 在调用 nfsd 命令之前必须停止之前启动的（不管是否使用了选项）nfsd 守护进程。 要更改 nfsd 的启动参数，请使用 sharectl(1M) 命令。
SMF 管理	nfsd 服务由服务管理工具 smf(5) 进行管理，其服务标识符为： <pre>svc:/network/nfs/server</pre> 可以使用 svcadm(1M) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 svcs(1) 命令来查询服务的状态。 如果通过 SIGTERM 终止了 nfsd，它将无法由服务管理工具来重新启动。但是，可以通过其他信号（例如 SIGINT）来重新启动 nfsd。 可以使用 sharectl(1M) 命令来管理与 nfsd 相关的所有参数。下面是当前支持的参数： <pre>server_versmax=num server_versmin=num</pre> NFS 服务器仅使用由这些变量指定的范围中的 NFS 版本。有效的值或版本为：2、3 和 4。如果这些参数中的一个或者两个都没有设置，则缺省的最小版本为 2，而缺省的最大版本为 4。 <pre>server_delegation=on off</pre> 缺省情况下，此变量为 on 并且 NFS 服务器向客户机提供委托。用户可以通过将此变量设置为 off（区分大小写）为所有导出的文件系统关闭委托。此变量仅适用于 NFS 版本 4。 <pre>max_connections=num</pre> 设置并发的、面向连接的连接的最大数目。缺省设置是无限制，并且可通过将其设置为 -1 来实现。等效于 nfsd 中的 -c 选项。 <pre>listen_backlog=num</pre> 为面向连接的传输上的 NFS 设置连接队列长度。缺省值为 32，表示队列中有 32 个条目。等效于 nfsd 中的 -l 选项。

protocol=ALL

仅在指定的协议上启动 `nfsd`。等效于 `nfsd` 中的 `-p` 选项。ALL 等效于 `nfsd` 命令行上的 `-a`。与 NFS SMF 参数 `device` 是互斥的。NFS SMF 参数 `device` 和 `protocol` 不得同时指定。如果同时指定了两者，`nfs/server` 服务将进入维护模式。对于 UDP 协议，只会建立版本 2 和版本 3 服务。NFS 版本 4 不支持 UDP 协议。等效于 `-p` 选项。

device=devname

仅为由指定的设备指定的传输启动 NFS 守护进程。等效于 `nfsd` 中的 `-t` 选项。与 NFS SMF 参数 `protocol` 是互斥的。NFS SMF 参数 `device` 和 `protocol` 不得同时指定。

servers=num

并发 NFS 请求的最大数目。等效于 `nfsd` 命令行上最后的数字参数。缺省值为 1024。等效于 `nservers` 操作数。

请参见下文的“示例”部分。

选项

所支持的选项如下：

- a** 在所有可用的连接和面向连接的传输（包括 UDP 和 TCP）上启动 NFS 守护进程。等效于在 SMF 中使用 `sharectl(1M)` 命令为 NFS 将 `protocol` 参数设置为 ALL。
- c #_conn** 这设置面向连接的传输上的 NFS 服务器允许的最大连接数。缺省情况下，连接数目是无限制的。等效于在 SMF 中使用 `sharectl(1M)` 命令为 NFS 设置 `max_connections` 参数。
- l** 为面向连接的传输上的 NFS TCP 设置连接队列长度。缺省值为 32 个条目。等效于在 SMF 中使用 `sharectl(1M)` 命令为 NFS 设置 `listen_backlog` 参数。
- p protocol** 在指定的协议上启动 NFS 守护进程。等效于在 SMF 中使用 `sharectl(1M)` 命令为 NFS 设置 `protocol` 参数。
- t device** 为由给定设备指定的传输启动 NFS 守护进程。等效于在 SMF 中使用 `sharectl(1M)` 命令为 NFS 设置 `device` 参数。

操作数

支持下列操作数：

- `nservers`** 这设置服务器可以处理的并发 NFS 请求的最大数目。此并发性是由根据需要在内核中创建的最多 `nservers` 个线程来实现的。`nservers` 应当基于该服务器上的预期负荷而设置。`nservers` 的数字通常为 16。如果未指定 `nservers`，则并发 NFS 请求的最大数目将缺省设置为 1。等效于在 SMF 中使用 `sharectl(1M)` 命令为 NFS 设置 `servers` 参数。

用法	如果在 <code>/etc/system</code> 中设置了 <code>NFS_PORTMON</code> 变量，则客户机必须使用特权端口（小于 <code>IPPORT_RESERVED</code> 的端口）来获取 NFS 服务。缺省情况下，该变量等于零。该变量已从 "nfs" 移到了 "nfssrv" 模块。要设置该变量，请编辑 <code>/etc/system</code> 文件并添加以下条目： <pre>set nfssrv:nfs_portmon = 1</pre>
示例	示例1 关闭委托 这些示例中指定的 <code>nfsd</code> 属性在上文中的“SMF 管理”部分进行了介绍。 委托是一项 NFSv4 功能，服务器通过该功能将文件的管理委托给客户机。例如，服务器可以向客户机授予（或不授予）读取委托或写入委托。下面的命令执行该操作，将 <code>server_delegation</code> 属性设置为 <code>off</code>。 <pre># sharectl set -p server_delegation=off nfs</pre> 示例2 确定委托的值 下面的命令获取 <code>server_delegation</code> 属性的当前值。 <pre># sharectl get -p server_delegation nfs server_delegation=on</pre> 示例3 设置并发请求的最大数目 下面的命令设置并发 NFS 请求的最大数目。 <pre># sharectl set -p servers=32 nfs</pre> 示例4 设置连接队列长度 下面的命令为面向连接的传输上的 NFS 设置最大队列长度。 <pre># sharectl set -p listen_backlog=48 nfs</pre>
退出状态	0 守护进程启动成功。 1 守护进程启动失败。
文件	<code>.nfsXXX</code> 打开但未链接的文件的客户机指针。 <code>/etc/system</code> 系统配置信息文件。 <code>/var/nfs/v4_state</code> <code>/var/nfs/v4_oldstate</code> 服务器用来管理客户机状态信息的目录。不应删除这些目录。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	service/file-system/nfs

另请参见 [ps\(1\)](#)、[svcs\(1\)](#)、[mountd\(1M\)](#)、[share\(1M\)](#)、[sharectl\(1M\)](#)、[svcadm\(1M\)](#)、[sharetab\(4\)](#)、[system\(4\)](#)、[《Oracle Solaris 管理：网络服务》](#)

附注 建议不要手动启动和重新启动 `nfsd`。如果必须要这样做，请使用 `svcadm` 来启用或禁用 `nfs` 服务 (`svc:/network/nfs/server`)。如果该服务被禁用，它将由 [share_nfs\(1M\)](#) 来启用，除非它的 `application/auto_enable` 属性设置为 `false`。有关更多信息，请参见 [《Oracle Solaris 管理：网络服务》](#) 和 [svcadm\(1M\)](#)。

引用名	nfslogd – nfs 日志记录守护进程
用法概要	/usr/lib/nfs/nfslogd
描述	nfslogd 守护进程向 Solaris NFS 服务器提供操作日志记录。nfslogd 守护进程的工作是通过分析 NFS 服务器处理的 RPC 操作来生成活动日志。仅会为在导出时启用了日志记录的文件系统生成日志。这是在文件系统导出时通过 share_nfs(1M) 命令指定的。 使用 NFS 版本 4 的 Solaris 计算机不支持 NFS 服务器日志记录。 日志文件中的每条记录包括一个时间戳、客户机系统的 IP 地址（或主机名，如果可解析）、对其执行操作的文件或目录名称，以及操作的类型。在基本格式中，操作可以是一个输入 (i) 或输出 (o) 操作。NFS 服务器日志的基本格式与 Washington University FTPd 守护进程生成的日志格式兼容。可以对日志格式进行扩展以包括目录修改操作，例如 mkdir、rmdir 和 remove。扩展的格式与 Washington University FTPd 守护进程格式不兼容。有关详细信息，请参见 nfslog.conf(4)。 NFS 服务器日志记录机制分为两个阶段。第一个阶段是由 NFS 内核模块执行的，它将原始 RPC 请求及其结果记录在由永久存储支持的工作缓冲区中。工作缓冲区的位置是在 /etc/nfs/nfslog.conf 文件中指定的。有关更多信息，请参见 nfslog.conf(4)。第二个阶段涉及到 nfslogd 用户级守护进程，它定期读取工作缓冲区，解释原始 RPC 信息，将相关的 RPC 操作分组到单个事务记录中，并生成输出日志。然后，nfslogd 守护进程将进入休眠状态，等待更多的信息记录到工作缓冲区中。可以通过在 /etc/default/nfslogd 中修改 IDLE_TIME 参数来配置该守护进程的休眠时间量。工作缓冲区供 nfslogd 守护进程内部使用。 NFS 操作使用文件句柄而不是路径名作为参数。因此，nfslogd 守护进程需要维护包含文件句柄到路径映射的数据库，以便记录与操作关联的路径名而不是记录对应的文件句柄。当客户机执行查找操作或执行会向客户机返回文件句柄的其他 NFS 操作时，将向该数据库中添加一个文件句柄条目。 在 NFS 客户机从服务器获取文件句柄后，它可以无限期地保留该文件句柄，以便以后将其用作对文件或目录执行的 NFS 操作的参数。NFS 客户机甚至在服务器重新引导多次后还可以使用该句柄。因为该数据库在服务器重新引导后仍然需要保留，所以该数据库由永久存储来支持。该数据库的位置是由 /etc/nfs/nfslog.conf 文件中的 <i>fh</i>table 参数指定的。该数据库供 nfslogd 守护进程内部使用。 为了将文件句柄映射数据库的大小保持为可管理的，nfslogd 会定期删改该数据库。它会删除超过指定的时间量未访问的文件句柄条目。/etc/default/nfslogd 中的 PRUNE_TIMEOUT 可配置参数指定删改进程的两次后续运行之间的时间间隔长度。如果某个文件句柄记录自上次执行删改进程后未被使用，则会删除该文件句柄记录。通过将 PRUNE_TIMEOUT 设置为 INT_MAX 可以有效地禁用该数据库的删改。 启用了删改功能时，始终会存在以下风险：客户机可能会将某个文件句柄保留超过 PRUNE_TIMEOUT 的时间并且在匹配的数据库中的匹配记录被删除后对该文件句柄执行 NFS 操作。在这种情况下，将无法解析文件句柄的路径名，并且日志中将包括文件句柄而不是路径名。

有各种可配置的参数可以影响 `nfslogd` 守护进程的行为。可以在 `/etc/default/nfslogd` 中找到这些参数，其描述如下：

UMASK	为日志文件、工作缓冲区文件和文件句柄映射数据库设置文件模式。
MIN_PROCESSING_SIZE	指定在处理工作信息并向日志文件进行写入之前，缓冲区文件必须达到的最小大小（以字节为单位）。MIN_PROCESSING_SIZE 的值必须介于 1 和 ulimit 之间。
IDLE_TIME	指定在等待更多信息放置到缓冲区文件中时，守护进程应休眠的时间量（以秒为单位）。IDLE_TIME 还决定了重新读取配置文件的频率。IDLE_TIME 的值必须介于 1 和 INT_MAX 之间。
MAX_LOGS_PRESERVE	nfslogd 会定期回收其日志。MAX_LOGS_PRESERVE 指定要保存的最大日志文件数目。当达到 MAX_LOGS_PRESERVE 时，在创建新日志文件时将覆盖最旧的文件。这些文件将以编号式扩展名保存，从 filename.0 开始。最旧的文件将具有最高的编号式扩展名，最大值等于为 MAX_LOGS_PRESERVE 配置的值。MAX_LOGS_PRESERVE 的值必须介于 1 和 INT_MAX 之间。
CYCLE_FREQUENCY	指定回收日志文件的频率（以小时为单位）。CYCLE_FREQUENCY 用来确保日志文件不会变得过大。CYCLE_FREQUENCY 的值必须介于 1 和 INT_MAX 之间。
MAPPING_UPDATE_INTERVAL	指定对文件句柄到路径映射表中的记录进行更新的时间间隔（以秒为单位）。不是在每次访问记录时都更新该记录的 atime，只有基于此参数判定它已老化时才会对其进行更新。删改例程使用记录访问时间来确定是否应从数据库中删除记录。此参数的值必须介于 1 和 INT_MAX 之间。
PRUNE_TIMEOUT	指定数据库记录何时超时（以小时为单位）。如果自上次访问记录以来经过的时间超过了 PRUNE_TIMEOUT，则可以从数据库中删改该记录。PRUNE_TIMEOUT 的缺省值是 168 小时（7 天）。PRUNE_TIMEOUT 的值必须介于 1 和 INT_MAX 之间。

退出状态

- 将返回下列退出值：
- 0 守护进程启动成功。
 - 1 守护进程启动失败。

文件

`/etc/nfs/nfslogtab`
`/etc/nfs/nfslog.conf`
`/etc/default/nfslogd`

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/file-system/nfs

另请参见

[share_nfs\(1M\)](#)、[nfslog.conf\(4\)](#)、[attributes\(5\)](#)

引用名	nfsmapid – NFS 用户和组 ID 映射守护进程
用法概要	<code>/usr/lib/nfs/nfsmapid</code>
描述	nfsmapid 守护进程可来回映射 NFS 版本 4 客户机和服务器使用的 NFS 版本 4 owner 和 owner_group 标识属性以及本地 UID 和 GID 编号。 nfsmapid 使用 <code>/etc/nsswitch.conf</code> 文件中的 <code>passwd</code> 和 <code>group</code> 条目指示其如何执行映射。 nfsmapid 守护进程没有客户可访问的外部接口。但您可以使用以下方法之一对 nfsmapid 进行管理配置： ■ 使用 sharectl(1M) 命令在 NFS 的 SMF 中指定 <code>nfsmapid_domain</code> 参数。 ■ 指定 <code>_nfsv4idmapdomain</code> DNS 资源记录。 文件 <code>/var/run/nfs4_domain</code> 中包含当前选定的 NFSv4 域。 有关详细信息，请参见《Oracle Solaris 管理：网络服务》《Oracle Solaris 管理：网络服务》。
SMF 管理	nfsmapid 服务由服务管理工具 smf(5) 管理，其服务标识符为： <pre>svc:/network/nfs/mapid</pre> 可以使用 svcadm(1M) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 svcs(1) 命令来查询服务的状态。 如果该服务处于禁用状态，可通过 mount_nfs(1M)、share_nfs(1M) 和 automountd(1M) 启用它，除非其 <code>application/auto_enable</code> 属性设置为 <code>false</code>。 nfsmapid 会缓存用户的 UID 和 GID。如果用户随后使用下面所列的实用程序之一更改了 UID 或 GID，nfsmapid 高速缓存将过时。此时，获取或设置属性的任何 NFS 操作都会导致交换此过时信息。要解决此问题，请重新启动 nfsmapid，如下所示： <pre># svcadm restart svc:/network/nfs/mapid:default</pre> 可使用 sharectl(1M) 命令处理指定域名 (<code>nfsmapid_domain</code>) 的启动 SMF 参数。 nfsmapid_domain NFS SMF 参数 <code>nfsmapid_domain</code> 的设置将覆盖 nfsmapid 分别用于构建和比较传出和传入属性字符串的域。此外，此设置还会覆盖用于设置 NFSv4 域的任何其他机制。如果缺少 <code>nfsmapid_domain</code> 设置，nfsmapid 守护进程将按以下方式确定 NFSv4 域： ■ 如果存在正确配置的 <code>/etc/resolv.conf</code>（请参见 resolv.conf(4)），nfsmapid 将为该域查询指定的名称服务器。 ■ 如果存在正确配置的 <code>/etc/resolv.conf</code>（请参见 resolv.conf(4)），但查询到的名称服务器没有正确的域名记录，nfsmapid 将尝试通过 BIND 接口获取域名（请参见 resolver(3RESOLV)）。

- 如果不存在 `/etc/resolv.conf`，`nfsmapid` 将重新使用配置的域名（请参见 [domainname\(1M\)](#)），返回该域名时将删除首位上的域后缀。例如，对于 `widgets.sales.acme.com`，将返回 `sales.acme.com`。
- 如果不存在 `/etc/resolv.conf`，且未配置域名（或者不存在 `/etc/defaultdomain`），那么在主机名包含完全限定域名 (fully qualified domain name, FQDN) 的情况下，`nfsmapid` 将重新从主机名中获取域名。

如果按照前面所有的步骤仍然无法获取域名，`nfsmapid` 将不配置任何域。这会导致以下行为：

- `owner` 和 `owner_group` 传出属性字符串将编码为文字 ID。例如，UID 12345 将编码为 12345。
- `nfsmapid` 会忽略传入属性字符串的 `domain` 部分，并且只针对用户或组执行名称服务查找。如果用户/组在本地系统名称服务数据库中存在，则即使未配置任何域，也会映射正确的 UID/GID。

该行为意味着 NFSv4 客户机和服务器之间存在相同的管理用户/组域（也就是说，客户机和服务器上存在具有相同 UID/GID 的用户/组）。如果存在重叠的 ID 空间，传入属性字符串可能会映射到错误的 ID。这在功能上与将传入字符串映射到 `nobody` 没有什么不同，但可提供更大的灵活性。请参见下文的“示例”部分。

允许更改 UID 和 GID 的实用程序包括：

- [usermod\(1M\)](#)
- [userdel\(1M\)](#)
- [groupmod\(1M\)](#)
- [groupdel\(1M\)](#)

文件 `/var/run/nfs4_domain` 包含当前由 NFSv4 使用的域名。

示例 示例1 设置域名

以下命令使用 `sharectl` 设置域名。

```
# sharectl set -p nfsmapid_domain=sun.com nfs
```

下面的“说明”部分中介绍了 `nfsmapid_domain` 属性。

示例2 获取域名

以下命令使用 `sharectl` 获取当前域名。

```
# sharectl get -p nfsmapid_domain nfs
nfsmapid_domain=sun.com
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/ntfsprogs

《Oracle Solaris 管理：网络服务》

引用名	nfsref – NFS 引用实用程序
用法概要	<pre>nfsref [-t <i>svc_type</i>] add <i>path location</i> [<i>location</i> ...] nfsref [-t <i>svc_type</i>] remove <i>path</i> nfsref [-t <i>svc_type</i>] lookup <i>path</i></pre>
描述	nfsref 命令管理 NFS 版本 4 引用，这些引用为服务器端指针，用于将客户机重定向到文件系统的实际位置。引用以重解析点为基础（请参见 reparsed(1M) 和 libreparse(3LIB)）。所有该用法概要形式的路径参数都是指重解析点符号链接的路径。此命令当前实现了一个“服务类型”—nfs-basic，它是缺省值。 第一种形式的命令 <code>nfsref add</code> 可创建指向指定位置的引用。如果不存在重解析点，将会创建一个。如果存在重解析点，将添加 NFS 服务数据或替换现有 NFS 服务数据。每个位置都具有 <code>host :/path</code> 格式。路径可包含空格，且空格必须进行转义以确保正确进行 shell 解析。 第二种形式的命令 <code>nfsref delete</code> 可删除 NFS 引用。它将从指定的重解析点删除 NFS 服务数据，如果不存在其他类型的服务数据，它将删除重解析点。 第三种形式的命令 <code>nfsref lookup</code> 可显示指定的 NFS 引用所指向的位置。
选项	所支持的选项如下： <pre>-t <i>svc_type</i></pre> 指定服务类型。当前，仅支持 <code>nfs-basic</code>。
示例	示例 1 添加 NFS 引用 以下命令将在服务器路径 <code>/pool/home/bob</code> 中创建 NFS 引用，以指向资源 <code>homeserver:/homepool/bob</code>。 <pre># nfsref add /pool/home/bob homeserver:/homepool/bob Created reparse point /pool/home/bob</pre> <pre># nfsref lookup /pool/home/bob homeserver:/homepool/bob</pre> 示例 2 删除 NFS 引用 以下命令将在服务器路径 <code>/pool/home/bob</code> 中删除 NFS 引用。 <pre># nfsref remove /pool/home/bob homeserver:/homepool/bob Removed svc_type 'nfs-basic' from /pool/home/bob</pre>
退出状态	<pre>0</pre> 成功完成。 <pre>>0</pre> 出现错误。

文件 /usr/lib/reparsed/*.so.1
 reparsed 的每服务插件。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/file-system/nfs
接口稳定性	Consolidation Private（整合专用）

另请参见 [reparsed\(1M\)](#)、[libreparsed\(3LIB\)](#)、[attributes\(5\)](#)

引用名	nfsstat – NFS 统计信息
用法概要	<pre>nfsstat [-cnrsza] [-T u d] [-v version] [interval [count]] nfsstat -m [pathname]...</pre>
描述	nfsstat 显示关于 NFS 和 RPC（Remote Procedure Call，远程过程调用）、内核接口的统计信息。它还可用于对此信息进行重新初始化。如果没有给出选项，缺省设置如下： <pre>nfsstat -csnra</pre> 缺省情况下显示全部信息，但是不进行重新初始化。
选项	-a 显示 NFS_ACL 信息。 -c 显示客户机信息。仅输出客户端 NFS、RPC 和 NFS_ACL 信息。可以结合 -n、-r 和 -a 选项一起使用，仅输出客户端 NFS、RPC 和 NFS_ACL 信息。 -m [pathname...] 显示每个 NFS 挂载文件系统的统计信息。如果不指定 <i>pathname</i>，显示所有 NFS 挂载文件系统的统计信息。如果指定了 <i>pathname</i>，显示 <i>pathname</i> 指示的 NFS 挂载文件系统的统计信息。 这包括服务器名称和地址、挂载标志、当前读取和写入大小、重新传输计数、属性高速缓存超时值、故障转移信息，以及用于动态重新传输的计时器。动态重新传输计时器仅在使用动态重新传输时显示。缺省情况下，通过 TCP 协议的 NFS 挂载和通过 TCP 或 UDP 的 NFS 版本 3 挂载不使用动态重新传输。 如果指定 -m 选项，这将是 nfsstat 使用的唯一选项。如果除 -m 外还指定了其他选项，您将接收到一条错误消息，警告您 -m 标志无法与其他选项一起使用。 -n 显示 NFS 信息。输出客户端和服务端端的 NFS 信息。可以与 -c 和 -s 选项一起使用，仅输出客户机或服务端的 NFS 信息。 -r 显示 RPC 信息。 -s 显示服务器信息。 -T u d 显示时间戳。 指定 u 表示时间的内部表示形式的印刷表示形式。请参见 time(2)。指定 d 表示标准日期格式。请参见 date(1)。

-v *version*

指定要输出其统计信息的 NFS 版本。后跟可选 *version* 参数 (2|3|4) 时，可以指定该版本的统计信息。缺省情况下，输出所有版本的统计信息。

-z

将统计信息归零（重新初始化）。该选项仅由超级用户使用，且可以与以上任何选项一起使用，在输出完特定统计信息集合后将其归零。

操作数

支持下列操作数：

count

仅显示计数报告

interval

每隔一定间隔（秒）报告一次。

pathname

指定要显示统计信息的 NFS 挂载文件系统中文件的路径名。

显示

服务器 RPC 显示屏包括以下字段：

badcalls

RPC 层拒绝的调用总数（badlen 和 xdr call 的总和，如下面所定义）。

badlen

长度小于最小 RPC 调用的 RPC 调用数。

calls

接收到的 RPC 调用总数。

dupchecks

在重复请求高速缓存中查找的 RPC 调用数。

dupreqs

发现重复的 RPC 调用数。

nullrecv

RPC 调用在认为已被接收时却无法执行的次数。

xdr call

头无法被 XDR 解码的 RPC 调用数。

服务器 NFS 显示屏显示收到 (calls) 和拒绝 (badcalls) 的 NFS 调用数，以及进行的各种调用的计数和百分比。

服务器 NFS_ACL 显示屏显示进行的各种调用的计数和百分比。

客户机 RPC 显示屏包括以下字段：

calls

进行的 RPC 调用总数。

badcalls

RPC 层拒绝的调用总数。

badverfs

调用由于响应中检验器出错而失败的次数。

badxids

从服务器接收到的、不对应于任何未完成调用的回复次数。

cantconn

调用由于无法连接到服务器而失败的次数。

cantsend

客户机尝试却无法通过无连接传输发送 RPC 请求的次数。

interrupts

调用在完成前被信号中断的次数。

newcreds

必须刷新验证信息的次数。

nomem

调用由于无法分配内存而失败的次数。

retrans

由于等待服务器回复而超时，必须重新传输调用的次数。仅适用于无连接传输的 RPC。

timeouts

调用等待服务器回复时超时的次数。

timers

计算的超时值大于等于为调用指定的最小超时值的次数。

客户机 NFS 显示屏显示发送和拒绝调用的次数，以及接收 CLIENT 句柄的次数 (clgets)，CLIENT 句柄高速缓存没有未使用条目的次数 (cltoomany)，以及各种调用的计数及其各自百分比。

客户机 NFS_ACL 显示屏显示进行的各种调用的计数和百分比。

-m 选项包括有关由挂载选项设置的挂载标志、系统的内部挂载标志的信息以及其他挂载信息。请参见 [mount_nfs\(1M\)](#)。

以下挂载标志由挂载选项设置：

forcedirectio

直接在客户机和服务器之间传输的数据，无需在客户机上缓冲。

grpuid

BSD 组 ID 继承。请参见 [mount_nfs\(1M\)](#) 中的说明。

hard

硬挂载。

intr

硬挂载上允许的中断。

llock

正在使用的本地锁定（无锁管理器）。请注意，这是一个专用接口。

noac

客户机不缓存属性。

nointr

硬挂载上不允许中断。

nocto

没有关闭-打开一致性。

retrans

NFS 重新传输。

rpctimesync

RPC 时间同步。

rsize

读取缓冲区大小（以字节为单位）。

sec

sec 具有下列值之一：

dh

des 式验证（加密时间戳）。

krb5

kerberos v5 式验证。

krb5i

kerberos v5 式完整性验证。

krb5p

kerberos v5 式保密性验证。

none

无验证。

short

Short hand UNIX 式验证。

sys

UNIX 式验证（UID、GID）。

soft

软挂载。

timeo

初始 NFS 超时值：十分之一秒。

wsiz

写入缓冲区大小（以字节为单位）。

以下挂载标志为系统内部使用：

acl

服务器支持 NFS_ACL。

down

服务器关机。

dynamic

动态传输大小调整。

link

服务器支持链接。

mirrormount

通过 mirrormount 机制自动挂载。

printed

输出 "Not responding"（无响应）消息。

readdironly

使用 readdir 而非 readdirplus。

referral

通过引用机制自动挂载。

symlink

服务器支持符号链接。

以下标志与其他挂载信息有关：

proto

协议。

vers

NFS 版本。

-m 选项还提供属性高速缓存超时值。-m 输出中的以下字段提供属性高速缓存的超时值：

acdirmax

保留高速缓存目录属性的最长秒数。

acdirmin

保留高速缓存目录属性的最短秒数。

- acregmax**
保留高速缓存文件属性的最长秒数。
- acregmin**
保留高速缓存文件属性的最短秒数。
- m** 输出中的以下字段提供故障转移信息：
 - currserver**
当前提供 NFS 服务的服务器。有关更多详细信息，请参见 《Oracle Solaris 管理：IP 服务》。
 - failover**
已经选择某台新服务器的次数。
 - noresponse**
服务器已响应失败的次数。
 - remap**
文件已重新评估到新服务器的次数。

下面显示的 **-m** 输出中的字段提供有关动态重新传输的信息。仅在使用动态重新传输时才显示这些项。

- cur**
当前的后退重新传输值，以毫秒为单位。
- dev**
估计偏差，以毫秒为单位。
- srtt**
平滑往返时间值，以毫秒为单位。

退出状态 将返回下列退出值：

- 0**
成功完成。
- >0**
出现错误。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/ntfsprogs

另请参见 [mount_nfs\(1M\)](#)、[attributes\(5\)](#)

- 《安装 Oracle Solaris 11 系统》
- 《Oracle Solaris 管理：IP 服务》

引用名	nscadm – 网络存储控制实用程序
用法概要	<pre>nscadm freeze device nscadm unfreeze device nscadm isfrozen device</pre>
描述	nscadm 命令可执行多个网络存储控制功能。 nscadm freeze 命令可关闭对指定设备的现有引用，并阻止未来的访问。这可在不关闭系统的情况下执行虚拟卷设备驱动程序（例如 RAID 0、RAID 1、RAID 5）的维护。 nscadm unfreeze 命令可针对指定设备执行与 nscadm freeze 相反的操作。 nscadm isfrozen 命令可返回指定设备的当前状态。
选项	nscadm 命令支持以下选项。 -h 显示用法菜单。
操作数	nscadm 命令行支持下列操作数：。 <i>device</i> 指定 nscadm 要对其执行操作的存储设备。
退出状态	对于 freeze 和 unfreeze，子命令 nscadm 将返回以下退出值： 0 成功 255 错误 对于 isfrozen 子命令，nscadm 将返回以下退出值： 0 设备当前处于冻结状态。 1 设备当前未冻结。 255 错误
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	storage/avs/avs-cache-management
接口稳定性	Committed（已确定）

另请参见 [scmadm\(1M\)](#)、[attributes\(5\)](#)

引用名 nscd – 名称服务高速缓存守护进程

用法概要 `/usr/sbin/nscd [-f configuration-file] [-g] [-e cachename, yes
| no] [-i cachename]`

描述 nscd 守护进程是一个为大多数名称服务请求提供高速缓存的进程。svc:/system/name-service/cache SMF 服务的服务属性决定了高速缓存守护进程的行为。请参见 [nscd.conf\(4\)](#)。

通过标准 libc 接口（如

[gethostbyname\(3NSL\)](#)、[getipnodebyname\(3SOCKET\)](#)、[gethostbyaddr\(3NSL\)](#)和其他），nscd 为

[passwd\(4\)](#)、[group\(4\)](#)、[hosts\(4\)](#)、[ipnodes\(4\)](#)、[exec_attr\(4\)](#)、[prof_attr\(4\)](#)、[user_attr\(4\)](#)、[ethers\(4\)](#)数据库提供缓存。将不会专门高速缓存影子文件。[getspnam\(3C\)](#) 调用将保持未缓存状态。

每个高速缓存的数据都具有一个单独的生存时间。缺省情况下，修改本地数据库（/etc/hosts、/etc/passwd 等）会导致下次调用 nscd 时此高速缓存无效。

更新和刷新 svc:/system/name-service/cache 选择性依赖的任何服务（如下所示）会导致重新启动 nscd，这会有效地清除所有高速缓存。

- svc:/network/dns/client，请参见 [resolv.conf\(4\)](#)
- svc:/network/nis/client，请参见 [ypbind\(1M\)](#) 和 [ypfiles\(4\)](#)
- svc:/network/ldap/client，请参见 [ldapclient\(1M\)](#)
- svc:/system/name-service/switch，请参见 [nsswitch.conf\(4\)](#)

请参见 [nscd.conf\(4\)](#)。

nscd 还可充当其自己的管理工具。如果 nscd 的实例已在运行，命令将透明地传递至正在运行的版本。

当在启用每用户查找的情况下运行时（请参见 [nscd.conf\(4\)](#)），nscd 将代表发出请求的用户派生一个且仅一个子进程（即每用户 nscd）。每用户 nscd 将使用用户凭证打开到为每用户形式的查找配置的名称信息库的每用户连接。查找将在子进程中执行。结果将缓存在进程中，且仅可用于同一用户。此高速缓存的管理方式与主 nscd 守护进程管理其自己的高速缓存的方式完全相同。每用户 nscd 终止之前，将一直处理来自用户的后续请求。每用户 nscd 使用可配置的非活动生存时间 (time-to-live, TTL) 值，并在非活动 TTL 到期之后终止其自身。

可由主 nscd 创建的每用户 nscd 的最大数目是可配置的（请参见 [nscd.conf\(4\)](#)）。创建最大数目的每用户 nscd 之后，主 nscd 将使用 LRU 算法根据需要终止不太活动的子 nscd。

主 nscd 守护进程创建、监视并管理所有子 nscd。它将在接收用户的第一次每用户查找时创建用户自己的 nscd。启动 nscd 守护进程时，如果启用每用户查找，它将进行检查以确保在准备创建每用户 nscd 之前满足所有条件。停止此守护进程时，它将终止其控制下的所有每用户 nscd。

每用户 `nscd` 使用与主 `nscd` 相同的配置。它们读取并使用同一缺省配置文件或使用 `-f` 命令行选项指定的配置文件。读取配置后，每用户 `nscd` 将在其整个生命周期中使用它。

选项

下面所述的若干选项需要指定 *cachename*。 *cachename* 支持的值包

括：`passwd`、`group`、`hosts`、`ipnodes`、`exec_attr`、`prof_attr`、`user_attr`、`ethers`、`rpc`、`pr` 和 `project`。

-f *configuration-file*

导致 `nscd` 在指定文件中读取其配置数据。该选项已过时，在将来的发行版中将被删除。

-g

将当前配置和统计信息打印到标准输出中。这是非超级用户可执行的唯一选项。

-e *cachename*, *yes|no*

启用或禁用指定的高速缓存。

-i *cachename*

使指定的高速缓存失效。

示例

示例1 停止并重新启动 `nscd` 守护进程。

```
example# svcadm disable system/name-service/cache
```

```
example# svcadm enable system/name-service/cache
```

文件

`/etc/nscd.conf` 已过时。之前确定的高速缓存守护进程的行为

属性

有关下列属性的描述，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见

[svcs\(1\)](#)、[svcadm\(1M\)](#)、[ypbind\(1M\)](#)、[getspnam\(3C\)](#)、[gethostbyname\(3NSL\)](#)、[getipnodebyname\(3NSL\)](#)

附注

`-g` 选项到 `nscd` 的输出可能会发生更改。请不要依赖它作为编程接口。

`nscd` 服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

```
svc:/system/name-service/cache
```

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

保留了过时的服务 `svc:/system/name-service-cache`，以便实现与可能引用该服务的脚本的向后兼容性。它的唯一用途就是用作服务 `svc:/system/name-service/cache` 的 `optional_all` 依赖项。过时的服务名称在将来的发行版中将被删除。

引用名	nscfg – import, export name service configurations
用法概要	<pre> /usr/sbin/nscfg <i>command</i> [<i>cmd_options</i>] [<i>operands</i>] /usr/sbin/nscfg import [-fnvq] <i>FMRI</i> /usr/sbin/nscfg export [-nvq] <i>FMRI</i> /usr/sbin/nscfg unconfig [-nvq] <i>FMRI</i> /usr/sbin/nscfg validate [-vq] <i>FMRI</i> /usr/sbin/nscfg help </pre>
描述	The <code>nscfg</code> utility imports or exports legacy name service configuration files into or out of the SMF repository. Given a valid SMF configuration and corresponding FMRI, <code>nscfg</code> will regenerate the legacy naming service configuration files, such as <code>nsswitch.conf</code>, <code>resolv.conf</code>, <code>nscd.conf</code>, and so forth, into their legacy locations. Alternatively, <code>nscfg</code> can import those same configuration files populating the SMF repository if it is currently unpopulated. The following are SMF services that can be operands to <code>nscfg</code>. In many instances, the legacy configuration file that is replaced by the SMF service is listed. <pre> svc:/system/name-service/switch:default Legacy file: /etc/nsswitch.conf Name service switch configuration (used by nscd). svc:/system/name-service/cache:default Legacy file: /etc/nscd.conf Name service cache (nscd). svc:/network/dns/client:default Legacy file: /etc/resolv.conf DNS naming service. svc:/network/nis/domain:default Legacy file: /etc/defaultdomain Legacy file: /var/yp/binding/\$DOMAIN/* Shared NIS domain configuration. Used by all NIS services. Also (historical) shared use with LDAP naming services. Must be enabled when using <code>nis/client</code> or <code>ldap/client</code>. svc:/network/nis/client:default NIS client naming service (ypbind and related) svc:/network/ldap/client:default Legacy file: /var/ldap/* LDAP client naming service (<code>ldap_cachemgr</code> and related). </pre>


```

svc:/network/nis/server:default
NIS server service (ypserv).

svc:/network/nis/passwd:default
NIS server passwd service (rpc.yppasswd).

svc:/network/nis/xfr:default
NIS server xfr service (ypxfrd)

svc:/network/nis/update:default
NIS server update service (rpc.yupdated)

svc:/system/name-service/upgrade:default
Import legacy configuration files into SMF service.

```

子命令

The `nscfg` utility supports the subcommands described below. Options are described in the context of the subcommands.

`import [-fvq] FMRI`

If none of the SMF repository properties for the specified FMRI are currently populated, import the legacy configuration files associated with the specified FMRI into the SMF repository.

With `-f`, force the repopulation of the SMF repository with the legacy configuration, even if currently populated.

With `-v`, issue verbose progress messages during the requested operation. With `-q`, issue no error or other messages during the requested operation.

`export [-vq] FMRI`

Export the SMF configuration for the specified FMRI to legacy configuration files. This operation removes any existing affected legacy file(s) and generates new one(s) using the SMF configuration.

`unconfig [-vq] FMRI`

Unconfigure the SMF configuration for the specified FMRI. This operation resets the specified FMRI and any existing legacy file(s) to their initial unconfigured state.

With `-v`, issue verbose progress messages during the requested operation. With `-q`, issue no error or other messages during the requested operation.

`validate [-vq] FMRI`

Validate the SMF configuration for the specified FMRI. This operation checks the current configuration and verifies whether any errors exist.

With `-v`, issue verbose progress messages during the requested operation. With `-q`, issue no error or other messages during the requested operation.

示例**示例 1** Importing DNS Client Configuration

The following command imports the DNS client configuration, stored in `resolv.conf`, into the SMF repository.

```
# nscfg import svc:/network/dns/client:default
```

示例 2 Exporting SMF LDAP Client Configuration

The following command exports the SMF LDAP client configuration to the legacy configuration files in `/var/ldap`.

```
# nscfg export svc:/network/ldap/client:default
```

示例 3 Resetting Name Service Switch Configuration

The following command resets the name service switch configuration to its initial unconfigured state. The command generates no output and exits quietly with status only.

```
# nscfg unconfig -q svc:/system/name-service/switch:default
```

示例 4 Validating LDAP Client Configuration

The following command validates the LDAP client configuration for errors or inconsistencies. The command generates no output and exits quietly with status only.

```
# nscfg validate -q svc:/network/ldap/client:default
```

退出状态

- 0
Command successfully executed.
- 1
An error occurred.
- 2
Configuration unmodified, no change necessary.
- 3
No configuration to import.

文件

- `/etc/default/{nss, yppasswd}`
- `/etc/defaultdomain`
- `/etc/nscd.conf`
- `/etc/nsswitch.conf`
- `/etc/resolv.conf`
- `/var/ldap/ldap_client_cred`
- `/var/ldap/ldap_client_file`
- `/var/yp/binding/{nisdomainname}/ypservers`
- `/var/yp/NISLDAPmapping`
- `/var/yp/securenets`
- `/var/yp/updaters`

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Uncommitted

另请参见 [svcs\(1\)](#), [ldapclient\(1M\)](#), [nscd\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [ypinit\(1M\)](#),
[defaultdomain\(4\)](#), [nsswitch.conf\(4\)](#), [resolv.conf\(4\)](#), [ypfiles\(4\)](#), [attributes\(5\)](#)

引用名	nslookup – query Internet name servers interactively
用法概要	nslookup [-option] [name -] [server]
描述	The nslookup utility is a program to query Internet domain name servers. It has two modes: interactive and non-interactive. Interactive mode allows the user to query name servers for information about various hosts and domains or to print a list of hosts in a domain. Non-interactive mode is used to print just the name and requested information for a host or domain.

参数	Interactive mode is entered in the following cases: 1. No arguments are given (the default name server is used). 2. The first argument is a hyphen (-) and the second argument is the host name or Internet address of a name server. Non-interactive mode is used when the name or Internet address of the host to be looked up is given as the first argument. The optional second argument specifies the host name or address of a name server.

Options can also be specified on the command line if they precede the arguments and are prefixed with a hyphen. For example, to change the default query type to host information, and the initial timeout to 10 seconds, type:

```
nslookup -query=hinfo -timeout=10
```

**Interactive
Commands**

host [server]	Look up information for host using the current default server or using server, if specified. If host is an Internet address and the query type is A or PTR, the name of the host is returned. If host is a name and does not have a trailing period, the search list is used to qualify the name. To look up a host not in the current domain, append a period to the name.
server domain	
lserver domain	Change the default server to <i>domain</i>; lserver uses the initial server to look up information about <i>domain</i>, while server uses the current default server. If an authoritative answer can't be found, the names of servers that might have the answer are returned.
root	Not implemented.
finger	Not implemented.
ls	Not implemented.
view	Not implemented.

help

Not implemented.

?

Not implemented.

exit

Exits the program.

set *keyword*[=*value*]

This command is used to change state information that affects the lookups. Valid keywords are:

all

Prints the current values of the frequently used options to set. Information about the current default server and host is also printed.

class=*value*

Change the query class to one of:

IN

the Internet class

CH

the Chaos class

HS

the Hesiod class

ANY

wildcard

The class specifies the protocol group of the information.

(Default = IN; abbreviation = cl)

[no]debug

Turn on or off the display of the full response packet and any intermediate response packets when searching.

(Default = nodebug; abbreviation = [no]deb)

[no]d2

Turn debugging mode on or off. This displays more about what nslookup is doing.

(Default = nod2)

domain=*name*

Sets the search list to *name*.

[no]search
If the lookup request contains at least one period but doesn't end with a trailing period, append the domain names in the domain search list to the request until an answer is received.

(Default = search)

port=*value*
Change the default TCP/UDP name server port to *value*.

(Default = 53; abbreviation = po)

querytype=*value*
type=*value*
Change the top of the information query.

(Default = A; abbreviations = q, ty)

[no]recurse
Tell the name server to query other servers if it does not have the information. (Default = recurse; abbreviation = [no]rec)

retry=*number*
Set the number of retries to number.

timeout=*number*
Change the initial timeout interval for waiting for a reply to number seconds.

[no]vc
Always use a virtual circuit when sending requests to the server.

(Default = novc)

[no]fail
Try the next nameserver if a nameserver responds with SERVFAIL or a referral (nofail) or terminate query (fail) on such a response.

(Default = nofail)

文件 /etc/resolv.conf
resolver configuration file

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	network/dns/bind
Interface Stability	Volatile

另请参见

[dig\(1M\)](#), [host\(1M\)](#), [named\(1M\)](#), [attributes\(5\)](#)

See the BIND 9 Administrator's Reference Manual. As of the date of publication of this man page, this document is available at <https://www.isc.org/software/bind/documentation>.

附注

BIND 9 nslookup is deprecated and not as full featured as its BIND 8 version. For more features and functionality refer to [dig\(1M\)](#).

nslookup and [dig\(1M\)](#) now report “Not Implemented” as NOTIMP rather than NOTIMPL. This will have impact on scripts that are looking for NOTIMPL.

引用名	nsupdate – 动态 DNS 更新实用程序														
用法概要	<pre>nsupdate [-dv] [-y keyname:secret -k keyfile] [-t timeout] [-u udptimeout] [-r udpretries] [filename]</pre>														
描述	nsupdate 实用程序向名称服务器提交动态 DNS 更新请求（如 RFC 2136 中定义）。该实用程序允许从区域添加或删除资源记录，无需手动编辑区域文件。单一更新请求可以包含添加或删除多个资源记录的请求。 nsupdate 或 DHCP 服务器动态控制下的区域不应手动编辑。手动编辑会与动态更新冲突，从而导致数据丢失。 通过 nsupdate 动态添加或删除的资源记录必须在相同区域中。请求将发送到区域 SOA 记录的 MNAME 字段标识的区域主服务器。 事务签名可用于对使用 RFC 2845 中描述的 TSIG 资源记录类型的动态 DNS 更新进行验证。签名依赖于应仅由 nsupdate 和名称服务器知道的共享秘密。当前，TSIG 唯一支持的加密算法是 HMAC-MD5，该算法在 RFC 2104 中定义。一旦为 TSIG 定义其他算法，应用程序在彼此验证时将需要确保它们选择了适当的算法和密钥。例如，适用的 key 和 server 语句将会添加到 <code>/etc/named.conf</code>，这样，名称服务器可以将适当的密钥和算法与将要使用 TSIG 验证的客户机应用程序的 IP 地址关联。nsupdate 实用程序不会读取 <code>/etc/named.conf</code>。 nsupdate 实用程序使用 -y 或 -k 选项提供共享秘密，生成用于验证动态 DNS 更新请求的 TSIG 记录时需要该秘密。这些选项是互斥的。请参见“选项”。														
选项	所支持的选项如下： <table><tr><td>-d</td><td>在调试模式下操作。这可以提供有关提出的更新请求以及从名称服务器接收到的回复的跟踪信息。</td></tr><tr><td>-k keyfile</td><td>从文件 <i>keyfile</i> 读取共享秘密，其名称格式为 <code>K{name}+.+157.+.+{random}.private</code>。由于历史原因，还必须存在文件 <code>K{name}+.+157.+.+{random}.key</code>。</td></tr><tr><td>-r udpretries</td><td>设置 UDP 重试次数。缺省为重试 3 次。如果将 <i>udpretries</i> 设置为零，仅进行一次更新请求。</td></tr><tr><td>-t timeout</td><td>设置中止更新前的 <i>timeout</i> 间隔（以秒为单位）。缺省值是 300 秒。设置为零会禁用超时。</td></tr><tr><td>-u udptimeout</td><td>以秒为单位设置 UDP 重试之间的间隔，缺省设置是 3 秒。如果设置为零，会根据超时 (-t) 和 UDP 重试次数 (-r) 计算间隔时间。</td></tr><tr><td>-v</td><td>使用 TCP 连接。进行批量更新请求时，使用 TCP 连接更合适。缺省情况下，nsupdate 使用 UDP 向名称服务器发送更新请求。</td></tr><tr><td>-y keyname:secret</td><td>从 <i>keyname:secret</i> 生成签名，其中 <i>keyname</i> 是密钥名称，<i>secret</i> 是 base64 编码的共享秘密。</td></tr></table>	-d	在调试模式下操作。这可以提供有关提出的更新请求以及从名称服务器接收到的回复的跟踪信息。	-k keyfile	从文件 <i>keyfile</i> 读取共享秘密，其名称格式为 <code>K{name}+.+157.+.+{random}.private</code> 。由于历史原因，还必须存在文件 <code>K{name}+.+157.+.+{random}.key</code> 。	-r udpretries	设置 UDP 重试次数。缺省为重试 3 次。如果将 <i>udpretries</i> 设置为零，仅进行一次更新请求。	-t timeout	设置中止更新前的 <i>timeout</i> 间隔（以秒为单位）。缺省值是 300 秒。设置为零会禁用超时。	-u udptimeout	以秒为单位设置 UDP 重试之间的间隔，缺省设置是 3 秒。如果设置为零，会根据超时 (-t) 和 UDP 重试次数 (-r) 计算间隔时间。	-v	使用 TCP 连接。进行批量更新请求时，使用 TCP 连接更合适。缺省情况下， nsupdate 使用 UDP 向名称服务器发送更新请求。	-y keyname:secret	从 <i>keyname:secret</i> 生成签名，其中 <i>keyname</i> 是密钥名称， <i>secret</i> 是 base64 编码的共享秘密。
-d	在调试模式下操作。这可以提供有关提出的更新请求以及从名称服务器接收到的回复的跟踪信息。														
-k keyfile	从文件 <i>keyfile</i> 读取共享秘密，其名称格式为 <code>K{name}+.+157.+.+{random}.private</code> 。由于历史原因，还必须存在文件 <code>K{name}+.+157.+.+{random}.key</code> 。														
-r udpretries	设置 UDP 重试次数。缺省为重试 3 次。如果将 <i>udpretries</i> 设置为零，仅进行一次更新请求。														
-t timeout	设置中止更新前的 <i>timeout</i> 间隔（以秒为单位）。缺省值是 300 秒。设置为零会禁用超时。														
-u udptimeout	以秒为单位设置 UDP 重试之间的间隔，缺省设置是 3 秒。如果设置为零，会根据超时 (-t) 和 UDP 重试次数 (-r) 计算间隔时间。														
-v	使用 TCP 连接。进行批量更新请求时，使用 TCP 连接更合适。缺省情况下， nsupdate 使用 UDP 向名称服务器发送更新请求。														
-y keyname:secret	从 <i>keyname:secret</i> 生成签名，其中 <i>keyname</i> 是密钥名称， <i>secret</i> 是 base64 编码的共享秘密。														

不建议使用 `-y` 选项，因为共享秘密作为命令行参数以明文形式提供，可以在 [ps\(1\)](#) 的输出中或用户 `shell` 维护的历史文件中看到。

输入格式

`nsupdate` 实用程序从 *filename* 读取输入或读取标准输入。每个命令都以单行输入提供。有些命令用于管理目的。其他则是对区域内容的更新说明或先决条件检查。这些检查设置某些名称或资源记录集合 (RRset) 在区域中存在或不存在的条件。要成功执行整个更新请求，必须满足这些条件。如果对先决条件的测试失败，更新将被拒绝。

每个更新请求都包含零或多个先决条件以及零或多个更新。如果某些指定资源记录在区域中存在或缺失，该条件允许适当验证的更新请求继续执行。空白输入行（或 `send` 命令）将累积的命令作为一个动态 DNS 更新请求发送到名称服务器。

命令格式及其意义如下：

`server servername [ port ]`

将全部动态更新请求发送到名称服务器 *servername*。未提供 `server` 语句时，`nsupdate` 将更新发送至正确区域的主服务器。该区域 SOA 记录的 `MNAME` 字段标识该区域的主服务器。*port* 参数是动态更新请求发送到的 *servername* 上的端口号。如果未指定端口号，会使用缺省 DNS 端口号 53。

`local address [ port ]`

使用本地 *address* 发送所有动态更新请求。未提供 `local` 语句时，`nsupdate` 使用系统选择的地址和端口发送更新。*port* 参数也可用于从特定端口提出请求。如果没有指定端口号，系统会分配一个。

`zone zonename`

指定要对区域 *zonename* 做出所有更新。如果未提供 `zone` 语句，`nsupdate` 会根据输入的剩余部分尝试确定要更新的正确区域。

`class classname`

指定缺省类。如果未指定类，缺省类为 `IN`。

`key name secret`

指定所有更新将使用 *name secret* 对进行 TSIG 签名。`key` 命令覆盖通过 `-y` 或 `-k` 选项在命令行上指定的所有密钥。

`prereq nxdomain domain-name`

要求不能存在名称为 *domain-name* 的任何类型的资源记录。

`prereq yxdomain domain-name`

要求存在 *domain-name*（至少有一个资源记录，无论类型如何）。

`prereq nxrrset domain-name [ class ] type`

要求不能存在指定 *type*、*class* 和 *domain-name* 的资源记录。如果省略 *class*，则假定是 `IN` (internet)。

`prereq yxrrset domain-name [ class ] type`

要求必须存在指定 *type*、*class* 和 *domain-name* 的资源记录。如果省略 *class*，则假定是 IN (internet)。

`prereq yxrrset domain-name [ class ] type data...`

来自共享相同 *type*、*class* 和 *domain-name* 的格式的每一组先决条件的 *data* 将组合构成一个 RR 集合。该 RR 集合必须与区域中给定 *type*、*class* 和 *domain-name* 的现有 RR 集合完全匹配。*data* 以资源记录的 RDATA 的标准文本表示来书写。

`update delete domain-name [ ttl ] [ class ] [ type [ data... ] ]`

删除名为 *domain-name* 的任何资源记录。如果提供了 *type* 和 *data*，只会删除匹配的资源记录。如果未提供 *class*，则假定为 internet 类。将忽略 *ttl*，其仅为兼容性而提供。

`update add domain-name ttl [ class ] type data...`

添加具有指定 *ttl*、*class* 和 *data* 的新资源记录。

`show`

显示当前消息，包含上次发送以来指定的所有先决条件和更新。

`send`

发送当前消息。这等效于输入空白行。

`answer`

显示回答。

以分号开始的行是注释，将被忽略。

示例

示例 1 在区域中插入和删除资源记录

以下示例显示如何使用 `nsupdate` 在 `example.com` 区域插入和删除资源记录。请注意，每个示例中的输入包括一个结尾空白行，这样，一组命令会作为一个动态更新请求发送到 `example.com` 主名称服务器。

```
# nsupdate
> update delete oldhost.example.com A
> update add newhost.example.com 86400 A 172.16.1.1
> send
```

`oldhost.example.com` 的所有 A 记录都会删除。会为 IP 地址为 172.16.1.1 的 `newhost.example.com` 添加 A 记录。新添加的记录有 1 天的 TTL（86400 秒）。

示例 2 仅当无记录存在时添加 CNAME

以下命令仅在其不存在记录时添加 CNAME。

```
# nsupdate
> prereq nxdomain nickname.example.com
> update add nickname.example.com 86400 CNAME somehost.example.com
> send
```

示例 2 仅当无记录存在时添加 CNAME （续）

先决条件使名称服务器检查 `nickname.example.com` 是否有任何类型的资源记录。如果有，更新请求失败。如果该名称不存在，会为其添加 `CNAME`。该操作确保添加 `CNAME` 时，它不会与 RFC 1034 中的以下长期存在的规则冲突：如果一个名称作为 `CNAME` 存在，就不能作为任何其他的记录类型存在。（该规则已经为 RFC 4035 中的 DNSSEC 进行了更新，以允许 `CNAME` 具有 `RSIG`、`DNSKEY` 和 `NSEC` 记录。）

文件	<code>/etc/resolv.conf</code>	用于标识缺省名称服务器
	<code>K{name}+.157.#{random}.key</code>	<code>dnssec-keygen(1M)</code> 创建的 HMAC-MD5 密钥的 base-64 编码。
	<code>K{name}+.157.#{random}.private</code>	<code>dnssec-keygen(1M)</code> 创建的 HMAC-MD5 密钥的 base-64 编码

已知问题
TSIG 密钥分别存储在两个单独文件中。这是为加密操作使用 DST 库的 `nsupdate` 结果，在未来的发行版中可能会有所变化。

属性
有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	<code>service/network/dns/bind</code>
接口稳定性	<code>Volatile</code> （可变）

另请参见
[named\(1M\)](#)、[dnssec-keygen\(1M\)](#)、[attributes\(5\)](#)

RFC 2136, RFC 3007, RFC 2104, RFC 2845, RFC 1034, RFC 2535, RFC 2931, RFC 4035

引用名 ntfscat – 在标准输出中显示 NTFS 文件和流

用法概要 ntfscat [*options*] *device* [*file*]

描述 ntfscat 命令读取 NTFS 卷中的文件或流，并在标准输出中显示内容。

传递给 ntfscat 的文件名将忽略大小写。

选项 下面列出了支持的选项。大多数选项都有单字母和全名两种形式。不接受参数的单字母选项可以多个组合到一起。例如，`-fv` 等效于 `-f -v`。全名选项可以缩写为其名称的唯一前缀。

-a, --attribute *type*
显示特定属性类型的内容。缺省情况下，将显示未命名的 `$DATA` 属性。可使用十进制或十六进制格式的编号或者名称指定属性。

Hex	Decimal	Name
0x10	16	"\$STANDARD_INFORMATION"
0x20	32	"\$ATTRIBUTE_LIST"
0x30	48	"\$FILE_NAME"
0x40	64	"\$OBJECT_ID"
0x50	80	"\$SECURITY_DESCRIPTOR"
0x60	96	"\$VOLUME_NAME"
0x70	112	"\$VOLUME_INFORMATION"
0x80	128	"\$DATA"
0x90	144	"\$INDEX_ROOT"
0xA0	160	"\$INDEX_ALLOCATION"
0xB0	176	"\$BITMAP"
0xC0	192	"\$REPARSE_POINT"
0xD0	208	"\$EA_INFORMATION"
0xE0	224	"\$EA"
0xF0	240	"\$PROPERTY_SET"
0x100	256	"\$LOGGED_UTILITY_STREAM"

指定属性名称时可以不带首位上的美元符号 (\$)。如果使用 \$ 符号，必须将名称放在引号内以禁止 shell 解释该名称。

- f, --force**
覆盖一些明显的缺省值，如不使用挂载的卷。请谨慎使用此选项。
- h, --help**
显示选项列表，其中包含每个选项的简短说明。
- i, --inode *num***
通过 inode 编号而不是名称指定文件。
- n, --attribute-name *name***
显示由 *name* 标识的属性。
- q, --quiet**
抑制某些调试、警告和错误消息。

- V, --version
显示版本号、版权和许可信息。
- v, --verbose
显示更多调试、警告和错误消息。

示例

示例1 显示根目录中的文件的内容
以下命令显示 NTFS 卷的根目录中的文件内容。

```
# ntfscat /dev/dsk/c0d0p1 boot.ini
```

示例2 显示子目录中的文件的内容
以下命令显示 NTFS 卷的子目录中的文件内容。

```
# ntfscat /dev/dsk/c0d0p1 /winnt/system32/drivers/etc/hosts
```

示例3 显示属性的内容
以下命令显示根目录 (inode 5) 的 \$INDEX_ROOT 属性的内容。

```
# ntfscat /dev/dsk/c0d0p1 -a INDEX_ROOT -i 5
```

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/ntfsprogs
接口稳定性	Uncommitted（未确定）

另请参见

[ntfsls\(1M\)](#)、[ntfsprogs\(1M\)](#)、[parted\(1M\)](#)、[attributes\(5\)](#)
<http://wiki.linux-ntfs.org>

作者

ntfscat 由 Richard Russon、Anton Altaparmakov 和 Szabolcs Szakacsits 编写。

引用名 ntfclone – 克隆、映像、还原或挽救 NTFS

用法概要

```
ntfclone [options] source
ntfclone --save-image [options] source
ntfclone --resotore-image [options] source
ntfclone --metadata [options] source
```

描述 ntfclone 实用程序可以有效地将 NTFS 文件系统克隆（包括复制、保存、备份和还原操作）或挽救到一个稀疏文件、映像、设备（分区）或到标准输出。它在磁盘扇区级别工作，仅复制写入的数据（即不包括空白空间）。未使用的磁盘空间将变为零（克隆至稀疏文件时）、以控制代码编码（以特殊映像格式保存时）、保留不变（克隆至磁盘/分区时），或以零填充（克隆至标准输出时）。

ntfclone 在制作备份（拍摄 NTFS 文件系统的确切快照）以及在以后还原该备份时很有用。它还可用于测试 NTFS 读取/写入功能，并且使得您可以使用克隆文件排除用户的问题，而不会有破坏原始文件系统的风险。

如果未使用“特殊映像格式”（请参见下面的同名部分），克隆文件将完全是原始 NTFS 文件系统的确切副本，逐个扇区都是相同的。这样，它也可以像原始 NTFS 文件系统一样挂载。例如，如果克隆到一个文件，且内核具有回送设备和 NTFS 支持，那么文件可以通过以下方式挂载：

```
# mount -t ntfs -o loop ntfclone.img
```

Windows 克隆 将系统或引导分区复制、移动或还原到另一台计算机、另一个磁盘或分区时（例如，从 /dev/dsk/c0d0p1 到 /dev/dsk/c0d0p2，从 /dev/dsk/c0d0p1 到 /dev/dsk/c0d1p1 或到另一个磁盘扇区偏移）要特别小心。

大多数情况下，要使 Windows 能够进行引导，您必须将 NTFS 复制、移动或还原到与原始分区和磁盘具有相同的下列特征的分区：

- 从同一扇区开始
- 在相同类型的磁盘上
- 具有相同的 BIOS 传统柱面设置

ntfclone 实用程序保证 NTFS 的确切副本，但是不处理引导问题。这是设计使然：ntfclone 是一个文件系统，而不是一个系统实用程序。其目标仅仅是 NTFS 克隆，而不是 Windows 克隆。因此，可以使用 ntfclone 作为一个非常快速可靠的构造块进行 Windows 克隆，但不能将其用作完整的解决方案。您可以在 NTFS 网站 <http://wiki.linux-ntfs.org> 上找到有关 NTFS 克隆的有用技巧。

稀疏文件 包含未分配块（空洞）的文件称为“稀疏文件”。为此类文件报告的大小始终高于它们在磁盘上实际占用空间的大小。du(1) 命令报告稀疏文件使用的实际磁盘空间。空洞始终读取为零。所有主要的 Linux 文件系统，如 ext2、ext3、reiserfs、Reiser4、JFS 和 XFS 都支持稀疏文件。不过，有些文件系统不支持稀疏文件，ISO 9600 CD-ROM 文件系统就是其中一例。

特殊映像格式	建议您将 NTFS 文件系统保存为特殊映像格式。该格式不将未分配的块表示为空洞，而是使用控制代码对它们进行编码。这样，映像可以保存空间，而不需要稀疏文件支持。映像格式非常适用于通过网络对文件系统映像进行流式传输。特殊映像格式的缺点是，您不能直接挂载映像，而是必须先还原它。 要使用特殊映像格式保存映像，请使用 <code>-s</code> 或 <code>--save-image</code> 选项。要还原映像，请使用 <code>-r</code> 或 <code>--restore-image</code> 选项。注意，通过使用连字符 (-) 作为源文件，可以从标准输入还原映像。
仅元数据克隆	使用 <code>-m</code> 或 <code>--metadata</code> 选项，<code>ntfscclone</code> 可以仅保存 NTFS 元数据，且克隆文件仍然可以挂载。在此类使用中，所有非元数据内容都将丢失；回读数据也只能得到零。 仅元数据的映像可以压缩到很小，通常大小是 1 到 8 MB。传输此类映像用于调查或故障排除非常方便。 在仅元数据模式下，<code>ntfscclone</code> 不保存用户数据，包括嵌入在元数据中的驻留用户数据。它们都填充为零。而且，所有的文件时间戳，元数据内被删除的和未使用的空间都填充为零。因此，此模式不适合某些情况，如司法鉴定。 注意，文件名不会被删除。因为文件名可能包含敏感信息，在发送仅元数据映像之前请考虑破坏安全性或保密性的各种可能。
选项	下面列出了支持的选项。大多数选项都有单字母和全名两种形式。不接受参数的单字母选项可以多个组合到一起。例如，<code>-fv</code> 等效于 <code>-f -v</code>。全名选项可以缩写为其名称的唯一前缀。

`-f`、`--force`

强制 `ntfscclone` 继续，即使文件系统在一致性检查后被标记为“脏”。

`-h`、`--help`

显示选项的列表及每个选项的简要说明。

`-i`、`--ignore-fs-check`

忽略文件系统的一致性检查结果。此选项仅可与 `--meta-data` 选项一起使用。导致不一致的簇将被保留。

`-m`、`--metadata`

仅克隆元数据。使用此选项时，必须仅克隆到一个文件。

`-o`、`--output file`

将 NTFS 克隆到不存在的 *file*。如果 *file* 是一个连字符 (-)，将克隆到标准输出。

`-O`、`--overwrite file`

将 NTFS 克隆到 *file*，如果 *file* 已经存在，则将其覆盖。

`--rescue`

忽略磁盘读取错误，使得有坏扇区的磁盘（如故障磁盘）也可以被挽救且对磁盘造成的影响最小。在该模式下，`ntfscclone` 在最低的扇区级别上工作，从而使更多的数据得以挽救。不可读取的扇区的内容将以问号 (?) 字符填充，此类扇区的开头以字符串 `BadSector` 标记。

-r、--restore-image *source*

从由 *source* 指定的特殊映像格式进行还原。如果 *source* 是一个连字符 (-)，则从标准输入读取映像。

-s、--save-image

保存为特殊映像格式。在空间使用和速度方面，如果在标准输出中生成映像，这是最高效的选项。此选项非常适用于对映像进行压缩、加密或通过网络进行流式传输。

示例

示例1 使用覆盖选项进行克隆

以下命令使用 **--overwrite** 选项进行克隆。

```
# ntfscclone --overwrite /dev/dsk/c0d2p1 /dev/dsk/c0d0p1
```

示例2 保存为特殊映像格式

以下命令将文件系统以特殊映像格式克隆到其原始分区。

```
# ntfscclone --save-image --output backup.img /dev/dsk/c0d0p1
```

示例3 从特殊映像文件还原

以下命令从特殊映像文件还原 NTFS。

```
# ntfscclone --restore-image --overwrite /dev/dsk/c0d0p1 backup.img
```

示例4 保存为压缩的映像

以下命令将 NTFS 保存为压缩的映像文件。

```
# ntfscclone --save-image -o - /dev/dsk/c0d0p1
```

示例5 从压缩的映像还原

以下命令从压缩的映像文件还原 NTFS 卷。

```
# gunzip -c backup.img.gz | \
ntfscclone --restore-image --overwrite /dev/dsk/c0d0p1 -
```

示例6 使用 ssh 备份至远程主机

以下命令使用 **ssh(1)** 备份至远程主机。请注意，ssh 可能需要口令。

```
# ntfscclone --save-image --output - /dev/dsk/c0d0p1 | \
gzip -c | ssh host 'cat > backup.img.gz'
```

示例7 使用 ssh 从远程主机还原

以下命令使用 **ssh(1)** 从远程主机恢复。请注意，ssh 可能需要口令。

```
# ssh host 'cat backup.img.gz' | gunzip -c | \
ntfscclone --restore-image --overwrite /dev/dsk/c0d0p1 -
```


示例 8 从 Web 服务器流式传输映像文件
以下命令从 Web 服务器流式传输映像文件，并将其还原至分区。

```
# wget -q0 - http://server/backup.img | \
ntfscclone --restore-image --overwrite /dev/dsk/c0d0p1 -
```

示例 9 克隆至新文件
以下命令将 NTFS 卷克隆到不存在的文件。

```
# ntfscclone --output ntfs-clone.img /dev/dsk/c0d0p1
```

示例 10 打包 NTFS 元数据
以下命令将 NTFS 元数据打包至映像文件。请注意，bzip2 比 gzip 花费更长时间，但是生成的归档大小不到后者的十分之一。

```
# ntfscclone --metadata --output ntfsmeta.img /dev/dsk/c0d0p1
bzip2 ntfsmeta.img
```

示例 11 解压缩 NTFS 元数据
以下命令将 NTFS 元数据解压缩至稀疏文件。

```
# bunzip2 -c ntfsmeta.img.bz2 | \
cp --sparse=always /proc/self/fd/0 ntfsmeta.img
```

退出状态 如果成功，返回代码为零，如果不成功，则返回代码为非零值。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/ntfsprogs
接口稳定性	Uncommitted（未确定）

另请参见 [du\(1\)](#)、[ssh\(1\)](#)、[ntfsresize\(1M\)](#)、[ntfsundelete\(1M\)](#)、[parted\(1M\)](#)、[attributes\(5\)](#)
<http://wiki.linux-ntfs.org>

作者 ntfscclone 由 Szabolcs Szakacsits 编写，Per Olofsson（特殊映像格式支持）和 Anton Altaparmakov 为其提供了帮助。

引用名	ntfscluster – 标识 NTFS 卷的指定区域中的文件
用法概要	ntfscluster [<i>options</i>] <i>device</i>
描述	ntfscluster 实用程序具有三种操作模式：<code>info</code>、<code>sector</code> 和 <code>cluster</code>，如下所述。 Info 缺省模式 <code>info</code> 当前尚未实现。实现后它将显示 NTFS 卷的一般信息。 Sector <code>sector</code> 模式显示在指定扇区范围内具有数据的文件的列表。此模式通过 <code>--sector</code> 选项实现。 Cluster <code>cluster</code> 模式显示在指定簇范围内具有数据的文件的列表。当簇大小是一个扇区时，等同于 <code>sector</code> 操作模式。此模式通过 <code>--cluster</code> 选项实现。
选项	下面列出了支持的选项。大多数选项都有单字母和全名两种形式。不接受参数的单字母选项可以多个组合到一起。例如，<code>-fv</code> 等效于 <code>-f -v</code>。全名选项可以缩写为其名称的唯一前缀。 <code>-c, --cluster <i>range</i></code> 显示其数据位于此簇范围内的任何文件。 <code>-F, --filename <i>filename</i></code> 显示有关 <i>filename</i> 的信息。 <code>-f, --force</code> 覆盖一些明显的缺省值，如不使用挂载的卷。请谨慎使用此选项。 <code>-h, --help</code> 显示选项列表，其中包含每个选项的简短说明。 <code>-I, --inode <i>num</i></code> 显示有关此 inode 的信息。 <code>-i, --info</code> 此选项尚未实现。 <code>-q, --quiet</code> 抑制某些调试、警告和错误消息。 <code>-s, --sector <i>range</i></code> 显示其数据位于此扇区范围内的任何文件。 <code>-V, --version</code> 显示版本号、版权和许可信息。 <code>-v, --verbose</code> 显示更多调试、警告和错误消息。

示例

示例1 显示有关卷的信息
以下命令将显示有关卷 `/dev/dsk/c0d0p1` 的信息。

```
# ntfscluster /dev/dsk/c0d0p1
```

示例2 显示簇范围内的文件列表
以下命令将在 `/dev/dsk/c0d0p1` 的前 500 个簇中查找文件。

```
# ntfscluster -c 0-500 /dev/dsk/c0d0p1
```

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/ntfsprogs
接口稳定性	Uncommitted（未确定）

另请参见

[ntfsinfo\(1M\)](#)、[ntfsprogs\(1M\)](#)、[parted\(1M\)](#)、[attributes\(5\)](#)
<http://wiki.linux-ntfs.org>

作者

ntfscluster 由 Richard Russon 编写，并包含 Anton Altaparmakov 所做的贡献。

引用名 ntfscmp – 比较两个 NTFS 文件系统并报告差异

用法概要 ntfscmp [*options*] *device1 device2*

描述 ntfscmp 实用程序可比较两个 NTFS 文件系统的所有方面，并报告其发现的所有差异。文件系统可位于块设备中，也可位于映像文件中。ntfscmp 可用于卷验证。但其主要目的是成为一种有效的开发工具，用于快速定位、标识和检查对 NTFS 所做的元数据更改的正确性。

如果仅关注 NTFS 元数据更改，则比较通过使用 ntfscmp(1M) 的 --metadata 选项创建的元数据映像会更方便，因为这样可以消除通常并不感兴趣的时间戳更改。

ntfscmp 的扼要输出是有意的，因为所提供的信息足以确定确切的差异。通过使用 diff(1) 可获取更多输出，以比较每个报告的 inode 的 ntfscmp(1M) 的详细输出。

选项 下面列出了支持的选项。这些选项具有单字母形式和全名形式。

- h, --help
显示帮助并退出。
- P, --no-progress-bar
不显示进度指示条。
- v, --verbose
显示更多调试、警告和错误消息。

退出状态 成功时退出码为 0，否则为非零值。

属性 有关以下属性的说明，请参见 attributes(5)：

属性类型	属性值
可用性	system/file-system/ntfsprogs
接口稳定性	Uncommitted（未确定）

另请参见 diff(1)、ntfscat(1M)、ntfscmp(1M)、ntfsinfo(1M)、ntfsprogs(1M)、parted(1M)、attributes(5)、<http://wiki.linux-ntfs.org>

作者 ntfscmp 由 Szabolcs Szakacsits 编写。

引用名	ntfscp – 将文件复制到 NTFS 卷
用法概要	ntfscp [<i>options</i>] <i>device source_file destination</i>
描述	ntfscp 实用程序用于将文件复制到 NTFS 卷。 <i>destination</i> （请参见“说明”）可以是一个文件，也可以是一个目录。如果 <i>destination</i> 是由名称指定的目录， <i>source_file</i> 将复制到此目录中。如果 <i>destination</i> 是由 inode 编号指定的目录，将为此 inode 创建一个未命名的数据属性，并且 <i>source_file</i> 将复制到此目录中。在通过 inode 编号指定目录之前，请考虑可能的不利后果：目录中包含未命名的数据流是不寻常的。
数据流	NTFS 中的所有数据都存储在流中，这些流可以具有名称。一个文件可以有多个数据流，但必须有且只有一个数据流不带名称。文件的大小是其未命名数据流的大小。通常，未指定流名称时，可设法访问未命名的数据流。如果要访问命名的数据流，需要将 <i>:stream_name</i> 添加至文件名。例如，通过打开 <i>some.mp3:artist</i> ，您将打开 <i>some.mp3</i> 中的 <i>artist</i> 。在禁止访问命名数据流的操作系统（如 Windows）中，需要使用一些如 FAR 的程序或 cygwin 中的实用程序访问这些流。
选项	下面列出了支持的选项。大多数选项都有单字母和全名两种形式。不接受参数的单字母选项可以多个组合到一起。例如，<i>-fv</i> 等效于 <i>-f -v</i>。全名选项可以缩写为其名称的唯一前缀。 <i>-a, --attribute num</i> 写入由 <i>num</i> 指定的属性。 <i>-f, --force</i> 覆盖一些明显的缺省值，如不使用挂载的卷。请谨慎使用此选项。 <i>-h, --help</i> 显示选项的列表及每个选项的简要说明。 <i>-i, --inode</i> 将 <i>destination</i>（请参见“用法概要”部分）视为 inode 编号。 <i>-N, --attr-name name</i> 写入具有此名称的属性。 <i>-n, --no-action</i> 使用此选项在执行实际复制操作之前进行测试运行。卷将以只读方式打开，并且不进行任何写入。 <i>-q, --quiet</i> 抑制某些调试、警告和错误消息。 <i>-V, --version</i> 显示版本号、版权和许可信息。 <i>-v, --verbose</i> 显示更多调试、警告和错误消息。

示例

示例1 从起始目录复制到根目录

以下命令将 new_boot.ini 作为 boot.ini 从 /home/user 复制到 /dev/dsk/c0d0p1 NTFS 卷的根目录中。

```
# ntfscp /dev/dsk/c0d0p1 /home/user/new_boot.ini boot.ini
```

示例2 复制流

以下命令将 myfile 复制到 C:\some\path\myfile:stream 中（假设 /dev/dsk/c0d0p1 驱动器指示符为 c）。

```
# ntfscp -N stream /dev/dsk/c0d0p1 myfile /some/path
```

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/ntfsprogs
接口稳定性	Uncommitted（未确定）

另请参见

[ntfsresize\(1M\)](#)、[ntfsprogs\(1M\)](#)、[parted\(1M\)](#)、[attributes\(5\)](#)

<http://wiki.linux-ntfs.org>

作者

ntfscp 由 Yura Pakhuchiy 编写，并包含 Anton Altaparmakov 和 Hil Liao 所做的贡献。

引用名	ntfsfix – 修复常见错误，并强制操作系统检查 NTFS						
用法概要	ntfsfix [<i>options</i>] <i>device</i>						
描述	ntfsfix 实用程序可修复一些常见的 NTFS 问题。请注意它不是 <code>chkdsk</code> 的某一版本。它将修复一些基本的 NTFS 不一致性，重置 NTFS 日志文件并在下次重新引导操作系统时安排进行 NTFS 一致性检查。 如果认为某个 NTFS 卷已被操作系统或某种其他方式损坏且无法挂载，请对其运行 <code>ntfsfix</code>。						
选项	下面列出了支持的选项。 <code>-h, --help</code> 显示选项列表，其中包含每个选项的简短说明。 <code>-V, --version</code> 显示版本号、版权和许可信息。						
属性	有关以下属性的说明，请参见 attributes(5) ：						
<table><tr><th>属性类型</th><th>属性值</th></tr><tr><td>可用性</td><td>system/file-system/ntfsprogs</td></tr><tr><td>接口稳定性</td><td>Uncommitted（未确定）</td></tr></table>		属性类型	属性值	可用性	system/file-system/ntfsprogs	接口稳定性	Uncommitted（未确定）
属性类型	属性值						
可用性	system/file-system/ntfsprogs						
接口稳定性	Uncommitted（未确定）						
另请参见	mkntfs(1M) 、 ntfsprogs(1M) 、 parted(1M) 、 attributes(5) http://wiki.linux-ntfs.org						
作者	ntfsfix 由 Anton Altaparmakov 编写，并包含 Szabolcs Szakacsits 所做的贡献。						

引用名

ntfsinfo – 转储文件的属性

用法概要

ntfsinfo [options] device

描述

ntfsinfo 实用程序可转储 inode *inode-number* 或文件 *path-filename* 的属性和/或有关 MFT（-m 选项）的信息。运行不带参数的 ntfsinfo 可获取选项的完整列表。

选项

下面列出了支持的选项。大多数选项都有单字母和全名两种形式。不接受参数的单字母选项可以多个组合到一起。例如，-fv 等效于 -f -v。全名选项可以缩写为其名称的唯一前缀。

-F, --file *file*
显示有关 *file* 的信息。

-f, --force
覆盖一些明显的缺省值，如不覆盖现有文件。请谨慎使用此选项。

-h, --help
显示选项列表，其中包含每个选项的简短说明。

-i, --inode *num*
显示由 *num* 标识的 inode 的相关信息。

-m, --mft
显示有关卷的信息。

-q, --quiet
抑制某些调试、警告和错误消息。

-t, --notime
不在输出中显示时间戳。

-V, --version
显示版本号、版权和许可信息。

-v, --verbose
显示更多调试、警告和错误消息。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/ntfsprogs
接口稳定性	Uncommitted（未确定）

另请参见

[ntfsprogs\(1M\)](#)、[ntfsundelete\(1M\)](#)、[parted\(1M\)](#)、[attributes\(5\)](#)
<http://wiki.linux-ntfs.org>

作者

ntfsinfo 由 Matthew J. Fanto、Anton Altaparmakov、Richard Russon、Szabolcs Szakacsits、Yuval Fledel、Yura Pakhuchiy 和 Cristian Klein 编写。

引用名

用法概要

描述

选项

属性

ntfslabel – 显示或更改 NTFS 文件系统上的标签

ntfslabel [options] device [new_label]

ntfslabel 实用程序可显示或更改位于 *device* 中的 NTFS 文件系统上的文件系统标签（请参见“用法概要”部分）。

如果可选参数 *new_label* 不存在，ntfslabel 将显示当前文件系统标签。

如果可选参数 *new_label* 存在，ntfslabel 将文件系统标签设置为 *new_label*。NTFS 文件系统标签长度最多可为 128 个 Unicode 字符；如果 *new_label* 的长度大于 128 个 Unicode 字符，ntfslabel 会将其截断，并显示警告消息。

也可以在创建文件系统时使用 [mkntfs\(1M\)](#) 的 -L 选项设置文件系统标签。

下面列出了支持的选项。大多数选项都有单字母和全名两种形式。不接受参数的单字母选项可以多个组合到一起。例如，-fv 等效于 -f -v。全名选项可以缩写为其名称的唯一前缀。

-f, --force
覆盖一些明显的缺省值，如不使用挂载的卷。请谨慎使用此选项。

-h, --help
显示选项列表，其中包含每个选项的简短说明。

-n, --no-action
不将新标签实际写入磁盘。

-q, --quiet
抑制某些调试、警告和错误消息。

-V, --version
显示版本号、版权和许可信息。

-v, --verbose
显示更多调试、警告和错误消息。

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/ntfsprogs
接口稳定性	Uncommitted（未确定）

另请参见 [mkntfs\(1M\)](#)、[ntfsprogs\(1M\)](#)、[parted\(1M\)](#)、[attributes\(5\)](#)

<http://wiki.linux-ntfs.org>

作者

`ntfslabel` 由 Matthew J. Fanto 编写，并包含 Anton Altaparmakov 和 Richard Russon 所做的贡献。

引用名	ntfsls – 列出 NTFS 文件系统目录内容
用法概要	ntfsls [<i>options</i>] <i>device</i>
描述	ntfsls 实用程序可列出由 PATH 选项（缺省情况下为根目录）指定的文件的相关信息。 <i>device</i> （请参见“用法概要”部分）是对应于设备（例如， <code>/dev/dsk/c0d0p0</code> <i>num</i> ）的特殊文件或 NTFS 映像文件。
选项	下面列出了支持的选项。大多数选项都有单字母和全名两种形式。不接受参数的单字母选项可以多个组合到一起。例如，<code>-fv</code> 等效于 <code>-f -v</code>。全名选项可以缩写为其名称的唯一前缀。 -a, --all 显示所有文件。如果未指定此选项，将不显示 POSIX 名称空间中的文件名。 -F, --classify 将下面所示的指示符之一附加到条目中。 <code>* / = @ </code> -f, --force 强制执行。例如，在存储于正常文件的 NTFS 分区中运行时需要此选项。 -h, --help 显示用法信息并退出。 -i, --inode 显示每个文件的 inode 编号。这是 NTFS 术语中的 MFT 参考号。 -l, --long 使用长列表格式。 -p, --path <i>path</i> 要列出其内容的目录或要显示其相关信息的文件（包括路径）。 -q, --quiet 抑制某些调试、警告和错误消息。 -R, --recursive 显示指定目录下的所有目录的内容。 -s, --system 除非指定此选项，否则将不会全部列出以美元符号开头的文件，因为这些文件通常是系统文件。 -v, --verbose 显示更多调试、警告和错误消息。 -V, --version 显示 ntfsls 版本号并退出。

-x, --dos
显示短文件名（即 DOS 名称空间中的文件），而不是长文件名（即 WIN32 名称空间中的文件）。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/ntfsprogs
接口稳定性	Uncommitted（未确定）

另请参见 [ntfsprogs\(1M\)](#)、[parted\(1M\)](#)、[attributes\(5\)](#)

<http://wiki.linux-ntfs.org>

作者 此版本的 `ntfsls` 由 Lode Leroy、Anton Altaparmakov、Richard Russon、Carmelo Kintana 和 Giang Nguyen 编写。

引用名	ntfsprogs – NTFS 工具列表
用法概要	ntfsprogs
描述	ntfsprogs 是基于共享库的一套 NTFS 实用程序的名称。这些工具可以免费使用，并提供完整的源代码。下面列出了这些工具。 mkntfs(1M) 创建 NTFS 文件系统。 ntfscat(1M) 将文件内容转储至标准输出。 ntfscclone(1M) 有效克隆、备份、恢复或挽救 NTFS。 ntfscluster(1M) 找到使用指定扇区或簇的文件。 ntfscmp(1M) 比较两个 NTFS 文件系统并报告差异。 ntfscp(1M) 将文件复制到 NTFS 卷。 ntfsfix(1M) 检查和修复一些常见错误，清除日志文件并在下次引导操作系统时执行全面检查。 ntfsinfo(1M) 显示有关 NTFS 或其中一个文件或目录的信息。 ntfslabel(1M) 显示或设置 NTFS 文件系统的卷标签。 ntfsls(1M) 列出有关驻留在 NTFS 上的目录中的文件的信息。 ntfsmount(1M)（非 SunOS 手册页） 读写 NTFS 用户空间驱动程序。 ntfsresize(1M) 在不丢失数据的情况下调整 NTFS 的大小。 ntfsundelete(1M) 恢复 NTFS 中的已删除文件。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	system/file-system/ntfsprogs

属性类型	属性值
接口稳定性	Uncommitted（未确定）

另请参见 [parted\(1M\)](#)、[attributes\(5\)](#)
<http://wiki.linux-ntfs.org>

作者 这些工具由 Anton Altaparmakov、Carmelo Kintana、Cristian Klein、Erik Sornes、Giang Nguyen、Holger Ohmacht、Lode Leroy、Matthew J. Fanto、Per Olofsson、Richard Russon、Szabolcs Szakacsits、Yura Pakhuchiy 和 Yuval Fledel 编写。

引用名	ntfsresize – 在不丢失数据的情况下调整 NTFS 文件系统的大小
用法概要	<pre>ntfsresize [options] --info device ntfsresize [options] [--size size[k M G]] device</pre>
描述	ntfsresize 程序可在不丢失数据的情况下安全调整 Windows XP、Windows Server 2003、Windows 2000、Windows NT4 和 Longhorn NTFS 文件系统的大小。支持 32 位和 64 位 Windows“操作系统”使用的所有 NTFS 版本。在调整大小之前不需要取消分段，因为 ntfsresize 可根据需要重新定位任何数据而不会危及数据完整性。 ntfsresize 可用于缩小或扩大位于未挂载设备（通常为磁盘分区）上的任何 NTFS 文件系统。新文件系统将具有指定的大小。大小参数可具有以下可选修饰符之一：k、M、G，分别表示千字节、兆字节或千兆字节。ntfsresize 通过支持 $k=10^3$、$M=10^6$ 和 $G=10^9$ 满足 SI、ATA 和 IEEE 标准以及磁盘制造商的要求。 如果同时省略 <code>---info</code> 和 <code>---size</code> 选项，NTFS 文件系统将扩大至底层设备的大小。 要调整分区中文件系统的大小，必须通过编辑磁盘上的分区表调整文件系统和分区的大小。与其他命令行文件系统大小调整工具类似，ntfsresize 并不处理分区的大小。要处理分区大小，必须使用磁盘分区工具，如 fdisk(1M)。或者，也可以使用一种内部使用 ntfsresize 的用户友好的分区工具。这类分区工具有很多，其中包括 Mandriva 的 DiskDrake、QTParted、SUSE/Novell 的 YaST Partitioner、IBM 的 EVMS、GParted 或 Debian/Ubuntu 的 Partman。 请在使用任何分区工具之前备份您的数据和分区表。对于 NTFS 文件系统，可使用 ntfsclone(1M) 作为一种备份方式。 要缩小 NTFS 分区，请首先使用 ntfsresize 缩小文件系统的大小。然后使用实用程序（如 fdisk(1M)）通过删除分区并重新创建较小分区来缩小分区大小。分区大小不得小于 NTFS 的新大小；否则，您将不能从该分区进行引导。如果错误地创建了过小的分区，您必须重新创建分区，使其与新的 NTFS 文件系统大小相同。 要扩大 NTFS 文件系统，必须首先扩大底层分区的大小。您可以使用 fdisk(1M) 删除分区并重新创建较大的分区。请确保新调整大小的分区不与任何其他分区重叠。然后使用 ntfsresize 扩大文件系统。 重新创建分区时，请确保使用与所替换分区中相同的起始扇区和分区类型。否则，将无法访问您的文件系统。使用 <code>fdisk u</code> 命令从缺省的柱面单元切换至可靠的扇区单元。此外，如果在旧分区中设置了可引导标志，请确保在重新创建的分区中也设置此标志。否则，可能无法从新分区进行引导。
扩展描述	使用 ntfsresize 时存在一些极为少见的限制。例如，存储在具有未知坏扇区的磁盘上的文件系统就会出现此类限制。不支持重定位第一个 MFT 范围以及将大小调整至 \$MFTMirr 范围的中间。检测到这些情况后，调整大小操作将限定为安全大小，或显示最接近的安全大小。

完成大小的调整后，`ntfsresize` 将安排 NTFS 一致性检查。在 Windows 中，此检查由 `chkdsk` 执行。随后首次重新引导进入 Windows 时，您会注意到 `chkdsk` 在蓝色背景下运行。这是正常现象。Windows 可能会在一致性检查之后强制执行快速重新引导。根据您的硬件配置，Windows 可能会警告您系统设置发生更改，并建议或要求进行重新引导。确认此消息并再次进行重新引导。

选项

下面列出了支持的选项。大多数选项都有单字母和全名两种形式。不接受参数的单字母选项可以多个组合到一起。例如，`-fv` 等效于 `-f -v`。全名选项可以缩写为其名称的唯一前缀。

`-b, --bad-sectors`

缺省情况下，`ntfsresize` 会在遇到坏扇区时退出。此选项允许实用程序在遇到此类扇区时仍继续运行。

使用此选项之前，强烈建议您使用 `ntfsclone(1M)` 及 `--rescue` 选项进行备份，然后在 Windows 中从命令行运行 `chkdsk /f /r volume:`。如果磁盘担保显示为有效，请进行更换，因为磁盘有缺陷。请注意，没有软件可以修复坏扇区错误。唯一能做的就是解决这些缺陷。

如果磁盘没有坏扇区，此选项将不起作用。

`-f, --force`

`ntfsresize` 在实际进行（即未使用 `--no-action`）调整大小操作之前始终将文件系统标记为进行一致性检查，这是为了获得额外的安全性。因此，如果 NTFS 文件系统由 `ntfsresize` 标记，使用此选项会很安全。如果需要多次调整大小，而不想在每个调整大小步骤之间引导至 Windows，则必须使用此选项。

`-h, --help`

显示用法信息并退出。

`-i, --info`

在缩小文件系统时使用。使 `ntfsresize` 确定支持的最小的缩小文件系统大小。多数时候，最小大小是文件系统中已使用的空间大小。`ntfsresize` 缩小文件系统后的大小不会小于此选项返回的大小。由于多种因素，它可能无法缩小至这一理论大小。虽然数据的完整性并不存在风险，但仍强烈建议在实际调整大小之前使用 `--no-action` 选项进行测试运行。

根据测试，可获得的最小大小约为文件系统中使用的空间外加 20–200 MB。另请注意，Windows 可能另外需要 50–100 MB 进行安全引导。

此选项不会导致文件系统的任何更改；分区将以只读方式打开。

`-n, --no-action`

使用此选项在执行调整大小操作之前进行测试运行。卷将以只读方式打开，`ntfsresize` 将显示在调整文件系统大小时要执行的操作。仅当测试运行通过时，才继续进行实际的调整大小操作。

`-P, --no-progress-bar`

在 `ntfsresize` 操作期间不显示进度指示条。

- s, --size size[k|M|G]
将文件系统大小调整至 size 字节。新文件系统将具有指定的大小。大小参数可具有以下可选修饰符之一：k、M、G，分别表示千字节、兆字节或千兆字节。ntfsresize 通过支持 $k=10^3$ 、 $M=10^6$ 和 $G=10^9$ 满足 SI、ATA 和 IEEE 标准以及磁盘制造商的要求。在执行实际的调整大小之前，首先运行 ntfsresize 并附带 --no-action 选项和此选项。
- v, --verbose
显示大量输出。
- V, --version
显示 ntfsresize 的版本号。

退出状态 成功时显示零，否则显示非零值。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/ntfsprogs
接口稳定性	Uncommitted（未确定）

另请参见 [fdisk\(1M\)](#)、[ntfsclone\(1M\)](#)、[parted\(1M\)](#)、[attributes\(5\)](#)

<http://wiki.linux-ntfs.org>

附注 在 Linux 版本 2.6 中，对于通过 [parted\(1M\)](#) 处理的分区，使用 ntfsresize 之后会发生分区表损坏，进而导致 Windows 系统不可引导。即使 NTFS 文件系统是一致的，也会发生这种情况。此问题与 NTFS 无关，因此也与 ntfsresize 无关。此外，ntfsresize 从不修改分区表。在上述情况下，可以在 BIOS 中将 Disk Access Mode（磁盘访问模式）更改为 LBA 以重新获得引导功能。有关此情况的进一步讨论，请参见位于 <http://mlf.linux.rulez.org/mlf/ezaz/ntfsresize.html> 的 ntfsresize 常见问题解答。

作者 ntfsresize 由 Szabolcs Szakacsits 编写，Anton Altaparmakov 和 Richard Russon 也对其做出了贡献。

引用名	ntfsundelete – 恢复 NTFS 卷中已删除的文件
用法概要	ntfsundelete [<i>options</i>] <i>device</i>
描述	在适当情况下，ntfsundelete 实用程序可恢复 NTFS 卷中已删除的文件。该命令有三种操作模式： Scan 缺省模式 scan 仅读取 NTFS 卷并查找已删除的文件。然后它将显示一个列表，给出每个已删除文件的 inode 编号、名称和大小。 Undelete “取消删除”模式将获取与正则表达式匹配的文件（选项 -m）或由 <i>inode-expressions</i> 指定的文件，并恢复尽可能多的数据。它将结果保存至另一位置。 Copy “向导”选项。将 MFT 的一部分保存至文件，这在调试 ntfsundelete 时会很有用。 在很多情况下 ntfsundelete 无法恢复文件。例如，请考虑以下情况。当删除某个文件时，MFT 记录标记为未使用，并会更新表示磁盘使用情况的位图。如果未立即关闭电源，此文件用于驻留的空闲空间可能会被覆盖。更糟糕的是，MFT 记录可能会重新用于其他文件。如果发生这种情况，将无法获知此文件在磁盘上的位置。 即使并非文件的所有簇都在使用中，也不能保证它们不会被某些临时文件覆盖。 ntfsundelete 无法恢复压缩文件或加密文件。在扫描过程中，它会将这类文件显示为 0% 可恢复。
语言环境	在 NTFS 中，所有文件名均存储为 Unicode。ntfsundelete 会将文件名转换为当前语言环境以进行显示。该实用程序可成功显示中国象形文件名，然后正确恢复这些文件名。
扩展 MFT 记录	在极少数情况下，单个 MFT 记录不够大，无法容纳描述文件的元数据（仅当文件具有数百个段时才会发生这种情况）。在这种情况下，一条 MFT 记录将存储文件名，另一条 MFT 记录将存储有关数据的信息。ntfsundelete 不会尝试将这类记录合并在一起。它只会列出未命名的文件及数据。
恢复的文件大小和创建日期	要恢复某个文件，ntfsundelete 必须读取该文件的元数据。遗憾的是，删除某个文件时，元数据可能会处于不一致的状态。例如，文件大小可能记录为零；文件的创建日期可能设置为删除它时的时间或随机时间。在这种情况下，ntfsundelete 将选择其发现的最大的文件大小，并将其写入磁盘。它还会尝试将文件的创建日期设置为上次修改的日期。该日期可能是正确的上次修改日期，也可能是其他某个意想不到的日期。
选项	下面列出了支持的选项。大多数选项都有单字母和全名两种形式。不接受参数的单字母选项可以多个组合到一起。例如，-fv 等效于 -f -v。全名选项可以缩写为其名称的唯一前缀。 -b, --byte <i>num</i> 使用 <i>num</i> 表示的字节填入不可恢复的文件簇部分。缺省值为零。

-C, --case

在使用 **--match** 选项尝试进行匹配时使文件名搜索区分大小写。缺省的文件名搜索不区分大小写。

-c, --copy *range*

此“向导”选项将一个 MFT FILE 记录块写入文件中。缺省文件为 **mft**，该文件将在当前目录中创建。此选项可与 **--output** 和 **--destination** 选项结合使用。

-d, --destination *dir*

为 **--copy** 和 **--undelete** 选项指定输出文件的位置。

-f, --force

覆盖一些明显的缺省值，如不覆盖现有文件。请谨慎使用此选项。

-h, --help

显示选项的列表及每个选项的简要说明。

-i, --inodes *range*

恢复指定 inode 编号范围内的文件。*range* 可以是单个 inode 编号、由逗号分隔的多个编号或是由破折号 (-) 分隔的一个范围。

-m, --match *pattern*

通过仅查找与 *pattern* 匹配的文件名过滤输出。模式 (*pattern*) 可包含与一个字符匹配的通配符 **?**，或者与零个或多个字符匹配的通配符 *****。缺省情况下，匹配不区分大小写。要使搜索区分大小写，请使用 **--case** 选项。

-O, --optimistic

恢复文件部分，即使它们当前标记为处于使用状态。

-o, --output *file*

设置由 **--copy** 或 **--undelete** 选项创建的输出文件的名称。

-P, --parent

显示已删除文件的父目录。

-p, --percentage *num*

通过仅匹配具有百分之 *num* 可恢复内容的文件过滤 **--scan** 选项的输出。

-q, --quiet

将输出量减少到最低限度。此选项与 **--scan** 选项一起使用时不起作用。

-s, --scan

搜索 NTFS 卷，并显示可恢复的文件列表。这是 **ntfsundelete** 的缺省操作。此列表可通过分别使用 **--match**、**--size**、**--percent** 和 **--time** 选项按文件名、大小、可恢复百分比或上次修改时间进行过滤。

在此选项的输出中，**%age**（百分比）字段显示可在多大程度上恢复文件。

-S, --size *range*

通过查找文件大小的特定范围过滤 `--scan` 选项的输出。*range* 可指定为以连字符 (-) 分隔的两个数字。大小单位可使用缩写后缀 `k`、`m`、`g` 和 `t`，分别表示千字节、兆字节、千兆字节和兆兆字节。

-t, --time *since*

过滤 `--scan` 选项的输出。仅匹配于此时间之后更改的文件。时间必须采用数字加后缀 `d`、`w`、`m` 或 `y`（分别表示天、周、月或年）的形式。

-T, --truncate

`ntfsundelete` 的缺省行为是将文件大小向上舍入为最接近的簇（为 512 字节的倍数）。如果实用程序具有关于文件大小的完整数据，此选项会将文件恢复至此大小。

-u, --undelete

指定取消删除模式。可通过使用 `--match` 或 `--inodes` 选项指定要恢复的文件。此选项可与 `--output`、`--destination` 和 `--byte` 结合使用。

恢复文件后，如果不使用 `--output` 选项，文件将采用其原始名称。

-v, --verbose

增加 `ntfsundelete` 显示的输出量。

-V, --version

显示 `ntfsundelete` 的版本号、版权以及许可。

示例

示例1 搜索删除的文件

以下命令将搜索特定设备上已删除的文件。

```
# ntfsundelete /dev/dsk/c0d0p1
```

示例2 扫描与通配符匹配的文件

以下命令将搜索与 `*.doc` 匹配的已删除文件。

```
# ntfsundelete /dev/dsk/c0d0p1 -s -m '*.doc'
```

示例3 搜索特定大小的文件

以下命令将在 `/dev/dsk/c0d0p1` 中查找大小介于 5000 和 6000000 字节之间且至少可恢复 90% 数据的已删除文件。

```
# ntfsundelete /dev/dsk/c0d0p1 -S 5k-6m -p 90
```

示例4 搜索最近更改的文件

以下命令将搜索最近两天更改的已删除文件。

```
# ntfsundelete /dev/dsk/c0d0p1 -t 2d
```

示例 5 指定 inode 范围

以下命令将取消删除设备 /dev/sda1 的 inode 2、5 和 100-131。

```
# ntfsundelete /dev/sda1 -u -i 2,5,100-131
```

示例 6 指定输出文件和目录

以下命令将取消删除 inode 编号 3689，将文件命名为 work.doc，并将其存储在用户的起始目录中。

```
# ntfsundelete /dev/dsk/c0d0p1 -u -i 3689 -o work.doc -d ~
```

示例 7 保存 MFT 记录

以下命令将 MFT 记录 3689-3690 保存到文件 debug 中。

```
# ntfsundelete /dev/dsk/c0d0p1 -c 3689-3690 -o debug
```

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/ntfsprogs
接口稳定性	Uncommitted（未确定）

另请参见

[ntfsclone\(1M\)](#)、[ntfsresize\(1M\)](#)、[parted\(1M\)](#)、[attributes\(5\)](#)
<http://wiki.linux-ntfs.org>

作者

ntfsundelete 由 Richard Russon 和 Holger Ohmacht 编写，Anton Altaparmakov 也对此作出了贡献。

引用名 nwamd – 网络自动配置守护进程

用法概要 /lib/inet/nwamd

描述 nwamd 是用于管理网络接口的系统守护进程。

此守护进程会通过 network/physical:default 服务自动启动，不应直接调用。它不构成编程接口的一部分。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Volatile（可变）

另请参见 [svcs\(1\)](#)、[netcfgd\(1M\)](#)、[netadm\(1M\)](#)、[netcfg\(1M\)](#)、[svcadm\(1M\)](#)、[attributes\(5\)](#)、[nwam\(5\)](#)、[smf\(5\)](#)

另请参见 JDS/GNOME 手册页集中的 [nwam-manager\(1M\)](#)。

附注 联网服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

svc:/network/physical:default

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

引用名	obpsym – Kernel Symbolic Debugging for OpenBoot Firmware	
用法概要	<pre>modload -p misc/obpsym</pre>	
描述	obpsym is a kernel module that installs OpenBoot callback handlers that provide kernel symbol information to OpenBoot. OpenBoot firmware user interface commands use the callbacks to convert numeric <i>addresses</i> to kernel symbol names for display purposes, and to convert kernel symbol names to numeric <i>literals</i> allowing symbolic names to be used as input arguments to user interface commands. Once obpsym is installed, kernel symbolic names may be used anywhere at the OpenBoot firmware's user interface command prompt in place of a literal (numeric) string. For example, if obpsym is installed, the OpenBoot firmware commands <code>ct race</code> and <code>dis</code> typically display symbolic names and offsets in the form <i>modname:symbolname + offset</i>. User interface Commands such as <code>dis</code> can be given a kernel symbolic name such as <code>ufs:ufs_mount</code> instead of a numeric address. Placing the command <pre>forceload: misc/obpsym</pre> into the system(4) file forces the kernel module <code>misc/obpsym</code> to be loaded and activates the kernel callbacks during the kernel startup sequence. obpsym may be useful as a kernel debugger in situations where other kernel debuggers are not useful. For example, on SPARC machines, if obpsym is loaded, you may be able to use the OpenBoot firmware's <code>ct race</code> command to display symbolic names in the stack backtrace after a watchdog reset.	
Kernel Symbolic Name Syntax	The syntax for a kernel symbolic name is: <pre>[<i>module-name</i> :] <i>symbol-name</i></pre> Where <i>module-name</i> is the name of the kernel module that the symbol <i>symbol-name</i> appears in. A NULL module name is taken as “all modules, in no particular order” by obpsym. The module name <code>unix</code> is equivalent to a NULL module name, so that conflicts with words defined in the firmware's vocabulary can be avoided. Typically, OpenBoot firmware reads a word from the input stream and looks the word up in its internal <i>vocabulary</i> before checking if the word is a <i>literal</i>. Thus, kernel symbols, such as <code>reset</code> may be given as <code>unix:reset</code> to avoid the unexpected side effect of the firmware finding and executing a matching word in its vocabulary.	
文件	<pre>/etc/system</pre> <pre>/platform/<i>platform-name</i>/kernel/misc/obpsym</pre>	System configuration information file. Platform-specific kernel symbol information.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/library/processor

另请参见 [kldb\(1\)](#), [uname\(1\)](#), [kernel\(1M\)](#), [modload\(1M\)](#), [modunload\(1M\)](#), [system\(4\)](#), [attributes\(5\)](#)

《OpenBoot 2.x Command Reference Manual》

警告 Some OpenBoot firmware user interface commands may use system resources incompatibly with the way they are used by the Unix kernel. These commands and the use of this feature as a kernel debugger may cause interactions that the Unix kernel is not prepared to deal with. If this occurs, the Unix kernel and/or the OpenBoot firmware user interface commands may react unpredictably and may panic the system, or may hang or may cause other unpredictable results. For these reasons, the use of this feature is only minimally supported and recommended to be used only as a kernel debugger of last resort.

If a breakpoint or watchpoint is triggered while the console frame buffer is powered off, the system can crash and be left in a state from which it is difficult to recover. If one of these is triggered while the monitor is powered off, you will not be able to see the debugger output.

附注 *platform-name* can be found using the `-i` option of [uname\(1\)](#)

obpsym is supported only on architectures that support OpenBoot firmware.

On some systems, OpenBoot must be completely RAM resident so the obpsym symbol callback support can be added to the firmware, if the firmware doesn't include support for the symbol callbacks. On these systems, obpsym may complain that it requires that “you must use ramforth to use this module”.

See the 《OpenBoot 2.x Command Reference Manual》 for details on how to use the *ramforth* command, how to place the command into *nvrामrc*, and how to set *use-nvrामrc?* to `true`. On systems with version 1.x OpenBoot firmware, *nvrामrc* doesn't exist, and the *ramforth* command must be typed manually after each reset, in order to use this module.

Once installed, the symbol table callbacks can be disabled by using the following OpenBoot firmware command:

```
0 0 set-symbol-lookup
```

引用名

oplhpd – Hot plug daemon for SPARC Enterprise Server line

用法概要

/usr/platform/SUNW,SPARC-Enterprise/lib/sparcv9/lib/oplhpd

描述

The hot plug daemon for SPARC Enterprise Servers is a daemon process that runs on the SUNW,SPARC-Enterprise family of servers. The daemon is started by the service management facility (see [smf\(5\)](#)) and communicates with the service processor when hot plug PCI cassettes change their dynamic reconfiguration state.

The service FMRI for oplhpd is:

```
svc:/platform/sun4u/oplhpd:default
```

A domain supports only one running oplhpd process at a time.

错误

OPLHPD uses [syslog\(3C\)](#) to report status and error messages. All of the messages are logged with the LOG_DAEMON facility.

Error messages are logged with the LOG_ERR and LOG_NOTICE priorities, and informational messages are logged with the LOG_DEBUG priority. The default entries in the `/etc/syslog.conf` file log all of the OPLHPD error messages to the `/var/adm/messages` log.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/domain-configuration/sparc-enterprise, SUNWdcsr
Interface Stability	Committed

另请参见

[svcs\(1\)](#), [inetadm\(1M\)](#), [svcadm\(1M\)](#), [syslog\(3C\)](#), [syslog.conf\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

引用名	parted – 分区操作程序
用法概要	<code>parted [options] [device [command [options...]]...]</code>
描述	parted 是对磁盘进行分区和调整分区大小的程序。通过它可以创建、销毁、移动和复制 <code>ext2</code>、<code>linux-swaps</code>、<code>FAT</code>、<code>FAT32</code> 和 <code>reiserfs</code> 分区以及调整其大小。它可以创建和移动 <code>Macintosh HFS</code> 分区并调整其大小，还可以检测 <code>jfs</code>、<code>ntfs</code>、<code>ufs</code> 和 <code>xfs</code> 分区。它非常适用于为新操作系统创建空间、重新组织磁盘使用量以及将数据复制到新的硬盘。 本手册页只是 parted 的简要说明。完整的 parted 文档以“GNU Info”格式随软件包一起提供。 parted 通过一组顶级选项和一组子命令实现的，它们中的大多数都具有自己的选项和操作数。下面描述了这些子命令。parted 具有一个可选操作数： <i>device</i> 要使用的块设备。如果未指定，parted 将使用它找到的第一个块设备。 如果在不使用参数的情况下调用 parted，该程序会显示命令提示。
选项	所支持的选项如下： -a alignment-type, --align alignment-type 为新创建的分区设置对齐。有效的对齐类型包括： none 使用磁盘类型所允许的最小对齐。 cylinder 将分区对齐到柱面。 minimal 使用磁盘拓扑信息提供的最小对齐。该值和 <i>opt</i> 值将使用磁盘提供的布局信息，将逻辑分区表地址对齐到磁盘上的实际物理块。<i>min</i> 值是将分区正确对齐到物理块所需的最小对齐，这可避免性能降低。 optimal 使用磁盘拓扑信息提供的最佳对齐。此值可在保证最佳性能的情况下对齐到物理块大小的倍数。 -h, --help 显示帮助消息。 -l, --list 列出所有块设备上的分区布局。 -m, --machine 显示计算机可解析的输出。 -s, --script 从不提示用户介入。

-v, --version
显示版本号。

子命令

如果在 **parted** 命令行中忽略某个子命令，该实用程序将给出一个命令提示符。

check partition
对 *partition* 进行简单检查。

cp [source-device] source dest
将 *source-device*（如果未指定其他设备，则为当前设备）上的源分区的文件系统复制到当前设备上的 *dest* 分区。

help command
显示一般帮助，或有关某个命令的帮助（如果指定命令）。

mkfs partition fs-type
在分区上创建文件系统 *fs-type*。*fs-type* 可以是 *fat16*、*fat32*、*ext2*、*linux-wap* 或 *reiserfs* 其中之一。

mklabel label-type
创建一个类型为 *label-type* 的新磁盘标签（分区表）。*label-type* 应该是 *bsd*、*dhv*、*gpt*、*loop*、*mac*、*msdos*、*pc98* 或 *sun* 其中之一。

mkpart part-type [fs-type] start end
创建一个文件系统为 *fs-type*（如果指定）的 *part-type* 分区，以 *start* 开始，以 *end* 结束（缺省情况下，以兆字节为单位）。*fs-type* 可以是 *fat16*、*fat32*、*ext2*、*HFS*、*linux-swap*、*NTFS*、*reiserfs* 或 *ufs* 其中之一。*part-type* 应该是 *primary*、*logical* 或 *extended* 其中之一。

mkpartfs part-type fs-type start end
创建一个文件系统为 *fs-type* 的 *part-type* 分区，以 *start* 开始，以 *end* 结束（缺省情况下，以兆字节为单位）。

不建议使用此子命令。相反，应使用 **mkpart** 创建空分区，然后使用诸如 **mke2fs(8)**（Linux 的一部分）等外部工具创建文件系统。

move partition start end
移动分区，使其以 *start* 开始，以 *end* 结束。请注意，**move** 永远不会更改从设备号。

name partition name
将分区名称设置为 *name*。此选项仅适用于 Mac、PC98 和 GPT 磁盘标签。如有必要，可将名称放在引号中。

print
显示分区表。

quit
从 **parted** 退出。

rescue start end
 挽救位于 *start* 和 *end* 之间的丢失的分区。如果找到一个分区，**parted** 将询问是否要在分区表中为其创建一个条目。

resize partition start end
 调整 *partition* 上文件系统的大小，使其以 *start* 开始，以 *end* 结束（缺省情况下，以兆字节为单位）。

rm partition
 删除 *partition*。

select device
 选择 *device* 作为要编辑的当前设备。*device* 通常应该是 Solaris 或 Linux 硬盘设备，但是如有必要，也可以是分区、软件 **raid** 设备或者 SVM 或 LVM 逻辑卷。

set partition flag state
 将 *partition* 上的 *flag* 状态更改为 *state*。支持的标志包
 括：**boot**、**root**、**swap**、**hidden**、**raid**、**lvm**、**lba** 和 **palo**。*state* 应当是 **on**（开启）或 **off**（关闭。）

unit unit
 将 *unit* 设置为显示位置和大小时使用的单位，以及在没有明确的单位后缀时，用于解释用户提供的位置和大小。*unit* 可以是 **s**（扇区）、**B**（字节）、**kB**、**MB**、**GB**、**TB**、**%**（设备大小的百分比）、**cyl**（柱面）、**chs**（柱面、磁头、扇区）或 **compact**（对于输入为兆字节，对于输出为某种用户友好形式）其中之一。

version
 显示版本信息和版权消息。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/storage/parted
接口稳定性	Uncommitted（未确定）

另请参见 [fdisk\(1M\)](#)、[mkfs\(1M\)](#)、[attributes\(5\)](#)

[info\(1\)](#) 格式的 GNU 分区软件手册对 **parted** 程序进行了详细说明。

附注 当前不能使用 **ext3** 文件系统功能。要管理 **ext3** 类型的文件系统，请使用诸如 [resize2fs\(8\)](#) 或 [mke2fs\(8\)](#)（两者均为 Linux 的一部分）等工具。请注意，一旦 **ext3** 支持最终完成，当前支持的 **ext2** 文件系统将过时。另请注意，**ext3** 支持将包含有待定义的有限功能。请使用诸如 [resize2fs\(8\)](#) 和 [mke2fs\(8\)](#) 等工具管理这些类型的文件系统。

要手动调整 `ext3` 文件系统或分区的大小，请使用 `resize2fs(8)`、`fdisk(8)` 或类似工具。对于 LVM，需要使用 LVM 命令调整 LVM 元素的大小。

作者

此手册页是由 Timshel Knoll 针对 Debian GNU/Linux 系统编写的。这里针对 Solaris 操作系统进行了改编。

引用名	pbind – control and query bindings of processes or LWPs
用法概要	<pre> pbind -b <i>processor_id</i> <i>pid</i> [/lwpid]... pbind [-q] [<i>pid</i> [/lwpid]]... pbind -Q [<i>processor_id</i>]... pbind -u <i>pid</i> [/lwpid]... pbind -U [<i>processor_id</i>]... </pre>
描述	pbind controls and queries bindings of processes and LWPs (lightweight processes) to processors. pbind can also remove processor bindings that were previously established. When an LWP is bound to a processor, it will be executed only by that processor except when the LWP requires a resource that is provided only by another processor. The binding is not exclusive, that is, the processor is free to execute other LWPs as well. Bindings are inherited, so new LWPs and processes created by a bound LWP will have the same binding. Binding an interactive shell to a processor, for example, binds all commands executed by the shell. Superusers may bind or unbind any process or LWP, while other users can bind or unbind any process or LWP for which they have permission to signal, that is, any process that has the same effective user ID as the user.
选项	The following options are supported: -b <i>processor_id</i> Binds all or a subset of the LWPs of the specified processes to the processor <i>processor_id</i>. Specify <i>processor_id</i> as the processor ID of the processor to be controlled or queried. <i>processor_id</i> must be present and on-line. Use the <code>psrinfo</code> command to determine whether or not <i>processor_id</i> is present and on-line. See psrinfo(1M). -q Displays the bindings of the specified processes or of all processes. If a process is composed of multiple LWPs which have different bindings and the LWPs are not explicitly specified, the bindings of only one of the bound LWPs will be displayed. The bindings of a subset of LWPs can be displayed by appending “/lwpids” to the process IDs. Multiple LWPs may be selected using “-” and “,” delimiters. See EXAMPLES. -Q Displays the LWPs bound to the specified list of processors, or all LWPs with processor bindings. For processes composed of multiple LWPs, the bindings of individual LWPs will be displayed. -u Removes the bindings of all or a subset of the LWPs of the specified processes, allowing them to be executed on any on-line processor.

-U

Removes the bindings of all LWPs bound to the specified list of processors, or to any processor if no argument is specified.

操作数

The following operands are supported:

pid

The process ID of the process to be controlled or queried.

lwpid

The set of LWP IDs of the specified process to be controlled or queried. The syntax for selecting LWP IDs is as follows:

2,3,4-8	LWP IDs 2, 3, and 4 through 8
-4	LWPs whose IDs are 4 or below
4-	LWPs whose IDs are 4 or above

processor_id

The processor ID of the processor to be controlled or queried.

示例

示例1 Binding Processes

The following example binds processes 204 and 223 to processor 2:

```
example% pbind -b 2 204 223
process id 204: was 2, now 2
process id 223: was 3, now 2
```

示例2 Unbinding a Process

The following example unbinds process 204:

```
example% pbind -u 204
```

示例3 Querying Bindings

The following example queries bindings. It demonstrates that process 1 is bound to processor 0, process 149 has at least one LWP bound to CPU3, and process 101 has no bound LWPs.

```
example% pbind -q 1 149 101
process id 1: 0
process id 149: 3
process id 101: not bound
```

示例4 Querying LWP Bindings

The following example queries bindings of LWPs. It demonstrates that LWP 1 of process 149 is bound to CPU3, and LWP 2 of process 149 is not bound.

```
example% pbind -q 149/1-2
lwp id 149/1: 3
lwp id 149/2: not bound
```


示例 5 Querying LWP Bindings for Processor 2:
The following example queries all LWPs bound to processor 2:

```
example% pbind -Q 2
lwp id 149/4: 2
lwp id 149/5: 2
```

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTE VALUE
Availability	system/core-os

退出状态 The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

另请参见 [psradm\(1M\)](#), [psrinfo\(1M\)](#), [psrset\(1M\)](#), [processor_bind\(2\)](#), [processor_info\(2\)](#), [sysconf\(3C\)](#), [attributes\(5\)](#)

诊断

pbind: cannot query pid 31: No such process
The process specified did not exist or has exited.

pbind: cannot bind pid 31: Not owner
The user does not have permission to bind the process.

pbind: cannot bind pid 31: Invalid argument
The specified processor is not on-line.

引用名	pcitool – 中断路由工具
用法概要	<pre>/usr/sbin/pcitool -h</pre> x86: <pre>/usr/sbin/pcitool pci@unit-address -i cpu#,ino# all [-r [-c] -w cpu# [-g]] [-v] [-q]</pre> SPARC: <pre>/usr/sbin/pcitool pci@unit-address -i ino# all [-r [-c] -w cpu# [-g]] [-v] [-q]</pre> <pre>/usr/sbin/pcitool pci@unit-address -m msi# all [-r [-c] -w cpu# [-g]] [-v] [-q]</pre>
描述	PCIttool 是一个低级工具，提供一个用于获取和设置中断路由信息的工具。
中断路由	pcitool -i 命令可显示给定结点上 INO 的设备和 CPU 路由信息，并允许将给定 INO 或 INO 组重新路由至特定 CPU。 在 SPARC 平台上，INO 将映射到某个中断 mondo，而一个或多个 MSI/X 将映射到一个 INO。因此，INO 和 MSI/X 可单独重新确定目标。使用 -i 选项可以检索或重新路由给定的 INO；使用 -m 选项可以检索或重新路由 MSI/X。 在 x86 平台上可以指定 <i>cpu#</i>。与 <i>ino#</i> 配合使用可标识一个独占向量。SPARC 平台不支持 <i>cpu#</i> 参数。
所需特权	要访问中断信息，用户必须具有所有特权。一般用户在通过 su(1M) 成为超级用户之后便可以访问中断信息，或者，如果他在 <code>user_attr</code> 文件中被授予“维护和修复”权限，也可以访问中断信息。请参见 user_attr(4) 和 rbac(5) 。
选项	所支持的选项如下： -h 显示命令用法。 -q 没有以消息形式显示任何错误。但是，pcitool 仍会返回 Unix 错误代码。 -r [-c] 显示给定结点上 INO 的设备和 CPU 路由信息。显示每个显示的 INO 的每个设备的设备路径和实例编号。在某些平台上，专用于根联合体的中断由附加到其路径名的字符串 (Internal) 表示。 使用 -c，可转储中断控制器信息。 如果命令行上既未提供 -r，也未提供 -w，则假定使用 -r。请参见示例。 -v 详细输出。

-w *cpu#* [-g]

将给定 INO 或 MSI/X 路由到给定 CPU。显示新的和原始的路由信息。必须指定 INO 或 MSI/X。

在某些平台上（如 x86），单个功能的多个 MSI 中断需要一起重新路由。使用 **-g** 执行此操作。**-g** 选项仅适用于支持的平台，并且仅适用于 MSI 中断组。（可接受 1“组”。）使用 **-g** 时，提供的向量必须是组中编号最小的向量。组的大小在内部确定。请参见示例。

示例

示例 1 显示所有 INO

用于显示 `/pci@0,0` 上所有 INO 的命令为：

```
# pcitool /pci@0,0 -i all
```

示例 2 显示特定 INO 的输出

用于在根结点 `/pci@0,0` 上显示 INO `0x0,0x21` 的命令对于 x86 平台和 SPARC 平台略有不同。

在 x86 平台上：

```
# pcitool /pci@0,0 -i 0,21
0x0,0x21:          mpt                      0
/pci@7b,0/pci1022,7458@11/pci1000,3060@2
```

在 SPARC 平台上：

```
# pcitool /pci@0,0 -i 21
0x0,0x21:          mpt                      0
/pci@7b,0/pci1022,7458@11/pci1000,3060@2
```

上面显示的输出是一个示例，可能会与您的输出有所不同。

示例 3 显示特定 MSI 的输出

下面是用于在根结点 `/pci@0,0` 上显示 MSI `0x1` 的命令和样例输出。

```
# pcitool /pci@0,0 -m 0x1
0x0,0x1: pcieb      0      /pci@7b,0/pci10de,5d@e
```

示例 4 将 INO 从一个 CPU 重新路由到另一个 CPU

将 INO 21 从 CPU 0 成功重新路由到 CPU 1 将生成以下输出。

在 x86 平台上：

```
# pcitool /pci@0,0 -i 0,21 -w 1
0x0,0x21 -> 0x1,0x20
```

在 SPARC 平台上：

```
# pcitool /pci@0,0 -i 21 -w 1
0x0,0x21 -> 0x1,0x20
```

示例 5 将 MSI 从一个 CPU 重新路由到另一个 CPU
将 MSI 1 从 CPU 1 成功重新路由到 CPU 0 将生成以下输出。

```
# pcitool /pci@0,0 -m 1 -w 0
0x1,0x1 -> 0x0,0x1
```

示例 6 重新路由 INO 组
将以 24 开始的 INO 组从 CPU 0 成功重新路由到 CPU 1 将生成以下输出。

在 x86 平台上：

```
# pcitool /pci@0,0 -i 3,24 -w 1 -g
0x3,0x24 => 0x1,0x22
```

在 SPARC 平台上：

```
# pcitool /pci@0,0 -i 24 -w 1 -g
0x3,0x24 => 0x1,0x22
```

退出状态

- 0 无错误。
- EINVAL 传入的参数超出范围、对齐错误或是无效参数。
- ETIME 将中断更改为新 CPU 之前等待暂挂中断解决问题的超时时间。
- EIO 发生了 I/O 错误。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
体系结构	PCI-based systems（基于 PCI 的系统）
可用性	system/management/pcitool
接口稳定性	Volatile（可变）

另请参见

[su\(1M\)](#)、[pci\(4\)](#)、[user_attr\(4\)](#)、[attributes\(5\)](#)、[rbac\(5\)](#)

PCI 规范（可从 www.pcisig.org 获取）

附注

- 所有值均以十六进制输入。
- 并非所有命令均适用于所有平台。

引用名	pginfo – 显示有关处理器组的信息
用法概要	<pre>pginfo [-T] [-p] [-v] [-r string] [-R string] [pg ... -c processor_id ...] pginfo -s [-v] [-r string] [-R string] [pg ... -c processor_id ...] pginfo -c -I [-r string] [-R string] [pg ... -c processor_id ...] pginfo -h</pre>
描述	pginfo 显示有关处理器组 (Processor Group, PG) 分层结构、其内容及其特征的信息。PG 是一组按共同特征分组到一起的 CPU。 操作系统使用 PG 来表示某些 CPU，这些 CPU 共享与性能相关的硬件，如执行流水线、高速缓存等。这些 PG 组织到一个分层结构中，后者构成了计算机的处理器拓扑模型。在此分层结构中，每个 CPU（导线束）具有一个叶 PG，该 PG 表示与其共享大多数硬件的 CPU。叶 PG 的每个连续父项与该 CPU 共享的硬件逐渐减少，直至到达根 PG。根 PG 包含系统中的所有 CPU，并表示彼此之间共享最少硬件的 CPU 组。（有关 PG 分层结构的示例，请参见下面的示例。） 如果计算机没有任何性能相关硬件的共享关系，pginfo 将只显示包含系统中所有 CPU 的根 PG。 缺省情况下，pginfo 会显示系统中每个 PG 的相关信息，包括其 PG ID、共享关系以及联机和脱机 CPU。缺省情况下，它以深度优先顺序显示 PG，并使用缩排帮助显示 PG 如何彼此关联（请参见下文的“示例”部分）。 您可以通过指定选项执行下列操作： ■ 以图形方式显示 PG 分层结构 ■ 列出运行系统上存在的 PG 共享关系 ■ 提供当前 PG 使用信息，按 PG ID、CPU ID 或共享关系指定感兴趣的 PG ■ 指定仅显示 CPU 或 PG ID 此外，-p 选项可用于显示哪些 PG 包含与具有共同物理关系（如 system、chip 和 core）的 CPU 对应的 CPU。这些物理关系描述 CPU 的物理特征，并且有可能会封装与性能相关的处理器共享关系。 如果 pginfo 在获取系统数据快照时系统配置不断更改，pginfo 会显示一条错误消息，并以退出状态 1 终止。
选项	所支持的选项如下： -c <i>processor_id</i> ... 将参数解释为处理器 ID，并仅显示包含指定处理器的 PG 的相关信息。 与 -T 选项结合使用时，此选项将显示的 PG 分层结构限制为仅包含每个指定 CPU 的沿袭。按 PG ID 指定感兴趣的 PG 时不能使用此选项。 -C 仅显示属于 PG 的所有 CPU 的 CPU ID。此选项不能与 -I 选项同时使用。

- h
显示简短帮助消息，并以退出状态 0 退出。
- I
仅显示 PG 的 PG ID。此选项不能与 -c 选项同时使用。
- p
显示与 PG 对应的物理关系。如果某个 PG 具有与整个系统、某个处理器内核或某个芯片相同的 CPU，则会在 PG 共享关系之后的方括号("[]")中显示 `system`、`core` 或 `chip`。
- r *string1, string2,...*
仅显示其共享关系名称与任何指定字符串匹配的 PG 的相关信息。

每个指定字符串可以是完整的关系名称，也可以是一个或多个关系名称的一部分，并且字符串匹配不区分大小写。-s 选项显示的共享关系列表中列出了可能的关系名称。

您可以指定多个 -r 选项，这样可以匹配关系名称中包含任何指定字符串的所有 PG。与 -T 选项结合使用时，此选项将显示的 PG 分层结构限制为仅包含具有指定关系的每个 PG 的沿袭。
- R *string1, string2,...*
仅显示其共享关系名称与指定名称不同的 PG 的相关信息。

字符串匹配与上述 -r 选项相同。可以输入多个 -R 选项。
- s
显示运行系统上支持的指定 PG 的所有共享关系。-v 选项可与此选项结合使用，以获取每个共享关系的 PG 列表。
- T
在所获得的分层结构中，每个 CPU（硬件导线束）的沿袭的排列顺序是：从与 CPU 共享最多硬件的 PG 到与 CPU 共享最少硬件的 PG。如果指定了任何感兴趣的 CPU、PG 或关系，所获得的 PG 分层结构将限制为 PG 分层结构中具有指定 CPU、PG 或关系的 PG 的沿袭。
- v
详细模式。显示有关 PG 的附加信息。在不与 -s、-C 或 -I 结合使用时，此选项相当于同时使用 -T 和 -p 选项。与 -s 选项结合使用时，此选项将给出每个共享关系的 PG 列表。

操作数

通过指定一个或多个对应 ID 或关键字 `all`，可以在命令行上提供以下操作数。多个 ID 可以指定为空格分隔的列表（如 `1 3`）、某个数字范围（如 `5-8`）或者两者（如 `1 3 5-8 13-16`）。不能同时指定 PG 和 CPU。

pg 可以在命令行上通过 PG ID 指定感兴趣的 PG。

processor_id 输入 -c 选项时，可以在命令行上通过 CPU ID 指定感兴趣的 CPU。

如果指定的某个 PG 或 CPU 无效，pginfo 命令将显示一条标准错误消息，说明 ID 无效，并继续处理命令行上指定的其他 PG 或 CPU。如果指定的所有 PG 或 CPU 均无效，pginfo 将以退出状态 2 退出。

示例

在下面的示例中，系统包含一个具有 8 个内核和 32 个导线束的 UltraSPARC T1 处理器芯片。

示例 1 显示每个 PG 的相关信息

以下命令未使用参数，将显示有关每个 PG 的信息。

```
$ pginfo
PG RELATIONSHIP          CPUs
0   System                0-31
3   Data_Pipe_to_memory   0-31
2   Floating_Point_Unit   0-31
1   Integer_Pipeline      0-3
4   Integer_Pipeline      4-7
5   Integer_Pipeline      8-11
6   Integer_Pipeline      12-15
7   Integer_Pipeline      16-19
8   Integer_Pipeline      20-23
9   Integer_Pipeline      24-27
10  Integer_Pipeline      28-31
```

示例 2 显示有关所有共享关系的信息

以下命令显示有关所有共享关系的信息。

```
$ pginfo -s -v
RELATIONSHIP          PGs
-----
System                0
Data_Pipe_to_memory   3
Floating_Point_Unit   2
Integer_Pipeline      1 4-10
```

示例 3 显示 PG 分层结构

以下命令显示有关系统中所有 PG 的一般信息。输出将显示哪些 PG 属于芯片和内核。

```
$ pginfo -p -T
0 (System) CPUs: 0-31
'-- 3 (Data_Pipe_to_memory [system,chip]) CPUs: 0-31
    '-- 2 (Floating_Point_Unit [system,chip]) CPUs: 0-31
        |-- 1 (Integer_Pipeline [core]) CPUs: 0-3
        |-- 4 (Integer_Pipeline [core]) CPUs: 4-7
        |-- 5 (Integer_Pipeline [core]) CPUs: 8-11
        |-- 6 (Integer_Pipeline [core]) CPUs: 12-15
```

示例3 显示 PG 分层结构 (续)

```
|-- 7 (Integer_Pipeline [core]) CPUs: 16-19
|-- 8 (Integer_Pipeline [core]) CPUs: 20-23
|-- 9 (Integer_Pipeline [core]) CPUs: 24-27
'-- 10 (Integer_Pipeline [core]) CPUs: 28-31
```

示例4 显示具有特定标准的列表

以下命令显示与 CPU 0 共享整数流水线的 CPU 列表。该示例还说明了如何使用 -r 选项按共享关系名称过滤 PG。

```
$ pginfo -r integer_pipeline -C -c 0
0 1 2 3
```

示例5 使用选项按特定标准进行排除

以下命令列出除关系为 Integer_Pipeline 的 PG 之外的所有 PG。

```
$ pginfo -R Integer_Pipeline
PG RELATIONSHIP      CPUs
0 System              0-31
3 Data_Pipe_to_memory 0-31
2 Floating_Point_Unit 0-31
```

退出状态

0
成功完成。

1
出现错误。

2
语法无效。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Uncommitted（未确定）

命令行选项和输出为 Uncommitted（未确定）。

另请参见

[pgstat\(1M\)](#)、[attributes\(5\)](#)

引用名	pgstat – 报告处理器组的利用率统计信息
用法概要	<pre>pgstat [-A] [-C] [-p] [-s key -S key] [-t number] [-T u d] [-v] [-r string] [-R string] [-P pg ...] [-c processor_id...] [interval [count]] pgstat -h</pre>
描述	pgstat 显示处理器组 (Processor Group, PG) 的利用率统计信息。PG 是一组按共同特征分组到一起的 CPU。 操作系统使用 PG 来表示某些 CPU，这些 CPU 共享与性能相关的硬件，如执行流水线、高速缓存等。这些 PG 组织到一个分层结构中，后者构成了计算机的处理器拓扑模型。在此分层结构中，每个 CPU（导线束）具有一个叶 PG，该 PG 表示与其共享大多数硬件的 CPU。叶 PG 的每个连续父项与该 CPU 共享的硬件逐渐减少，直至到达根 PG。根 PG 包含系统中的所有 CPU，并表示彼此之间共享最少硬件的 CPU 组。（有关 PG 分层结构的示例，请参见下面的“示例”部分。） 如果计算机没有任何性能相关硬件的共享关系，pgstat 将只显示包含系统中所有 CPU 的根 PG。 缺省情况下，pgstat 执行以下操作： ■ 每隔一秒钟测量一次 PG 分层结构中所有 PG 的硬件和软件利用率。 ■ 按深度优先顺序显示 PG 的利用率，并使用缩排帮助显示 PG 如何彼此关联。 ■ 在每个间隔结束时显示每个 PG 的 ID、共享关系、硬件负载、软件负载和联机 CPU。 间隔和计数可作为参数提供，以指定抽样间隔中的秒数以及测量和显示指定 PG 利用率的次数。 您可以指定选项以进一步调整输出，以特定方式组织输出，以及指定感兴趣的 PG（有关详细信息，请参见下面的“选项”部分）。 不支持给定 PG 的利用率时，将显示连字符 ("-")，利用率不可用时将显示问号 (?)。对于需要使用 CPU 硬件性能计数器测量硬件利用率的系统，可能无法获得硬件利用率，因为 cpc(3CPC) 使用者（如 attributes(5)、cputrack(1)、dtrace(1M)）或者使用 libcpc(3LIB) 的另一个应用程序会使用这些计数器。
选项	所支持的选项如下： -A 当 pgstat 跨越多个间隔运行时显示利用率数据的汇总。 -c processor_id... 显示包含指定 CPU 的 PG 的利用率。可以通过逗号分隔的 CPU ID 列表的形式指定 CPU。使用连字符 ("-") 可以指定连续的 CPU ID 范围（例如，0-3）。 -C 显示每个 PG 中每个 CPU 的利用率。

- h
显示简短帮助消息，并以退出状态 0 退出。
- p
显示与 PG 对应的物理关系。如果某个 PG 具有与整个系统、某个处理器内核或某个芯片相同的 CPU，则会在 PG 共享关系之后的方括号 ("[]") 中显示 `system`、`core` 或 `chip`。
- P *pg...*
显示指定 PG 的利用率。可以通过逗号分隔的 PG ID 列表的形式指定多个 PG。使用连字符 ("-") 可以指定连续的 PG ID 范围（例如，`0-3`）。
- r *string1, string2,...*
仅显示其共享关系名称与任何指定字符串匹配的 PG 的利用率。字符串可以是完整的关系名称，也可以是一个或多个关系名称的一部分。字符串的匹配不区分大小写。

您可以输入多个 -r 选项，这样可以匹配关系名称与任何指定字符串匹配的所有 PG。
- R *string1, string2,...*
仅显示其共享关系名称与指定名称不同的 PG 的相关信息。

字符串匹配与上述 -r 选项相同。可以输入多个 -R 选项。
- s *key*
按指定关键字降序排序输出行。指定关键字可以是以下各项之一：
 - pg*
按 PG ID 排序。
 - hwload*
按硬件利用率排序。
 - swload*
按软件利用率排序。
 - user*
按用户时间排序。
 - sys*
按系统时间排序。
 - idle*
按空闲时间排序。
 - depth*
按 PG 树从根到叶降序排序，深度优先（缺省）。
 - breadth*
按 PG 树从根到叶降序排序，广度优先。

-S *key*

按指定关键字升序排序输出行。可能的关键字值与 -s 选项相同。

-t *number*

显示指定整数的最大 PG 数。

-T *u|d*

采用 Unix 时间（请参见 [time\(2\)](#)）或 [date\(1\)](#) 使用的标准日期格式显示每个抽样间隔的时间戳。

-v

显示每个 PG 的附加信息，包括硬件利用率和容量以及软件用户、系统和空闲时间。

输出标题

pgstat 会显示列标题，下面列出了这些标题及其含义。

PG

处理器组 ID。

RELATIONSHIP

PG 的共享关系。

HW

硬件负载百分比（针对间隔按 UTIL/CAP 计算）。

UTIL

在间隔内 PG 共享硬件组件的硬件利用率。该数字会很大，因此使用 K、M、B 和 T 分别表示千、百万、十亿和万亿。

CAP

在间隔内 PG 共享硬件组件的大约最大可能利用率。该数字会很大，因此使用 K、M、B 和 T 分别表示千、百万、十亿和万亿。

SW

软件负载百分比（按 $(USR + SYS) / (USR + SYS + IDLE)$ 计算）。

USR

间隔时间内软件线程在 PG 的 CPU 中以用户模式运行的时间百分比。

SYS

间隔时间内软件线程在 PG 的 CPU 中以系统模式运行的时间百分比。

IDLE

间隔时间内无任何软件线程在 PG 的 CPU 中运行的时间百分比。

CPUS

PG 中 CPU 的 CPU ID。

示例

在下面的示例中，系统包含一个具有 8 个内核和 32 个导线束的 UltraSPARC T1 处理器芯片。

示例1 显示指定时间段内的利用率
以下命令显示最后两秒内所有 PG 的利用率。

```
$ pgstat 1 2
PG RELATIONSHIP      HW      SW  CPUS
0  System              -    0.4% 0-31
3  Data_Pipe_to_memory -    0.4% 0-31
2  Floating_Point_Unit 0%    0.4% 0-31
1  Integer_Pipeline    0%    0%   0-3
4  Integer_Pipeline    0%    0%   4-7
5  Integer_Pipeline    0%    0%  8-11
6  Integer_Pipeline    0%    0.2% 12-15
7  Integer_Pipeline    0%    0%  16-19
8  Integer_Pipeline    2.8%  2.7% 20-23
9  Integer_Pipeline    0.1%  0.2% 24-27
10 Integer_Pipeline    0%    0%  28-31
PG RELATIONSHIP      HW      SW  CPUS
0  System              -    0.4% 0-31
3  Data_Pipe_to_memory -    0.4% 0-31
2  Floating_Point_Unit 0%    0.4% 0-31
1  Integer_Pipeline    0%    0.2% 0-3
4  Integer_Pipeline    0%    0%   4-7
5  Integer_Pipeline    0%    0%  8-11
6  Integer_Pipeline    0%    0%  12-15
7  Integer_Pipeline    0%    0%  16-19
8  Integer_Pipeline    3.1%  2.5% 20-23
9  Integer_Pipeline    0%    0%  24-27
10 Integer_Pipeline    0%    0.2% 28-31
```

示例2 显示有关整数流水线的信息
以下命令显示最后两秒内利用率最高的两条整数流水线的详细信息。

```
$ pgstat -v -t 2 -r 'Integer_Pipeline' 1 2
PG RELATIONSHIP      HW UTIL  CAP      SW      USR      SYS  IDLE CPUS
1  Integer_Pipeline   0.2% 2.2M 1.4B   0.2%   0.0%   0.2% 99.8% 0-3
4  Integer_Pipeline  13.1% 181M 1.4B  14.9%   0.0%  14.9% 85.1% 4-7
PG RELATIONSHIP      HW UTIL  CAP      SW      USR      SYS  IDLE CPUS
1  Integer_Pipeline   0.2% 1.9M 1.2B   0.2%   0.0%   0.2% 99.8% 0-3
4  Integer_Pipeline  13.1% 163M 1.2B  14.9%   0.0%  14.9% 85.1% 4-7
```

示例3 显示指定时间段内的内核利用率
以下命令显示最后两分钟内内核利用率的信息。

```
$ pgstat -A 60 2
PG RELATIONSHIP      HW      SW  CPUS
0  System              -   56.9% 0-31
3  Data_Pipe_to_memory -   56.9% 0-31
2  Floating_Point_Unit 0.0%   56.9% 0-31
```

示例 3 显示指定时间段内的内核利用率 (续)

1	Integer_Pipeline	36.7%	58.7%	0-3
4	Integer_Pipeline	41.9%	58.3%	4-7
5	Integer_Pipeline	31.0%	58.0%	8-11
6	Integer_Pipeline	30.7%	57.9%	12-15
7	Integer_Pipeline	30.1%	55.8%	16-19
8	Integer_Pipeline	40.2%	54.8%	20-23
9	Integer_Pipeline	35.0%	56.0%	24-27
10	Integer_Pipeline	40.3%	55.8%	28-31
PG	RELATIONSHIP	HW	SW	CPUS
0	System	-	10.7%	0-31
3	Data_Pipe_to_memory	-	10.7%	0-31
2	Floating_Point_Unit	0.0%	10.7%	0-31
1	Integer_Pipeline	9.0%	10.7%	0-3
4	Integer_Pipeline	9.6%	10.8%	4-7
5	Integer_Pipeline	8.6%	9.9%	8-11
6	Integer_Pipeline	10.5%	11.9%	12-15
7	Integer_Pipeline	9.1%	10.4%	16-19
8	Integer_Pipeline	9.6%	10.9%	20-23
9	Integer_Pipeline	8.9%	10.0%	24-27
10	Integer_Pipeline	9.5%	10.7%	28-31

SUMMARY: UTILIZATION OVER 120 SECONDS

		-----HARDWARE-----			-----SOFTWARE-----			
PG	RELATIONSHIP	MIN	AVG	MAX	MIN	AVG	MAX	CPUS
0	System	-	-	-	10.7%	10.7%	56.9%	0-31
3	Data_Pipe_to_memory	-	-	-	10.7%	10.7%	56.9%	0-31
2	Floating_Point_Unit	0.0%	0.0%	0.0%	10.7%	10.7%	56.9%	0-31
1	Integer_Pipeline	9.0%	8.5%	36.7%	10.7%	10.7%	58.7%	0-3
4	Integer_Pipeline	9.6%	9.1%	41.9%	10.8%	10.8%	58.3%	4-7
5	Integer_Pipeline	8.6%	8.1%	31.0%	9.9%	9.9%	58.0%	8-11
6	Integer_Pipeline	10.5%	10.0%	30.7%	11.9%	11.9%	57.9%	12-15
7	Integer_Pipeline	9.1%	8.6%	30.1%	10.4%	10.4%	55.8%	16-19
8	Integer_Pipeline	9.6%	9.1%	40.2%	10.9%	10.9%	54.8%	20-23
9	Integer_Pipeline	8.9%	8.4%	35.0%	10.0%	10.0%	56.0%	24-27
10	Integer_Pipeline	9.5%	8.9%	40.3%	10.7%	10.7%	55.8%	28-31

退出状态

将返回下列退出值：

- 0
成功完成。
- 1
无法从系统获取 PG 信息。
- 2
指定的间隔、计数或所有 CPU、PG 和共享关系无效。

3
语法无效。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Uncommitted（未确定）

命令行选项和输出为 Uncommitted（未确定）。

另请参见 [cputrack\(1\)](#)、[attributes\(5\)](#)、[dtrace\(1M\)](#)、[pginfo\(1M\)](#)、[cpc\(3CPC\)](#)、[libcpc\(3LIB\)](#)、[attributes\(5\)](#)

引用名	picld – PICL daemon
用法概要	/usr/lib/picl/picld
描述	The Platform Information and Control Library (PICL) provides a mechanism to publish platform-specific information for clients to access in a platform-independent way. <code>picld</code> maintains and controls access to the PICL information from clients and plug-in modules. The daemon is started in both single-user and multi-user boot mode. Upon startup, the PICL daemon loads and initializes the plug-in modules. These modules use the <code>libpicltree(3PICLTREE)</code> interface to create nodes and properties in the PICL tree to publish platform configuration information. After the plug-in modules are initialized, the daemon opens the PICL daemon door to service client requests to access information in the PICL tree.
PICL Tree	The PICL tree is the repository of all the nodes and properties created by the plug-in modules to represent the platform configuration. Every node in the PICL tree is an instance of a well-defined PICL class. The name of the base PICL class is <code>picl</code>, which defines a basic set of properties that all nodes in the tree must possess. Two of those properties are <code>name</code> and <code>_class</code>, where <code>name</code> contains the name of the node, and the <code>_class</code> contains the PICL class name of the node. Certain nodes in the PICL tree have well-known names. For example, the name of the root node of the PICL tree is <code>/</code> and the name of the root node of the sub-tree containing platform device nodes is <code>platform</code>.
PICL plug-in Modules	The PICL plug-in modules are shared objects that publish platform-specific data in the PICL tree. They are located in well-known directories so that the daemon can locate and load them. Plug-in modules are located in one of the following plug-in directories depending on the platform-specific nature of the data that they collect and publish: <pre> /usr/platform/`uname -i`/lib/picl/plugins /usr/platform/`uname -m`/lib/picl/plugins </pre> A plug-in module can specify its dependency on another plug-in module using the <code>-l</code> or <code>-R</code> linker option. The plug-ins are loaded by the daemon using <code>dlopen(3C)</code> according to the specified dependencies. Each plug-in module must define a <code>.init</code> section, which is executed when the plug-in module is loaded, to register themselves with the daemon. See <code>picld_plugin_register(3PICLTREE)</code> for additional information on plug-in registration. The plug-in modules use the <code>libpicltree(3PICLTREE)</code> interface to publish nodes and properties in the PICL tree so that clients can access them. When the PICL daemon invokes the initialization routine of the plug-in module, the plug-in collects the platform information and creates nodes and/or properties to represent the configuration in the PICL tree. A plug-in can create additional threads to monitor the platform configuration and update the PICL tree with any changes. This enables a PICL plug-in to operate as a daemon within the PICL framework.

An environmental monitor is an example of a plug-in module that uses a thread to monitor the temperatures and fan speeds of the platform, then publishes the environmental information in the PICL tree so clients can access them.

Clients use the `libpicl(3PICL)` interface to send requests to `picld` for accessing the PICL tree.

退出状态

`picld` does not return an exit status.

文件

`/var/run/picld_door` PICL daemon door
`/usr/lib/picl/picld` PICL daemon

属性

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/picl

另请参见

`svcs(1)`, `svcadm(1M)`, `dlopen(3C)`, `libpicl(3PICL)`, `libpicltree(3PICLTREE)`, `picld_log(3PICLTREE)`, `picld_plugin_register(3PICLTREE)`, `attributes(5)`, `smf(5)`

附注

The `picld` service is managed by the service management facility, `smf(5)`, under the service identifier:

`svc:/system/picl`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using `svcadm(1M)`. The service's status can be queried using the `svcs(1)` command.

引用名

ping – send ICMP (ICMP6) ECHO_REQUEST packets to network hosts

用法概要

```
/usr/sbin/ping host [timeout]

/usr/sbin/ping -s [-l | -U] [-abdLLnrRv] [-A addr_family]
               [-c traffic_class] [-g gateway [-g gateway...]]
               [-N next_hop_router] [-F flow_label] [-I interval]
               [-i interface] [-P tos] [-p port] [-t ttl] host
               [data_size] [npackets]
```

描述

The utility ping utilizes the ICMP (ICMP6 in IPv6) protocol's ECHO_REQUEST datagram to elicit an ICMP (ICMP6) ECHO_RESPONSE from the specified *host* or network *gateway*. If *host* responds, ping will display:

host is alive

...on the standard output and exit. Otherwise, after *timeout* seconds, it will write:

no answer from *host*

The default value of *timeout* is 20 seconds.

When you specify the *s* flag, sends one datagram per second (adjust with *-I*) and prints one line of output for every ECHO_RESPONSE that it receives. ping produces no output if there is no response. In this second form, ping computes round trip times and packet loss statistics; it displays a summary of this information upon termination or timeout. The default *data_size* is 56 bytes, or you can specify a size with the *data_size* command-line argument. If you specify the optional *npackets*, ping sends ping requests until it either sends *npackets* requests or receives *npackets* replies.

When using ping for fault isolation, first ping the local host to verify that the local network interface is running.

选项

The following options are supported:

-A addr_family

Specify the address family of the target host. *addr_family* can be either *inet* or *inet6*. Address family determines which protocol to use. For an argument of *inet*, IPv4 is used. For *inet6*, IPv6 is used.

By default, if the name of a host is provided, not the literal IP address, and a valid IPv6 address exists in the name service database, ping will use this address. Otherwise, if the name service database contains an IPv4 address, it will try the IPv4 address.

Specify the address family *inet* or *inet6* to override the default behavior. If the argument specified is *inet*, ping will use the IPv4 address associated with the host name. If none exists, ping will state that the host is unknown and exit. It does not try to determine if an IPv6 address exists in the name service database.

If the specified argument is *inet6*, ping uses the IPv6 address that is associated with the host name. If none exists, ping states that the host is unknown and exits.

-F *flow_label*

Specify the flow label of probe packets. The value must be an integer in the range from 0 to 1048575. This option is valid only on IPv6.

-D

Turn off fragmentation. For IPv4, this means setting the Don't Fragment bit. For IPv4 and IPv6, this means do not allow fragmentation as the datagrams are sent. If the *data_size* exceeds the MTU, then ping might report that sending failed due to Message too long.

-I *interval*

Turn on the statistics mode and specify the interval between successive transmissions. The default is one second. See the discussion of the -s option.

-L

Turn off loopback of multicast packets. Normally, members are in the host group on the outgoing interface, a copy of the multicast packets will be delivered to the local machine.

-N *next_hop_router*

Specify a next-hop router so that the probe packet goes through the specified router along its path to the target host. This option essentially bypasses the system routing table and leaves the probe packet header unmodified. Only one next-hop router can be specified.

-P *tos*

Set the type of service (*tos*) in probe packets to the specified value. The default is zero. The value must be an integer in the range from 0 to 255. Gateways also in the path can route the probe packet differently, depending upon the value of *tos* that is set in the probe packet. This option is valid only on IPv4.

-R

Record route. Sets the IPv4 record route option, which stores the route of the packet inside the IPv4 header. The contents of the record route are only printed if the -v and -s options are given. They are only set on return packets if the target host preserves the record route option across echos, or the -l option is given. This option is valid only on IPv4.

-U

Send UDP packets instead of ICMP (ICMP6) packets. ping sends UDP packets to consecutive ports expecting to receive back ICMP (ICMP6) PORT_UNREACHABLE from the target host.

-a

ping all addresses, both IPv4 and IPv6, of the multihomed destination. The output appears as if ping has been run once for each IP address of the destination. If this option is used together with -A, ping probes only the addresses that are of the specified address family. When used with the -s option and *npackets* is not specified, ping continuously probes the destination addresses in a round robin fashion. If *npackets* is specified, ping sends *npackets* number of probes to each IP address of the destination and then exits.

- b
Bypass the global IPsec policy and send and receive packets in the clear for this connection only. This option can be used to troubleshoot network connectivity independent of IPsec. Because this option bypasses system-wide policy for this connection, it can only be used by superuser or a user granted the `sys_net_config` privilege.
- c *traffic_class*
Specify the traffic class of probe packets. The value must be an integer in the range from 0 to 255. Gateways along the path can route the probe packet differently, depending upon the value of *traffic_class* set in the probe packet. This option is valid only on IPv6.
- d
Set the `SO_DEBUG` socket option.
- g *gateway*
Specify a loose source route gateway so that the probe packet goes through the specified host along the path to the target host. The maximum number of gateways is 8 for IPv4 and 127 for IPv6. Note that some factors such as the link MTU can further limit the number of gateways for IPv6.
- i *interface_address*
Specify the outgoing interface address to use for multicast packets for IPv4 and both multicast and unicast packets for IPv6. The default interface address for multicast packets is determined from the (unicast) routing tables. *interface_address* can be a literal IP address, for example, `10.123.100.99`, or an interface name, for example, `eri0`, or an interface index, for example `2`.
- l
Use to send the probe packet to the given host and back again using loose source routing. Usually specified with the `-R` option. If any gateways are specified using `-g`, they are visited twice, both to and from the destination. This option is ignored if the `-U` option is used.
- n
Show network addresses as numbers. `ping` normally does a reverse name lookup on the IP addresses it extracts from the packets received. The `-n` option blocks the reverse lookup, so `ping` prints IP addresses instead of host names.
- p *port*
Set the base UDP *port* number used in probes. This option is used with the `-U` option. The default base *port* number is 33434. The `ping` utility starts setting the destination port number of UDP packets to this base and increments it by one at each probe.
- r
Bypass the normal routing tables and send directly to a host on an attached network. If the host is not on a directly attached network, an error is returned. This option can be used to `ping` a local host through an interface that has been dropped by the router daemon. See [in.routed\(1M\)](#).

- s
Send one datagram per second and collect statistics.
- t *tll*
Specify the IPv4 time to live, or IPv6 hop limit, for unicast and multicast packets. The default time to live (hop limit) for unicast packets can be set with the [ipadm\(1M\)](#) `set-prop` subcommand, using the `icmp_ipv4_ttl` property for IPv4 and the `icmp_ipv6_hoplimit` property for IPv6. The default time to live (hop limit) for multicast is one hop. See EXAMPLES. For further information, see [ipadm\(1M\)](#).

注 – You might observe property names that begin with “_” (underscore). These properties are private to a protocol and are subject to change or removal. See [ipadm\(1M\)](#).
- v
Verbose output. List any ICMP (ICMP6) packets, other than replies from the target host.

操作数

host
The network host

示例

示例1 Using ping With IPv6

This example shows ping sending probe packets to all the IPv6 addresses of the host xyz, one at a time. It sends an ICMP6 ECHO_REQUEST every second until the user interrupts it.

```
istanbul% ping -s -A inet6 -a xyz
PING xyz: 56 data bytes
64 bytes from xyz (4::114:a00:20ff:ab3d:83ed): icmp_seq=0. time=0.479 ms
64 bytes from xyz (fec0::114:a00:20ff:ab3d:83ed): icmp_seq=1. time=0.843 ms
64 bytes from xyz (4::114:a00:20ff:ab3d:83ed): icmp_seq=2. time=0.516 ms
64 bytes from xyz (fec0::114:a00:20ff:ab3d:83ed): icmp_seq=3. time=4.943 ms
64 bytes from xyz (4::114:a00:20ff:ab3d:83ed): icmp_seq=4. time=0.485 ms
64 bytes from xyz (fec0::114:a00:20ff:ab3d:83ed): icmp_seq=5. time=2.201 ms
^C
----xyz PING Statistics----
6 packets transmitted, 6 packets received, 0% packet loss
round-trip (ms)  min/avg/stddev = 0.479/1.583/4.943/1.823
```

示例2 Using ipadm to Set Hop Limits

The following commands use [ipadm\(1M\)](#) to set IPv4 and IPv6 hop limits.

```
# ipadm set-prop -p _ipv6_hoplimit=100 icmp
# ipadm set-prop -p _ipv4_ttl=100 icmp
```

退出状态

The following exit values are returned:

0
Successful operation; the machine is alive.

non-zero

An error has occurred. Either a malformed argument has been specified, or the machine was not alive.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	network/ping

另请参见

[ifconfig\(1M\)](#), [in.routed\(1M\)](#), [ipadm\(1M\)](#), [ndd\(1M\)](#), [netstat\(1M\)](#), [rpcinfo\(1M\)](#), [traceroute\(1M\)](#), [attributes\(5\)](#), [icmp\(7P\)](#), [icmp6\(7P\)](#)

引用名	pkg2du – convert driver packages to Driver Update format
用法概要	<pre>/usr/bin/pkg2du [-f] [-v] [-d dir] [-o iso] [-l label] [-r release] pkg [pkg ...]</pre>
描述	The <code>/usr/bin/pkg2du</code> utility takes one or more packages as input and converts them to Driver Update (DU) format. If the <code>-d</code> option is specified, the resulting DU directory tree is placed in the directory <i>dir</i>. If the <code>-o</code> option is specified, a Solaris ISO image of the DU directory tree is written in the file <i>iso</i>. The ISO image can be burned onto CD/DVD using cdwr(1) or cdrecord(1) (not a SunOS man page) and used during Solaris installation. At least one of the <code>-d</code> and <code>-o</code> options must be specified. If both are specified, then both an ISO image and a directory tree are generated.
选项	The following options are supported: <code>-d dir</code> Directory where the DU directory should be created. <code>-o iso</code> Create a Solaris ISO image of the DU directory. <code>-f</code> If <i>dir</i>/DU or <i>iso</i> exists, remove it without asking first. <code>-l label</code> Label/volume name of the ISO image (if <code>-o</code> option is specified). <code>-r release</code> Solaris release number to use. It takes the form of the return from <code>uname -r</code> command, for example, 5.10. If unspecified, the release number of the currently running Solaris is used. <code>-v</code> Verbose. Multiple <code>-v</code> options increase verbosity.
操作数	The following operands are supported: <pre>pkg [pkg...]</pre> One or more packages to be converted to DU format.
示例	示例 1 Creating a DU CD/DVD The following commands create a DU CD or DVD containing packages <code>SUNWfoo</code> and <code>SUNWbar</code>. <pre># /usr/bin/pkg2du -r 5.10 -o my.iso SUNWfoo SUNWbar # /usr/bin/cdrw -i my.iso</pre>
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[cdrw\(1\)](#), [mkbootmedia\(1M\)](#), [attributes\(5\)](#)

[mkisofs\(8\)](#) (not a SunOS man page)

引用名

pkgadd – transfer software packages to the system

用法概要

```
pkgadd [-nv] [-a admin] [-G] [-x proxy]
      [ [-M] -R root_path] [-r response] [-k keystore]
      [-P passwd] [-V fs_file]
      [-d device | -d datastream pkginst | all]
      [pkginst | -Y category [, category]...]

pkgadd -s [-d device | -d datastream pkginst | all]
      [pkginst | -Y category [, category]...]
```

描述

pkgadd transfers the contents of a software package from the distribution medium or directory to install it onto the system. Used without the *-d device* source specifier, pkgadd looks in the default spool directory (*/var/spool/pkg*) for the package. Used with the *-s* option, it writes the package to a spool directory instead of installing it.

The pkgadd utility requires an amount of temporary space the size of the package that is being installed. pkgadd determines which temporary directory to use by checking for the existence of the *\$TMPDIR* environment variable. If *\$TMPDIR* is not defined, pkgadd uses *P_tmpdir* from *stdio.h*. *P_tmpdir* has a default of */var/tmp/*.

Certain unbundled and third-party packages are no longer entirely compatible with the latest version of pkgadd. These packages require user interaction throughout the installation and not just at the very beginning, or require that their request scripts be run as the root user.

To install these older packages (released prior to Solaris 2.4), set the following environment variable: *NONABI_SCRIPTS=TRUE*

As long as this environment variable is set, pkgadd permits keyboard interaction throughout the installation and package request scripts are run as root.

If you have package request scripts that require running as user root (instead of *noaccess* [the default] or user *install*), use the *rscript_alt* parameter in the [admin\(4\)](#) file to make an appropriate selection. See [admin\(4\)](#).

Note that, in Solaris 8 and Solaris 9, the default user when running a request script was either root or nobody, depending on the operating system's patch level. In the current release, the default user is *noaccess*.

When running pkgadd in the global zone (see [zones\(5\)](#)), a package that contains a request script (see [pkgask\(1M\)](#)) is added only to the global zone. The package is not propagated to any current or yet-to-be-installed non-global zone. This behavior mimics the effect of the *-G* option, described below.

Package commands are [largefile\(5\)](#)-aware. They handle files larger than 2 GB in the same way they handle smaller files. In their current implementations, pkgadd, [pkgtrans\(1\)](#) and other package commands can process a datastream of up to 4 GB.

The *-d*, *-Y*, and *pkginst* arguments shown in the SYNOPSIS are described under OPERANDS, following OPTIONS.

选项

The supported options are described as follows. The *-d device* source specifier is described under OPERANDS, below.

-a *admin*

Define an installation administration file, *admin*, to be used in place of the default administration file. The token *none* overrides the use of any *admin* file, and thus forces interaction with the user. Unless a full path name is given, pkgadd first looks in the current working directory for the administration file. If the specified administration file is not in the current working directory, pkgadd looks in the */var/sadm/install/admin* directory for the administration file.

-G

Add package(s) in the current zone only. When used in the global zone, the package is added to the global zone only and is not propagated to any existing or yet-to-be-created non-global zone. When used in a non-global zone, the package(s) are added to the non-global zone only.

This option causes package installation to fail if, in the *pkginfo* file for a package, *SUNW_PKG_ALLZONES* is set to true. See [pkginfo\(4\)](#).

-k *keystore*

Use *keystore* as the location from which to get trusted certificate authority certificates when verifying digital signatures found in packages. If no keystore is specified, then the default keystore locations are searched for valid trusted certificates. See KEYSTORE LOCATIONS for more information.

-M

Instruct pkgadd not to use the *\$root_path/etc/vfstab* file for determining the client's mount points. This option assumes the mount points are correct on the server and it behaves consistently with Solaris 2.5 and earlier releases.

-n

Installation occurs in non-interactive mode. Suppress output of the list of installed files. The default mode is interactive.

-P *passwd*

Password to use to decrypt keystore specified with *-k*, if required. See PASS PHRASE ARGUMENTS for more information about the format of this option's argument.

-r *response*

Identify a file or directory which contains output from a previous [pkgask\(1M\)](#) session. This file supplies the interaction responses that would be requested by the package in interactive mode. *response* must be a full pathname.

-R *root_path*

Define the full path name of a directory to use as the *root_path*. All files, including package system information files, are relocated to a directory tree starting in the specified *root_path*. The *root_path* may be specified when installing to a client from a server (for example, */export/root/client1*).

注 – The root file system of any non-global zones must not be referenced with the -R option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See [zones\(5\)](#).

-s *spool*

Write the package into the directory *spool* instead of installing it.

-v

Trace all of the scripts that get executed by pkgadd, located in the *pkginst/install* directory. This option is used for debugging the procedural and non-procedural scripts.

-V *fs_file*

Specify an alternative *fs_file* to map the client's file systems. For example, used in situations where the *\$root_path/etc/vfstab* file is non-existent or unreliable.

-x *proxy*

Specify a HTTP[S] proxy to use when downloading packages. The format of proxy is *host:port*, where *host* is the hostname of the HTTP[S] proxy, and *port* is the port number associated with the proxy. This switch overrides all other methods of specifying a proxy. See ENVIRONMENT VARIABLES for more information on alternate methods of specifying a default proxy.

When executed without options or operands, pkgadd uses */var/spool/pkg* (the default spool directory).

操作数

The following operands are supported:

Sources

By default, pkgadd looks in the */var/spool/pkg* directory when searching for instances of a package to install or spool. Optionally, the source for the package instances to be installed or spooled can be specified using:

-d *device*

-d *datastream pkgname,...* | *all*

Install or copy a package from *device*. *device* can be any of the following:

- A full path name to a directory or the identifiers for tape or removable medium (for example, */var/tmp*).
- A datastream created by *pkgtrans* (see [pkgtrans\(1\)](#)).
- A URL pointing to a datastream created by *pkgtrans*. The supported Universal Resource Identifiers (URIs) are *http:* and *https:*.

The second form of the -d specifier, above, indicates the syntax you use when specifying a datastream. In this case you must specify either a comma-separated list of package names or the keyword *all*.

Instances

By default, pkgadd searches the specified source, and presents an interactive menu allowing the user to select which package instances found on the source are to be installed. As an alternative, the package instances to be installed can be specified using:

pkginst

The package instance or list of instances to be installed. The token `all` may be used to refer to all packages available on the source medium. The format *pkginst*.`*` can be used to indicate all instances of a package.

The asterisk character (`*`) is a special character to some shells and may need to be escaped. In the C-Shell, the asterisk must be surrounded by single quotes (`'`) or preceded by a backslash (`\`).

-Y category[,category...]

Install packages based on the value of the CATEGORY parameter stored in the package's [pkginfo\(4\)](#) file. All packages on the source medium whose CATEGORY matches one of the specified categories will be selected for installation or spooling.

Keystore Locations Package such as `pkgadd` use a set of trusted certificates to perform signature validation on any signatures found within the packages. If there are no signatures included in the packages then signature validation is skipped. The certificates can come from a variety of locations. If `-k keystore` is specified, and *keystore* is a directory, then *keystore* is assumed to be the base directory of the certificates to be used. If *keystore* is a file, then the file itself is assumed to have all required keys and certificates. When `-k` is not specified, then `/var/sadm/security` is used as the base directory.

Within the specified base directory, the store locations to be searched are different based on the application doing the searching and the type of store being searched for. The following directories are searched in the specified order:

1. `<store_dir>/<app_name>/<store_type>`
2. `<store_dir>/<store_type>`

Where `<store_dir>` is the directory specified by `-k`, `<app_name>` is the name of the application doing the searching, and `<store_type>` is one of `keystore` (for private keys), `certstore` (for untrusted public key certificates), or `truststore` (for trusted certificate authority certificates).

For example, when `pkgadd` is run with `-k /export/certs`, then the following locations are successively searched to find the trust store:

1. `/export/certs/pkgadd/truststore`
2. `/export/certs/truststore`

This searching order enables administrators to have a single location for most applications, and special certificate locations for certain applications.

Keystore And Certificate Formats

The packaging utilities, such as `pkgtrans`, require access to a set of keys and certificates in order to sign, and optionally verify, packages.

The keystore files found by following the search pattern specified in KEYSTORE LOCATIONS must each be a self-contained PKCS#12-format file.

When signing a package with `pkgtrans`, if a `certstore` has more than one public key certificate, then each public key must have a `friendlyName` attribute in order to be identifiable and selectable with the `-a` option when signing packages. In addition, the public key certificate selected with `-a` and found in the `certstore` must have an associated private key in the keystore.

Several browsers and utilities can be used to export and import certificates and keys into a PKCS#12 keystore. For example, a trusted certificate can be exported from Mozilla, and then imported into a PKCS#12 keystore for use with `pkgadd` with the OpenSSL Toolkit.

Pass Phrase Arguments

`pkgtrans` and `pkgadd` accept password arguments, typically using `-p` to specify the password. These allow the password to be obtained from a variety of sources. Both of these options take a single argument whose format is described below. If no password argument is given and a password is required then the user is prompted to enter one: this will typically be read from the current terminal with echoing turned off.

`pass:password`

The actual password is *password*. Because the password is visible to utilities such as `ps` this form should only be used where security is not important.

`env:var`

Obtain the password from the environment variable *var*. Because the environment of other processes is visible on certain platforms this option should be used with caution.

`file:pathname`

The first line contained within *pathname* is the password. *pathname* need not refer to a regular file: it could, for example, refer to a device or named pipe. For example, to read the password from standard input, use `file:/dev/stdin`.

`console`

Read the password from `/dev/tty`.

示例

示例 1 Installing a Package from a Solaris DVD

The following example installs a package from a Solaris DVD. You are prompted for the name of the package you want to install.

```
example# pkgadd -d /cdrom/cdrom0/s0/Solaris_10/Product
```

示例 2 Installing a Set of Packages from a Datastream

The example command shown below installs all of the packages in the datastream specified by the `-d` source specifier. Prior to this command, this datastream must have been created with the `pkgtrans(1)` command.

```
example# pkgadd -d /var/tmp/datastream all
```

The keyword `all` specifies that all of the packages found in the designated datastream will be installed.

退出状态	0 Successful completion 1 Fatal error. 2 Warning. 3 Interruption. 4 Administration. 5 Administration. Interaction is required. Do not use pkgadd -n. 10 Reboot after installation of all packages. 20 Reboot after installation of this package.
环境变量	HTTPPROXY Specifies an HTTP proxy host. Overrides administration file setting, and http_proxy environment variable. HTTPPROXYPORT Specifies the port to use when contacting the host specified by HTTPPROXY. Ignored if HTTPPROXY is not set. http_proxy URL format for specifying proxy host and port. Overrides administration file setting.
文件	/var/sadm/install/logs/ Location where pkgadd logs an instance of software installation.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	package/svr4
Interface Stability	Committed

另请参见 [pkginfo\(1\)](#), [pkgmk\(1\)](#), [pkgparam\(1\)](#), [pkgproto\(1\)](#), [pkgtrans\(1\)](#), [installf\(1M\)](#), [pkgadm\(1M\)](#), [pkgask\(1M\)](#), [pkgchk\(1M\)](#), [pkgrm\(1M\)](#), [removef\(1M\)](#), [admin\(4\)](#), [pkginfo\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#), [zones\(5\)](#)

《Application Packaging Developer’s Guide》

<http://www.openssl.org>

附注

When transferring a package to a spool directory, the `-r`, `-n`, and `-a` options cannot be used.

The `-r` option can be used to indicate a directory name as well as a filename. The directory can contain numerous response files, each sharing the name of the package with which it should be associated. This would be used, for example, when adding multiple interactive packages with one invocation of `pkgadd`. In this situation, each package would need a response file. If you create response files with the same name as the package (for example, `pkinst1` and `pkinst2`), then name the directory in which these files reside after the `-r`.

The `-n` option causes the installation to halt if any interaction is needed to complete it.

If the default *admin* file is too restrictive, the administration file may need to be modified to allow for total non-interaction during a package installation. See [admin\(4\)](#) for details.

If a package stream is specified with `-d`, and a digital signature is found in that stream, the default behavior is to attempt to validate the certificate and signature found. This behavior can be overridden with *admin* file settings. See [admin\(4\)](#) for more information.

引用名	pkgadm – manage packaging and patching system
用法概要	<pre> pkgadm addcert [-ty] [-a <i>app</i>] [-k <i>keystore</i>] [-e <i>keyfile</i>] [-f <i>format</i>] [-n <i>name</i>] [-P <i>passarg</i>] [-p <i>import_passarg</i>] [-R <i>rootpath</i>] <i>certfile</i> pkgadm removecert [-a <i>app</i>] [-k <i>keystore</i>] -n <i>name</i> [-P <i>passarg</i>] [-R <i>rootpath</i>] pkgadm listcert [-a <i>app</i>] [-f <i>format</i>] [-k <i>keystore</i>] -n <i>name</i> [-P <i>passarg</i>] [-o <i>outfile</i>] [-R <i>rootpath</i>] pkgadm dbstatus [-R <i>rootpath</i>] pkgadm sync [-R <i>rootpath</i>] [-q] pkgadm -V pkgadm -?</pre>
描述	The pkgadm utility is used for managing the packaging and patching system. It has several subcommands that perform various operations relating to packaging. The pkgadm command includes subcommands for managing certificates and keys used.
Managing Keys and Certificates	pkgadm maintains the packaging-system-wide keystore in <code>/var/sadm/security</code>, and individual user's certificates in <code>~/.pkg/security</code>. The following subcommands operate on the package keystore database: addcert Add (import) a certificate into the database, with optional trust. Once added, trusted certificates can be used to verify signed packages and patches. Non-trusted user certificates and their associated keys can be used to sign packages and patches. Added user certificates are <i>not</i> used to build certificate chains during certificate verification. removecert Removes a user certificate/private key pair, or a trusted certificate authority certificate from the keystore. Once removed, the certificate and keys cannot be used. listcert Print details of one or more certificates in the keystore. sync Writes the contents file and rolls the contents log file. With use of the <code>-q</code> option, forces the contents file server to quit.
Internal Install Database	The Solaris operating system relies upon enhanced System V revision 4 (SVr4) packages as the basis for its software installation and revision management. The package maintenance software stores information about installed packages in an internal database. The pkgadm subcommand <code>dbstatus</code> is used to determine how the package internal database is implemented. The <code>dbstatus</code> command returns a string that indicates the type of internal database in use. In

the current implementation, the `dbstatus` command always returns the string `text`, which indicates that the [contents\(4\)](#) package database is in use. Future releases of Solaris might supply alternative database implementations.

选项

The following options are supported:

-a *app*

If this option is used, then the command only affects the keystore associated with a particular application. Otherwise, the global keystore is affected.

-e *keyfile*

When adding a non-trusted certificate/key combination, this option can be used to specify the file that contains the private key. If this option is not used, the private key must be in the same file as the certificate being added.

-f *format*

When adding certificates, this specifies the format to expect certificates and private keys in. Possible values when adding are:

`pem`

Certificate and any private key uses PEM encoding.

`der`

Certificate and any private key uses DER encoding.

When printing certificates, this specifies the output format used when printing. Acceptable values for format are:

`pem`

Output each certificate using PEM encoding.

`der`

Output each certificate using DER encoding.

`text`

Output each certificate in human-readable format.

-k *keystore*

Overrides the default location used when accessing the keystore.

-n *name*

Identifies the entity in the store on which you want to operate. When adding a user certificate, or removing certificates, this name is required. The name is associated with the certificate/key combination, and when adding, can be used later to reference the entity. When printing certificates, if no alias is supplied, then all keystore entities are printed.

-o *outfile*

Output the result of the command to *outfile*. Only used when examining (printing) certificates from the key store. Standard out is the default.

- P *passarg*
Password retrieval method to use to decrypt keystore specified with -k, if required. See PASS PHRASE ARGUMENTS in [pkgadd\(1M\)](#) for more information about the format of this option's argument. `console` is the default.
- p *import_passarg*
This option's argument is identical to -P, but is used for supplying the password used to decrypt the certificate and/or private key being added. `console` is the default.
- q
(Applies to sync subcommand.) Shuts down the contents file cache daemon.
- R *rootpath*
Defines the full name of a directory to use as the root (/) path. The default user location of the certificate operations is `${HOME}/.pkg`. If the -R option is supplied, the certificates and keys will be stored under `<altroot>/var/sadm/security`. Note that this operation fails if the user does not have sufficient permissions to access this directory. The `listcert` command requires read permission, while `addcert` and `removecert` require both read and write permission.

注 – The root file system of any non-global zones must not be referenced with the -R option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See [zones\(5\)](#).
- t
Indicates the certificate being added is a trusted CA certificate. The details of the certificate (including the Subject Name, Validity Dates, and Fingerprints) are printed and the user is asked to verify the data. This verification step can be skipped with -y. When importing a trusted certificate, a private key should not be supplied, and will be rejected if supplied. Once a certificate is trusted, it can be used as a trust anchor when verifying future untrusted certificates.
- V
Print version associated with packaging tools.
- y
When adding a trusted certificate, the details of the certificate (Subject name, Issuer name, Validity dates, Fingerprints) are shown to the user and the user is asked to verify the correctness before proceeding. With -y, this additional verification step is skipped.
- ?
Print help message.

操作数

The following operand is supported:

`certfile`

File containing the certificate and optional private key, used when adding a trust anchor or certificate/key combination. Certificates must be encoded using PEM or binary DER.

Keystore Aliases All keystore entries (user cert/key and trusted certificate entries) are accessed via unique aliases. Aliases are case-sensitive.

An alias is specified when you add an entity to a keystore using the `addcert` or `trustcert` subcommand. If an alias is not supplied for a trust anchor, the trust anchor's Common Name is used as the alias. An alias is required when adding a signing certificate or chain certificate. Subsequent `pkgcert` or other package tool commands must use this same alias to refer to the entity.

Keystore Passwords See the [pkgadd\(1M\)](#) man page for a description of the passwords supplied to the `pkgadm` utility.

示例

示例 1 Adding a Trust Anchor

The following example adds a well-known and trusted certificate to be used when verifying signatures on packages.

```
example% pkgadm addcert -t /tmp/certfile.pem
```

示例 2 Adding a Signing Certificate

The following example adds a signing certificate and associated private key, each of which is in a separate file, which can then be used to sign packages.

```
example% pkgadm addcert -a pkgtrans -e /tmp/keyfile.pem \
/tmp/certfile.pem
```

示例 3 Printing Certificates

The following example prints all certificates in the root keystore.

```
example% pkgadm listcert
```

退出状态

```
0
    successful completion

non-zero
    fatal error
```

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	package/svr4
Interface Stability	Committed

另请参见 [pkginfo\(1\)](#), [pkgmk\(1\)](#), [pkgparam\(1\)](#), [pkgproto\(1\)](#), [pkgtrans\(1\)](#), [svcs\(1\)](#), [installf\(1M\)](#), [pkgadd\(1M\)](#), [pkgask\(1M\)](#), [pkgrm\(1M\)](#), [removef\(1M\)](#), [svcadm\(1M\)](#), [admin\(4\)](#), [contents\(4\)](#), [exec_attr\(4\)](#), [pkginfo\(4\)](#), [attributes\(5\)](#), [rbac\(5\)](#), [smf\(5\)](#)

《Application Packaging Developer's Guide》

附注

The service for pkgadm is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/system/pkgserv`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	pkgask – stores answers to a request script
用法概要	pkgask [-d <i>device</i>] [-R <i>root_path</i>] -r <i>response</i> <i>pkginst...</i>
描述	pkgask allows the administrator to store answers to an interactive package (one with a request script, that is, a user-created file that must be named request). Invoking this command generates a response file that is then used as input at installation time. The use of this response file prevents any interaction from occurring during installation since the file already contains all of the information the package needs.
选项	The following options are supported -d <i>device</i> Run the request script for a package on <i>device</i>. <i>device</i> can be a directory pathname or the identifiers for tape or removable medium (for example, /var/tmp and /dev/dsk/c1d0s0). The default device is the installation spool directory. -R <i>root_path</i> Define the full path name of a directory to use as the <i>root_path</i>. All files, including package system information files, are relocated to a directory tree starting in the specified <i>root_path</i>. 注 – The root file system of any non-global zones must not be referenced with the -R option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See zones(5). -r <i>response</i> Identify a file or directory which should be created to contain the responses to interaction with the package. The name must be a full pathname. The file, or directory of files, can later be used as input to the pkgadd(1M) command.

操作数	The following operands are supported: <i>pkginst</i> Specify the package instance, or list of instances for which request scripts will be created. The token all may be used to refer to all packages available on the source medium.

退出状态	0 Successful completion. >0 An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [pkginfo\(1\)](#), [pkgmk\(1\)](#), [pkgparam\(1\)](#), [pkgproto\(1\)](#), [pkgtrans\(1\)](#), [installf\(1M\)](#), [pkgadd\(1M\)](#), [pkgchk\(1M\)](#), [pkgrm\(1M\)](#), [removef\(1M\)](#), [admin\(4\)](#), [attributes\(5\)](#)

《Application Packaging Developer's Guide》

附注

The `-r` option can be used to indicate a directory name as well as a filename. The directory name is used to create numerous response files, each sharing the name of the package with which it should be associated. This would be used, for example, when you will be adding multiple interactive packages with one invocation of [pkgadd\(1M\)](#). Each package would need a response file. To create multiple response files with the same name as the package instance, name the directory in which the files should be created and supply multiple instance names with the `pkgask` command. When installing the packages, you will be able to identify this directory to the [pkgadd\(1M\)](#) command.

If the default `admin` file is too restrictive, the administration file may need to be modified to allow for total non-interaction during a package installation. See [admin\(4\)](#) for details.

引用名 pkgchk – check package installation accuracy

用法概要

```
pkgchk [-l | -acfnqv] [-i file | -]
        [-p path... | -P partial-path...] [-R root_path]
        [-m pkgmap [-e envfile]] | pkginst... | -Y category,category...

pkgchk -d device [-l | -fv] [-i file | -] [-M] [-p path]...
        [-V fs_file]
        [pkginst... | -Y category[,category...]]
```

描述 pkgchk checks the accuracy of installed files or, by using the `-l` option, displays information about package files. pkgchk checks the integrity of directory structures and files. Discrepancies are written to standard error along with a detailed explanation of the problem.

The first synopsis defined above is used to list or check the contents and/or attributes of objects that are currently installed on the system, or in the indicated pkgmap. Package names may be listed on the command line, or by default, the entire contents of a machine will be checked.

The second synopsis is used to list or check the contents of a package which has been spooled on the specified device, but not installed. Note that attributes cannot be checked for spooled packages.

选项

The following options are supported:

`-a`

Audit the file attributes only and do not check file contents. Default is to check both.

`-c`

Audit the file contents only and do not check file attributes. Default is to check both.

`-d device`

Specify the device on which a spooled package resides. *device* can be a directory path name or the identifiers for tape or removable medium (for example, `/var/tmp`).

`-e envfile`

Request that the package information file named as *envfile* be used to resolve parameters noted in the specified pkgmap file.

`-f`

Correct file attributes if possible. If used with the `-x` option, this option removes hidden files. When pkgchk is invoked with this option, it creates directories, named pipes, links, and special devices if they do not already exist. If the `-d` option calls out an uninstalled package, the `-f` option will only take effect if the package is in directory (not stream) format. All file attributes will be set to agree with the entries in the pkgmap file except that setuid, setgid, and sticky bits will not be set in the mode.

- i *file* | -
Read a list of path names from *file* or from stdin (-) and compare this list against the installation software database or the indicated pkgmap file. Path names that are not contained in *file* or stdin are not checked.
- l
List information on the selected files that make up a package. This option is not compatible with the -a, -c, -f, -g, and -v options.
- m pkgmap
Check the package against the package map file, pkgmap.
- M
Instruct pkgchk not to use the *\$root_path/etc/vfstab* file for determining the client's mount points. This option assumes the mount points are correct on the server and it behaves consistently with Solaris 2.5 and earlier releases.
- n
Do not check volatile or editable files' contents. This should be used for most post-installation checking.
- p *path*
Check the accuracy only of the path name or path names listed. *path* can be one or more path names separated by commas (or by whitespace, if the list is quoted).

To specify a *path* that includes a comma, you must use the -i option, described above. See EXAMPLES.
- P *partial-path*
Check the accuracy of only the partial path name or path names listed. *partial-path* can be one or more partial path names separated by commas (or by whitespace, if the list is quoted). This option can be used instead of -p and is not compatible with the other option. This option matches any path name that contains the string contained in the partial path. See the note about paths that contain commas in the description of -p.
- q
Quiet mode. Do not give messages about missing files.
- R *root_path*
Define the full name of a directory to use as the *root_path*. All files, including package system information files, are relocated to a directory tree starting in the specified *root_path*. The *root_path* may be specified when installing to a client from a server (for example, /export/root/client1).

注 - The root file system of any non-global zones must not be referenced with the -R option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See [zones\(5\)](#).
- v
Verbose mode. Files are listed as processed.

-V *fs_file*

Specify an alternative *fs_file* to map the client's file systems. For example, used in situations where the *\$root_path/etc/vfstab* file is non-existent or unreliable.

-x

Search exclusive directories, looking for files which exist that are not in the installation software database or the indicated pkgmap file.

-Y *category*

Check packages based on the value of the CATEGORY parameter stored in the installed or spooled package's [pkginfo\(4\)](#) file.

操作数

pkginst

The package instance or instances to be checked. The format *pkginst.** can be used to check all instances of a package. The default is to display all information about all installed packages.

The asterisk character (*) is a special character to some shells and may need to be escaped. In the C-Shell, an asterisk must be surrounded by single quotes (') or preceded by a backslash (\);

partial-path

A portion of a path, such as a file or directory name.

示例

示例 1 Using pkgchk for Displaying Package Installation Information

The following example displays package installation information for */usr/bin/ls*:

```
example% pkgchk -l -p /usr/bin/ls
```

示例 2 Checking on Java Font Properties

The following example displays package installation information for all Java font properties installed on the system.

```
example% pkgchk -l -P font.properties
```

示例 3 Specifying a Path That Contains a Comma

Assume you want to specify the path:

```
/platform/SUNW,Netra-T12/lib
```

List this path in a file. Here is one way in which you can do that:

```
example% echo "/platform/SUNW,Netra-T12/lib" > /tmp/p
```

You can then enter:

```
example% pkgchk -i /tmp/p -l
```

```
Pathname: /platform/SUNW,Netra-T12/lib
```

```
Type: directory
```


示例 3 Specifying a Path That Contains a Comma (续)

Expected mode: 0755
Expected owner: root
Expected group: bin
Referenced by the following packages:
 system/core-osar
Current status: installed

退出状态 0
 Successful completion.

>0
 An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [pkginfo\(1\)](#), [pkgtrans\(1\)](#), [pkgadd\(1M\)](#), [pkgask\(1M\)](#), [pkgrm\(1M\)](#), [pkginfo\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

《Application Packaging Developer’s Guide》

附注 Package commands are [largefile\(5\)](#)-aware. They handle files larger than 2 GB in the same way they handle smaller files. In their current implementations, [pkgadd\(1M\)](#), [pkgtrans\(1\)](#) and other package commands can process a datastream of up to 4 GB.

引用名

pkgcond – determine type and capability of target

用法概要

```
/usr/bin/pkgcond [-nv] condition
```

描述

The pkgcond command allows you to determine the type of target being operated on (global zone, non-global zone, diskless client, and so forth) and the capabilities available for that type of client (can add a driver, path is writable, and so forth). The pkgcond command is intended to be invoked from package and patch scripts, but can also be used in situations that mimic the context of these scripts. See NOTES for further guidance.

pkgcond has one mandatory argument, a condition. The command tests whether the condition is true for the specified path. The condition can be one of the following:

- can_add_driver *[path]*
- can_remove_driver *[path]*
- can_update_driver *[path]*
- is_alternative_root *[path]*
- is_boot_environment *[path]*
- is_diskless_client *[path]*
- is_global_zone *[path]*
- is_mounted_miniroot *[path]*
- is_netinstall_image *[path]*
- is_nonglobal_zone *[path]*
- is_path_writable *path*
- is_running_system *[path]*
- is_sparse_root_nonglobal_zone *[path]*
- is_what *[path]*
- is_whole_root_nonglobal_zone *[path]*

The *path* argument usually denotes the root of the global zone or non-global zone, or alternate root. If path is optional and not specified, the default is */*.

The behavior of the *is_what* condition is somewhat special, because it displays results of all other conditions to standard output.

选项

The following options are supported:

- n
Negate return status (0 becomes 1 and 1 becomes 0). It negates results in the case of *is_what* condition.
- v
Verbose mode. Displays detailed data about intermediate checks performed.

示例

示例 1 Listing All Available Information

The following command lists all available information about the current running system, in a user-friendly way.

```
example# pkgcond -n is_what
```

示例 2 Determining if Target is an Alternate Root

The following command determines whether an alternate boot environment exists under /altroot_mount.

```
example# pkgcond is_alternative_root /altroot_mount
```

退出状态

- 0
Condition is true unless -n was specified.
- 1
Condition is false unless -n was specified.
- 2
Command line usage error.
- 3
Command failed to perform the test due to a fatal error.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Uncommitted

另请参见

[pkgtrans\(1\)](#), [pkgadd\(1M\)](#), [pkgask\(1M\)](#), [pkgchk\(1M\)](#), [pkgrm\(1M\)](#), [pkginfo\(4\)](#), [attributes\(5\)](#)

附注

Supported usage of pkgcond is subject to the following constraints:

1. Do not use pkgcond outside of the Solaris marketing release in which it is provided (for example, do not use Solaris 10 pkgcond against a Solaris 9 target).
2. Restrict use of the optional path argument according to the following rules:
 - The command `pkgcond condition $ROOTDIR` can be used in patch level scripts.
 - The command `pkgcond condition $PKG_INSTALL_ROOT` can be used in package level scripts.
 - A command of the form `pkgcond condition without` the optional path argument can be used in any context.

Use of pkgcond with an arbitrary path argument is *not* recommended or supported, as the results returned might not be accurate.

引用名 pkgrm – remove a package from the system

用法概要

```
pkgrm [-nv] [-a admin] [ [-A | -M] -R root_path]  
      [-V fs_file]  
      [pkginst... | -Y category[,category...]]  
  
pkgrm -s spool  
      [pkginst... | -Y category[,category...]]
```

描述

pkgrm will remove a previously installed or partially installed package from the system. A check is made to determine if any other packages depend on the one being removed. If a dependency exists, the action taken is defined in the admin file.

The default state for the command is in interactive mode, meaning that prompt messages are given during processing to allow the administrator to confirm the actions being taken. Non-interactive mode can be requested with the -n option.

The -s option can be used to specify the directory from which spooled packages should be removed.

Certain unbundled and third-party packages are no longer entirely compatible with the latest version of pkg rm. These packages require user interaction throughout the removal and not just at the very beginning.

To remove these older packages (released prior to Solaris 2.4), set the following environment variable:NONABI_SCRIPTS=TRUE pkg rm permits keyboard interaction throughout the removal as long as this environment variable is set.

选项

The following options are supported:

-a *admin*

Use the installation administration file, *admin*, in place of the default admin file. pkg rm first looks in the current working directory for the administration file. If the specified administration file is not in the current working directory, pkg rm looks in the /var/sadm/install/admin directory for the administration file.

-A

Remove the package files from the client's file system, absolutely. If a file is shared with other packages, the default behavior is to not remove the file from the client's file system.

-M

Instruct pkg rm not to use the *\$root_path/etc/vfstab* file for determining the client's mount points. This option assumes the mount points are correct on the server and it behaves consistently with Solaris 2.5 and earlier releases.

-n

Non-interactive mode. If there is a need for interaction, the command will exit.

Use of this option requires that at least one package instance be named upon invocation of the command. Certain conditions must exist for a package to be removed non-interactively or a non-restrictive admin file needs to be used.

-R *root_path*

Defines the full path name of a directory to use as the *root_path*. All files, including package system information files, are relocated to a directory tree starting in the specified *root_path*.

注 – The root file system of any non-global zones must not be referenced with the -R option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See [zones\(5\)](#).

-s *spool*

Remove the specified package(s) from the directory *spool*. The default directory for spooled packages is `/var/sadm/pkg`.

-v

Trace all of the scripts that get executed by pkgrm, located in the *pkginst/install* directory. This option is used for debugging the procedural and non-procedural scripts.

-V *fs_file*

Specify an alternative *fs_file* to map the client's file systems. Used in situations where the `$root_path/etc/vfstab` file is non-existent or unreliable.

-Y *category*

Remove packages based on the value of the CATEGORY parameter stored in the installed or spooled package's [pkginfo\(4\)](#) file. No package with the CATEGORY value of system can be removed from the file system with this option.

操作数

The following operand is supported:

pkginst

Specifies the package to be removed. The format *pkginst.** can be used to remove all instances of a package.

The asterisk character (*) is a special character to some shells and may need to be escaped. In the C-Shell, "*" must be surrounded by single quotes (') or preceded by a backslash (\).

示例

示例 1 Removing All Instances of SUNWjunk from client1

The following example removes all instances of SUNWjunk from client1:

```
example% pkgrm -R /export/root/client1 SUNWjunk*
```

Note the caveat on the use of the -R option in the description of that option, above.

退出状态

The following exit values are returned:

0

Successful completion.

- 1 Fatal error.
- 2 Warning.
- 3 Interruption.
- 4 Administration.
- 10 Reboot after removal of all packages.
- 20 Reboot after removal of this package.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [pkginfo\(1\)](#), [pkgmk\(1\)](#), [pkgparam\(1\)](#), [pkgproto\(1\)](#), [pkgtrans\(1\)](#), [installf\(1M\)](#), [pkgadd\(1M\)](#), [pkgask\(1M\)](#), [pkgchk\(1M\)](#), [removef\(1M\)](#), [admin\(4\)](#), [pkginfo\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

《Application Packaging Developer’s Guide》

附注 Package commands are [largefile\(5\)](#)-aware. They handle files larger than 2 GB in the same way they handle smaller files. In their current implementations, [pkgadd\(1M\)](#), [pkgtrans\(1\)](#) and other package commands can process a datastream of up to 4 GB.

引用名	plockstat – report user-level lock statistics
用法概要	<pre>plockstat [-vACHV] [-n count] [-s depth] [-e secs] [-x arg [=val]] command [arg]...</pre> <pre>plockstat [-vACHV] [-n count] [-s depth] [-e secs] [-x arg [=val]] -p pid</pre>
描述	The plockstat utility gathers and displays user-level locking statistics. By default, plockstat monitors all lock contention events, gathers frequency and timing data about those events, and displays the data in decreasing frequency order, so that the most common events appear first. plockstat gathers data until the specified command completes or the process specified with the -p option completes. plockstat relies on DTrace to instrument a running process or a command it invokes to trace events of interest. This imposes a small but measurable performance overhead on the processes being observed. Users must have the <code>dttrace_proc</code> privilege and have permission to observe a particular process with plockstat. Refer to the 《Solaris Dynamic Tracing Guide》 for more information about DTrace security features.
选项	The following options are supported: -A Watch all lock events. This option is equivalent to -CH. -C Watch contention events. -H Watch hold events. -e secs Exit after the number of seconds specified have elapsed. -n count Display only the specified number of entries for each output category. -s depth Record a stack trace rather than just the calling function. -p pid Specify a process ID from which plockstat is to gather data. -v Print out a message to indicate that tracing has started. -x arg[=val] Enable or modify a DTrace runtime option or D compiler option. The list of options is found in the 《Solaris Dynamic Tracing Guide》. Boolean options are enabled by specifying their name. Options with values are set by separating the option name and value with an equals sign (=). -V Print the Dtrace commands used to gather the data. The output can then be used directly with the dttrace(1M) command.
操作数	The following operands are supported: <i>arg</i> A string to be passed as an argument to <i>command</i>. <i>command</i> The name of a utility to be invoked.

- count* A positive integer value.
- pid* A process identifier for a process to be monitored.
- secs* Duration specified as a positive integer number of seconds.

Display Headers The following headers appear over columns of data in plockstat output.

- Count
 Number of times an event occurred.
- nsec
 Average duration of an event, in nanoseconds.
- Lock
 Address of a lock, displayed symbolically if possible.
- Caller
 Address of a caller, displayed symbolically if possible.

退出状态 The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/dtrace
Interface Stability	See below.

The command-line syntax is Committed. The human-readable output is Uncommitted.

另请参见 [dtrace\(1M\)](#), [lockstat\(1M\)](#), [mutex_init\(3C\)](#), [pthread_mutex_lock\(3C\)](#),
[pthread_rwlock_rdlock\(3C\)](#), [pthread_rwlock_wrlock\(3C\)](#), [pthread_rwlock_unlock\(3C\)](#),
[rwlock\(3C\)](#), [attributes\(5\)](#), [fasttrap\(7D\)](#)

《Solaris Dynamic Tracing Guide》

引用名

pntadm – DHCP network table management utility

用法概要

```
pntadm -C [-r resource] [-p path] [-u uninterpreted] network

pntadm -A name_IP_address [-c comment] [-e mm/dd/yyyy]
    [-f num | keywords] [-h client_hostname]
    [-i [-a] client_ID] [-m [-y] macro] [-s server]
    [-r resource] [-p path] [-u uninterpreted] network

pntadm -M name_IP_address [-c comment] [-e mm/dd/yyyy]
    [-f num | keywords] [-h client_hostname]
    [-i [-a] client ID] [-m [-y] macro]
    [-n new_client_IP_address] [-s server] [-r resource]
    [-p path] [-u uninterpreted] network

pntadm -D name_IP_address [-y] [-r resource] [-p path]
    [-u uninterpreted] network

pntadm -P [-v] [-x] [-r resource] [-p path]
    [-u uninterpreted] network

pntadm -R [-r resource] [-p path] [-u uninterpreted] network

pntadm -L [-r resource] [-p path] [-u uninterpreted]

pntadm -B [-v] [batchfile]
```

描述

The pntadm command is used to manage the Dynamic Host Configuration Protocol (DHCP) network tables. It is used to add and remove networks under DHCP management, and add, delete, or modify IP address records within network tables, or to view tables. For a description of the format of DHCP network tables, see [dhcp_network\(4\)](#).

pntadm can be run as root or by other users assigned to the DHCP Management profile. See [rbac\(5\)](#) and [user_attr\(4\)](#).

If the networks you want to add are subnetted, you need to update the [netmasks\(4\)](#) table.

One of the following options (function flags) must be specified with the pntadm command: -A, -B, -C, -D, -L, -M, -P, or -R.

The pntadm utility is obsolete and is subject to removal in a future release of Oracle Solaris.

选项

The following options are supported:

-A name_IP_address

Add a client entry with hostname or client IP address, *name_IP_address*, to the named DHCP network table.

The following sub-options are optional:

-c comment

Comment text. The default is NULL.

-e *mm/dd/yyyy*

Absolute lease. The default is 0.

-f *num* | *keywords*

Flag value. The default is 00.

The flag (-f) option can be specified either as a single number denoting the intended flag value, or as a series of the following keywords, combined using the plus (+) symbol:

DYNAMIC or 00

Server manager's assignment.

PERMANENT or 01

Lease on entry is permanent.

MANUAL or 02

Administrator managed assignment.

UNUSABLE or 04

Entry is not valid.

BOOTP or 08

Entry reserved for BOOTP clients.

For a more detailed description of the flag values, see [dhcp_network\(4\)](#).

-h *client_hostname*

Client hostname. The default is NULL.

When the -h option is used in this mode, the *client_hostname* is added to the hosts table within the resource used for storing host names (files or DNS). The command will fail if this *client_hostname* is already present in the hosts table.

-i *client_ID* [-a]

Client identifier [-a]. The default is 00.

The -i option modified with -a specifies that the client identifier is in ASCII format, and thus needs to be converted to hexadecimal format before insertion into the table.

-m *macro* [-y]

Macro name. Default is UNKNOWN.

The -m option modified with -y verifies the existence of the named macro in the `dhcptab` table before adding the entry.

-s *server*

Server IP or name. Default is system name (uname -n).

-B

Activate batch mode. pntadm will read from the specified file or from standard input a series of pntadm commands and execute them within the same process. Processing many

pntadm commands using this method is much faster than running an executable batchfile itself. Batch mode is recommended for using pntadm in scripts.

The following sub-option is optional:

-v
Display commands to standard output as they are processed.

-C
Create the DHCP network table for the network specified by *network*. See [操作数](#). For details, see [dhcp_network\(4\)](#) and [networks\(4\)](#).

-D *name_IP_address*
Delete the specified client entry with hostname or client IP address, *name_IP_address*, in the named DHCP network table. (See [dhcp_network\(4\)](#).)

The following sub-option is optional:

-y
Remove associated host table entry. The **-y** option requests that all hostnames associated with the IP address in the hosts table in the resource be removed.

-L
List the DHCP network tables presently configured, one per line, on standard output. If none are found, no output is printed and an exit status of 0 is returned.

-M *name_IP_address*
Modify the specified client entry with hostname or client IP address, *name_IP_address*, in the named DHCP network table. See [dhcp_network\(4\)](#). The default for the sub-options is what they currently are set to.

The following sub-options are optional.

-c *comment*
New comment text.

-e *mm/dd/yy*
New absolute lease expiration date. Time defaults to 12:00 AM of the day specified.

-f *num* | *keyboard*
New flag value, see explanation following the description of the **-A** option.

-h *host_name*
New client hostname.

The **-h** option allows you to change the current *hostname* associated with the IP address or to add a new *hostname* to the hosts table if an entry associated with this IP address does not exist.

-i *client_ID*
New client identifier [-a].

- m *macro* [-y]
Macro name defined in dhcptab.
- n *new_client_IP_address*
New IP address.
- s *server*
New server IP or name.

For more detailed description of the sub-options and flag values, see [dhcp_network\(4\)](#).

- P
Display the named DHCP network table.
The following sub-options are optional:

- v
Display lease time in full verbose format and resolve IP addresses for the clients and server to hostnames.
- x
Display lease time in raw format.

These flag codes are used with the -P sub-options:

-v	-x	Description
D	00	DYNAMIC
P	01	PERMANENT
M	02	MANUAL
U	04	UNUSABLE
B	08	BOOTP

See [dhcp_network\(4\)](#) for information on these sub-options and associated flag codes.

- p *path*
Override the [dhcpsvc.conf\(4\)](#) configuration value for data store resource path, *path* See [dhcpsvc.conf\(4\)](#)
- R
Remove the named DHCP network table. See [dhcp_network\(4\)](#).
- r *data_store_resource*
Override the `/etc/inet/dhcpsvc.conf` configuration value for RESOURCE= with the *data_store_resource* specified. See the [dhcpsvc.conf\(4\)](#) man page for more details on resource type, and the 《[Oracle Solaris DHCP Service Developer's Guide](#)》 for more information about adding support for other data stores.

-u uninterpreted

Data which will be ignored by pntadm, but passed to the currently configured public module to be interpreted by the data store. This might be used for a database account name or other authentication or authorization parameters required by a particular data store.

操作数

The following operand is supported:

network

The network address or network name which corresponds to the dhcp network table. See [dhcp_network\(4\)](#).

示例

示例 1 Creating a Table for the 10.0.0.0 DHCP Network

The following command creates a table for the 10.0.0.0 (subnetted to class C) DHCP network table. Note that if you have an alias for this network in your [networks\(4\)](#) table, you can use that value rather than the dotted Internet Address notation.

```
example# pntadm -C 10.0.0.0
```

示例 2 Adding an Entry to the 10.0.0.0 Table

The following command adds an entry to the 10.0.0.0 table in the files resource in the /var/mydhcp directory:

```
example# pntadm -r SUNWfiles -p /var/mydhcp -A 10.0.0.1 10.0.0.0
```

示例 3 Modifying the 10.0.0.1 Entry of the 10.0.0.0 Table

The following command modifies the 10.0.0.1 entry of the 10.0.0.0 table, changing the macro name to Green, setting the flags field to MANUAL and PERMANENT:

```
example# pntadm -M 10.0.0.1 -m Green -f 'PERMANENT+MANUAL' 10.0.0.0
```

示例 4 Changing the 10.0.0.1 Entry to 10.0.0.2

The following command changes the 10.0.0.1 entry to 10.0.0.2, making an entry in the [hosts\(4\)](#) table called myclient:

```
example# pntadm -M 10.0.0.1 -n 10.0.0.2 -h myclient 10.0.0.0
```

示例 5 Setting the Client ID as ASCII

The following command sets the client ID as ASCII aruba.foo.com for the myclient entry:

```
example# pntadm -M myclient -i 'aruba.foo.com' -a 10.0.0.0
```

示例 6 Deleting the myclientEntry from the 10.0.0.0 Table

The following command deletes the myclient (10.0.0.2) entry from the 10.0.0.0 table:

```
example# pntadm -D myclient 10.0.0.0
```

示例 7 Listing the Configured DHCP Network Tables

The following command lists the configured DHCP network tables:

```
example# pntadm -L
192.168.0.0
10.0.0.0
```

示例 8 Executing pntadm Commands in Batch Mode

The following command runs a series of pntadm commands contained in a batch file:

```
example# pntadm -B addclients
```

退出状态

- 0 Successful completion.
- 1 Object already exists.
- 2 Object does not exist.
- 3 Non-critical error.
- 4 Critical error.

文件

- /etc/inet/dhcpsvc.conf
- /etc/inet/hosts

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/dhcp
Interface Stability	Obsolete

另请参见

[dhcpconfig\(1M\)](#), [dhcpcmgr\(1M\)](#), [dhcp_network\(4\)](#), [dhcpsvc.conf\(4\)](#), [dhcptab\(4\)](#), [hosts\(4\)](#), [netmasks\(4\)](#), [networks\(4\)](#), [user_attr\(4\)](#), [attributes\(5\)](#), [dhcp\(5\)](#), [dhcp_modules\(5\)](#), [rbac\(5\)](#)

《Oracle Solaris DHCP Service Developer’s Guide》

《Oracle Solaris 管理：IP 服务》

Alexander, S., and R. Droms, DHCP Options and BOOTP Vendor Extensions, RFC 1533, Lachman Technology, Inc., Bucknell University, October 1993.

Droms, R., Interoperation Between DHCP and BOOTP, RFC 1534, Bucknell University, October 1993.

Droms, R., Dynamic Host Configuration Protocol, RFC 1541, Bucknell University, October 1993.

Wimer, W., Clarifications and Extensions for the Bootstrap Protocol, RFC 1542, Carnegie Mellon University, October 1993.

引用名 polkit-is-privileged – 检查 PolicyKit 特权

用法概要 polkit-is-privileged [-hvV] -u *user* -p *privilege* [-r *resource*]

描述 polkit-is-privileged 命令可查询系统策略，以确定是否允许用户具有给定的特权和使用给定的资源。资源名称可以忽略。在 Solaris 操作系统上，应使用 RBAC 授权名称作为特权名称。

当前，PolicyKit 唯一的使用者是 [hald\(1M\)](#)。

选项 所支持的选项如下：

-h, --help	显示选项列表并退出。
-p <i>privilege</i> , --privilege <i>privilege</i>	与用户关联的特权的名称。命令将测试此特权。
-r <i>resource</i> , --resource <i>resource</i>	与用户和特权关联的资源名称。命令将测试此资源。
-u <i>user</i> , --user <i>user</i>	要测试的用户名和用户 ID。
-v, --verbose	详细模式。
-V, --version	显示版本号。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/library/policykit
接口稳定性	Volatile（可变）

另请参见 [auths\(1\)](#)、[profiles\(1\)](#)、[hald\(1M\)](#)、[getauthattr\(3C\)](#)、[auth_attr\(4\)](#)、[policy.conf\(4\)](#)、[prof_attr\(4\)](#)

引用名	pooladm – 激活和取消激活资源池工具
用法概要	<pre> /usr/sbin/pooladm [-n] [-s] [-c] [filename] -x /usr/sbin/pooladm [-d -e] </pre>
描述	pooladm 命令提供有关池和集合的管理操作。pooladm 读取指定文件名，并尝试激活其中包含的池配置。 在更新当前的池运行时配置之前，pooladm 会验证配置的正确性。 不带选项的 pooladm 将打印当前正在运行的池配置。
选项	所支持的选项如下： -c 对给定位置的配置进行实例化。如果未指定文件名，将缺省使用 /etc/pooladm.conf。 -d 禁用池工具，以便不能再对池进行处理。 -e 启用池功能以对池进行处理。 -n 验证配置而实际上不更新当前的活动配置。确认没有语法错误，并且可以在当前系统上将配置实例化。不执行应用程序特定属性的验证。 -s 使用当前动态配置的详细信息更新指定位置。 此选项要求对要更新的配置具有更新权限。如果将此选项与 -c 选项结合使用，动态配置将在静态位置之前更新。 -x 删除当前活动的池配置。销毁所有定义的资源，并将之前已分区的所有组件返回到各自的缺省资源。
操作数	支持下列操作数： <i>filename</i> 使用此文件中包含的配置。
示例	示例1 实例化配置 以下命令将实例化 /home/admin/newconfig 中包含的配置： <pre>example# /usr/sbin/pooladm -c /home/admin/newconfig</pre> 示例2 验证配置而不对其实例化 以下命令将尝试实例化 /home/admin/newconfig 中包含的配置。命令将显示遇到的任何错误情况，但不会实际修改活动的配置。 <pre>example# /usr/sbin/pooladm -n -c /home/admin/newconfig</pre>

引用名	poolbind – bind processes, tasks, or projects or query binding of processes to resource pools																
用法概要	<pre>/usr/sbin/poolbind -p poolname -e command [arguments]... /usr/sbin/poolbind -p poolname [-i idtype] id... /usr/sbin/poolbind -q pid... /usr/sbin/poolbind -Q pid...</pre>																
描述	The poolbind command allows an authorized user to bind zones, projects, tasks, and processes to pools. With the -e option (see below), it can execute a command you specify, placing the executed command in a specified pool. It can also enable you to query a process to determine which pool a process is bound to.																
选项	The following options are supported: <table><tr><td>-e <i>command</i> [<i>arguments</i>...]</td><td>Executes <i>command</i>, bound to the pool you specify with -p.</td></tr><tr><td>-i <i>idtype</i></td><td>This option, together with the <i>idlist</i> arguments, specifies one or more processes to which the poolbind command is to apply. The interpretation of <i>idlist</i> depends on the value of <i>idtype</i>. The valid <i>idtype</i> arguments and corresponding interpretations of <i>idlist</i> are as follows: <table><tr><td><i>pid</i></td><td><i>idlist</i> is a list of process IDs. Binds the specified processes to the specified pool. This is the default behavior if no <i>idtype</i> is specified.</td></tr><tr><td><i>taskid</i></td><td><i>idlist</i> is a list of task IDs. Bind all processes within the list of task IDs to the specified pool.</td></tr><tr><td><i>projid</i></td><td><i>idlist</i> is a list of project IDs. Bind all processes within the list of projects to the specified pool. Each project ID can be specified as either a project name or a numerical project ID. See project(4).</td></tr><tr><td><i>zoneid</i></td><td><i>idlist</i> is a list of zone IDs. Bind all processes within the list of zones to the specified pool. Each zone ID can be specified as either a zone name or a numerical zone ID. See zones(5).</td></tr></table></td></tr><tr><td>-p <i>poolname</i></td><td>Specifies the name of a pool to which the specified zone, project, tasks, or processes are to be bound.</td></tr><tr><td>-q <i>pid</i> ...</td><td>Queries the pool bindings for a given list of process IDs. If the collection of resources associated with the process does not correspond to any currently existing pool, or if there are multiple pools with the set of resources that the process is bound to, the query fails for that particular process ID.</td></tr></table>	-e <i>command</i> [<i>arguments</i> ...]	Executes <i>command</i> , bound to the pool you specify with -p.	-i <i>idtype</i>	This option, together with the <i>idlist</i> arguments, specifies one or more processes to which the poolbind command is to apply. The interpretation of <i>idlist</i> depends on the value of <i>idtype</i> . The valid <i>idtype</i> arguments and corresponding interpretations of <i>idlist</i> are as follows: <table><tr><td><i>pid</i></td><td><i>idlist</i> is a list of process IDs. Binds the specified processes to the specified pool. This is the default behavior if no <i>idtype</i> is specified.</td></tr><tr><td><i>taskid</i></td><td><i>idlist</i> is a list of task IDs. Bind all processes within the list of task IDs to the specified pool.</td></tr><tr><td><i>projid</i></td><td><i>idlist</i> is a list of project IDs. Bind all processes within the list of projects to the specified pool. Each project ID can be specified as either a project name or a numerical project ID. See project(4).</td></tr><tr><td><i>zoneid</i></td><td><i>idlist</i> is a list of zone IDs. Bind all processes within the list of zones to the specified pool. Each zone ID can be specified as either a zone name or a numerical zone ID. See zones(5).</td></tr></table>	<i>pid</i>	<i>idlist</i> is a list of process IDs. Binds the specified processes to the specified pool. This is the default behavior if no <i>idtype</i> is specified.	<i>taskid</i>	<i>idlist</i> is a list of task IDs. Bind all processes within the list of task IDs to the specified pool.	<i>projid</i>	<i>idlist</i> is a list of project IDs. Bind all processes within the list of projects to the specified pool. Each project ID can be specified as either a project name or a numerical project ID. See project(4) .	<i>zoneid</i>	<i>idlist</i> is a list of zone IDs. Bind all processes within the list of zones to the specified pool. Each zone ID can be specified as either a zone name or a numerical zone ID. See zones(5) .	-p <i>poolname</i>	Specifies the name of a pool to which the specified zone, project, tasks, or processes are to be bound.	-q <i>pid</i> ...	Queries the pool bindings for a given list of process IDs. If the collection of resources associated with the process does not correspond to any currently existing pool, or if there are multiple pools with the set of resources that the process is bound to, the query fails for that particular process ID.
-e <i>command</i> [<i>arguments</i> ...]	Executes <i>command</i> , bound to the pool you specify with -p.																
-i <i>idtype</i>	This option, together with the <i>idlist</i> arguments, specifies one or more processes to which the poolbind command is to apply. The interpretation of <i>idlist</i> depends on the value of <i>idtype</i> . The valid <i>idtype</i> arguments and corresponding interpretations of <i>idlist</i> are as follows: <table><tr><td><i>pid</i></td><td><i>idlist</i> is a list of process IDs. Binds the specified processes to the specified pool. This is the default behavior if no <i>idtype</i> is specified.</td></tr><tr><td><i>taskid</i></td><td><i>idlist</i> is a list of task IDs. Bind all processes within the list of task IDs to the specified pool.</td></tr><tr><td><i>projid</i></td><td><i>idlist</i> is a list of project IDs. Bind all processes within the list of projects to the specified pool. Each project ID can be specified as either a project name or a numerical project ID. See project(4).</td></tr><tr><td><i>zoneid</i></td><td><i>idlist</i> is a list of zone IDs. Bind all processes within the list of zones to the specified pool. Each zone ID can be specified as either a zone name or a numerical zone ID. See zones(5).</td></tr></table>	<i>pid</i>	<i>idlist</i> is a list of process IDs. Binds the specified processes to the specified pool. This is the default behavior if no <i>idtype</i> is specified.	<i>taskid</i>	<i>idlist</i> is a list of task IDs. Bind all processes within the list of task IDs to the specified pool.	<i>projid</i>	<i>idlist</i> is a list of project IDs. Bind all processes within the list of projects to the specified pool. Each project ID can be specified as either a project name or a numerical project ID. See project(4) .	<i>zoneid</i>	<i>idlist</i> is a list of zone IDs. Bind all processes within the list of zones to the specified pool. Each zone ID can be specified as either a zone name or a numerical zone ID. See zones(5) .								
<i>pid</i>	<i>idlist</i> is a list of process IDs. Binds the specified processes to the specified pool. This is the default behavior if no <i>idtype</i> is specified.																
<i>taskid</i>	<i>idlist</i> is a list of task IDs. Bind all processes within the list of task IDs to the specified pool.																
<i>projid</i>	<i>idlist</i> is a list of project IDs. Bind all processes within the list of projects to the specified pool. Each project ID can be specified as either a project name or a numerical project ID. See project(4) .																
<i>zoneid</i>	<i>idlist</i> is a list of zone IDs. Bind all processes within the list of zones to the specified pool. Each zone ID can be specified as either a zone name or a numerical zone ID. See zones(5) .																
-p <i>poolname</i>	Specifies the name of a pool to which the specified zone, project, tasks, or processes are to be bound.																
-q <i>pid</i> ...	Queries the pool bindings for a given list of process IDs. If the collection of resources associated with the process does not correspond to any currently existing pool, or if there are multiple pools with the set of resources that the process is bound to, the query fails for that particular process ID.																

-Q pid ...

Queries the resource bindings for a given list of process IDs.
The resource bindings are each reported on a separate line.

示例

示例 1 Binding All Processes

The following command binds all processes in projects 5 and 7 to the pool web_app:

example# /usr/sbin/poolbind -p web_app -i projid 5 7

示例 2 Binding the Running Shell

The following command binds the running shell to the pool web_app:

example# /usr/sbin/poolbind -p web_app \$\$

示例 3 Querying the Pool Bindings

The following command queries the bindings to verify that the shell is bound to the given pool:

example# /usr/sbin/poolbind -q \$\$

示例 4 Querying the Resource Bindings

The following command queries the bindings to verify that the shell is bound to the given resources:

example# /usr/sbin/poolbind -Q \$\$

退出状态

The following exit values are returned:

0 Successful completion.

1 Requested operation could not be completed.

2 Invalid command line options were specified.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/resource-mgmt/resource-pools
Interface Stability	See below.

The invocation is Committed. The output is Uncommitted.

另请参见

[pooladm\(1M\)](#), [poolcfg\(1M\)](#), [libpool\(3LIB\)](#), [project\(4\)](#), [attributes\(5\)](#), [zones\(5\)](#)
《Oracle Solaris 管理：Oracle Solaris Zones、Oracle Solaris 10 Zones 和资源管理》

引用名	poolcfg – 创建和修改资源池配置文件
用法概要	<pre> /usr/sbin/poolcfg -c <i>command</i> [-d [<i>filename</i>]] /usr/sbin/poolcfg -f <i>command_file</i> [-d [<i>filename</i>]] /usr/sbin/poolcfg -h </pre>
描述	poolcfg 实用程序提供有关池和集合的配置操作。这些操作针对现有配置执行，并采取修改指定配置文件的形式。如果使用 -d 选项，则对内核状态进行修改。对所获得的配置的实际激活操作通过 pooladm(1M) 实用程序完成。 池配置文件必须是使用 poolcfg 本身或 libpool(3LIB) 直接构建的结构化文件。 如果动态位置已开放用于写入，则使用池动态位置和写入权限的 poolcfg 调用将挂起。 pooladm 可以使用此实用程序创建的配置在目标主机上对配置进行实例化。
选项	所支持的选项如下： -c <i>command</i> 将 <i>command</i> 指定为编辑命令。请参见 USAGE。 -d 直接对内核状态进行操作。不允许使用 <i>filename</i>。 -f <i>command_file</i> 从 <i>command_file</i> 中获取命令。<i>command_file</i> 由编辑命令组成，每行一个命令。 -h 显示有关编辑命令语法的扩展信息。
用法	
脚本	脚本包含以下编辑命令，每行一个命令： <pre> info [<i>entity-name</i>] 以用户可读的形式将配置（或指定部分）显示至标准输出。如果没有指定实体，将显示系统信息。因此，poolcfg -c 'info' afile 是等同于 poolcfg -c 'info system name' afile 的调用。 create <i>entity-name</i> [<i>property-list</i>] 创建具有指定类型和名称的实体。 destroy <i>entity-name</i> 删除指定的实体。 modify <i>entity-name</i> [<i>property-list</i>] 更改指定实体上列出的属性。 associate <i>pool-name</i> [<i>resource-list</i>] 将一个或多个资源连接到池，或者替换一个或多个现有连接。 transfer to [<i>resourcetype</i>] <i>name</i>[<i>component-list</i>] 将一个或多个离散组件传输到资源。 </pre>

`transfer [quantity] from [resourcetype] [src] to [tgt]`
将资源数量从 *src* 传输到 *tgt*。

`transfer [quantity] to [resourcetype] [tgt] from [src]`
将资源数量从 *tgt* 传输到 *src*。

discover
创建系统实体，其中一个池实体和资源与当前系统配置匹配。所发现的每个资源类型的所有资源均记录在文件中，其中单个池引用每个资源类型的缺省资源。

当 `poolcfg` 直接对内核进行操作时，此命令是一个 NO-OP。请参见 `-d` 选项。

应避免使用此命令。创建配置的首选方法是配合使用 `pooladm(1M)` 和 `-s` 选项来导出动态配置。

rename *entity-name* to *new-name*
将系统上实体的名称更改为新名称。

属性列表 属性列表按以下方式指定：

`( proptype name = value [ ; proptype name = value ]* )`

其中对于给定的属性类型 (*proptype*) 和名称 (*name*) 对，将采用序列中最后一个定义。要删除属性，请使用 `~ proptype name`。

资源列表 资源列表按以下方式指定：

`( resourcetype name [ ; resourcetype name ]* )`

其中对于某个资源，将采用序列中最后一个指定资源。没有适用于资源列表的删除语法。

组件列表 组件列表按以下方式指定：

`( componenttype name [ ; componenttype name ]* )`

其中对于某个组件，将采用序列中最后一个指定组件。没有适用于组件列表的删除语法。

识别的实体 `system` 计算机级别实体

`pool` 资源关联的命名集合

资源类型 `pset` 处理器集资源

属性类型 `boolean` 值为 `true` 或 `false` 之一。

`int` 64 位带符号的整数值。

`uint` 64 位不带符号的整数值。

`string` 字符串由引号 (") 分隔，并且支持以 `formats(5)` 定义的字符转义序列。

`float` 不支持科学记数法。

示例

示例1 编写 poolcfg 脚本

以下 poolcfg 脚本将创建一个名为 Accounting 的池和一个处理器集 small-1。首先创建处理器集，然后创建池，再将其与处理器集关联。

```
create pset small-1 ( uint pset.min = 1 ; uint pset.max = 4)
create pool Accounting
associate pool Accounting ( pset small-1 )
```

示例2 报告 pool_0

以下命令以用户可读形式在标准输出中提供有关 pool_0 的报告：

```
# poolcfg -c 'info pool pool_0' /etc/pooladm.conf
```

示例3 销毁 pool_0 及其关联

以下命令将销毁 pool_0 及关联，但不会销毁之前已关联的资源：

```
# poolcfg -c 'destroy pool pool_0' /etc/pooladm.conf
```

示例4 现实当前配置

以下命令将显示当前配置：

```
$ poolcfg -c 'info' /etc/pooladm.conf
system example_system
    int system.version 1
    boolean system.bind-default true
    string system.comment Discovered by libpool

    pool pool_default
        boolean pool.default true
        boolean pool.active true
        int pool.importance 5
        string pool.comment
        string pool.scheduler FSS
        pset pset_default

    pset pset_default
        int pset.sys_id -1
        string pset.units population
        boolean pset.default true
        uint pset.max 4294967295
        uint pset.min 1
        string pset.comment
        boolean pset.escapable false
        uint pset.load 0
        uint pset.size 2

    cpu
```

示例 4 现实当前配置 (续)

```
int cpu.sys_id 0
string cpu.comment

cpu
int cpu.sys_id 2
string cpu.comment
```

示例 5 将 ID 为 2 的 cpu 移动到内核中的处理器集 pset1
以下命令将 ID 为 2 的 cpu 移动到内核中的处理器集 pset1：
poolcfg -dc 'transfer to pset pset1 (cpu 2)'

示例 6 将 2 个 cpu 从内核中的处理器集 pset1 移动到处理器集 pset2 中
以下命令将 2 个 cpu 从内核中的处理器集 pset1 移动到处理器集 pset2 中：
poolcfg -dc 'transfer 2 from pset pset1 to pset2'

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/resource-mgmt/resource-pools
接口稳定性	请参见下文。

调用为 Committed（已确定）。输出为 Uncommitted（未确定）。

另请参见 [pooladm\(1M\)](#)、[poolbind\(1M\)](#)、[libpool\(3LIB\)](#)、[attributes\(5\)](#)、[formats\(5\)](#)
《Oracle Solaris 管理：Oracle Solaris Zones、Oracle Solaris 10 Zones 和资源管理》

引用名	poold – automated resource pools partitioning daemon												
用法概要	poold [-l level]												
描述	poold provides automated resource partitioning facilities. poold can be enabled or disabled using the Solaris Service Management Facility, smf(5). poold requires the Resource Pools facility to be active in order to operate. The dynamic resource pools service's fault management resource identifier (FMRI) is: <pre>svc:/system/pools/dynamic</pre> The resource pools service's FMRI is: <pre>svc:/system/pools</pre> poold's configuration details are held in a libpool(3LIB) configuration and you can access all customizable behavior from this configuration. poold periodically examines the load on the system and decides whether intervention is required to maintain optimal system performance with respect to resource consumption. poold also responds to externally initiated (with respect to poold) changes of either resource configuration or objectives. If intervention is required, poold attempts to reallocate the available resources to ensure that performance objectives are satisfied. If it is not possible for poold to meet performance objectives with the available resources, then a message is written to the log. poold allocates scarce resources according to the objectives configured by the administrator. The system administrator must determine which resource pools are most deserving of scarce resource and indicate this through the importance of resource pools and objectives.												
选项	The following options are supported: -l level Specify the verbosity level for logging information. Specify <i>level</i> as ALERT, CRIT, ERR, WARNING, NOTICE, INFO, and DEBUG. If <i>level</i> is not supplied, then the default logging level is INFO. <table><tr><td>ALERT</td><td>A condition that should be corrected immediately, such as a corrupted system database.</td></tr><tr><td>CRIT</td><td>Critical conditions, such as hard device errors.</td></tr><tr><td>ERR</td><td>Errors.</td></tr><tr><td>WARNING</td><td>Warning messages.</td></tr><tr><td>NOTICE</td><td>Conditions that are not error conditions, but that may require special handling.</td></tr><tr><td>INFO</td><td>Informational messages.</td></tr></table>	ALERT	A condition that should be corrected immediately, such as a corrupted system database.	CRIT	Critical conditions, such as hard device errors.	ERR	Errors.	WARNING	Warning messages.	NOTICE	Conditions that are not error conditions, but that may require special handling.	INFO	Informational messages.
ALERT	A condition that should be corrected immediately, such as a corrupted system database.												
CRIT	Critical conditions, such as hard device errors.												
ERR	Errors.												
WARNING	Warning messages.												
NOTICE	Conditions that are not error conditions, but that may require special handling.												
INFO	Informational messages.												

DEBUG	Messages that contain information normally of use only when debugging a program.

When invoked manually, with the `-l` option, all log output is directed to standard error.

示例

示例 1 Modifying the Default Logging Level

The following command modifies the default logging level to ERR:

```
# /usr/lib/pool/poold -l ERR
```

示例 2 Enabling Dynamic Resource Pools

The following command enables dynamic resource pools:

```
# /usr/sbin/svcadm enable svc:/system/pools/dynamic
```

属性

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/resource-mgmt/resource-pools
Interface Stability	See below.

The invocation is Committed. The output is Uncommitted.

另请参见

```
pooladm(1M), poolbind(1M), poolcfg(1M), poolstat(1M), svcadm(1M),  
pool_set_status(3POOL), libpool(3LIB), attributes(5), smf(5)
```

《Oracle Solaris 管理：Oracle Solaris Zones、Oracle Solaris 10 Zones 和资源管理》

引用名	poolstat – report active pool statistics																				
用法概要	<pre>poolstat [-p pool-list] [-r rset-list] [-T u d] [interval [count]] poolstat [-p pool-list] [-o format -r rset-list] [-T u d] [interval [count]]</pre>																				
描述	The poolstat utility iteratively examines all active pools on the system. It reports statistics based on the selected output mode. poolstat provides options to examine only specified pools and report resource set-specific statistics. Without options, poolstat examines all pools, reports basic statistics for their resource sets, and exits.																				
DISPLAY FORMATS	In default output format, poolstat outputs a header line and a line for each pool. The line begins with the pool ID and its name, followed by a column of statistical data for the processor set attached to the pool. The columns are defined as follows: <table><tr><td>id</td><td>Pool ID.</td></tr><tr><td>pool</td><td>Pool name.</td></tr><tr><td>rid</td><td>Resource set id.</td></tr><tr><td>rset</td><td>Resource set name.</td></tr><tr><td>type</td><td>Resource set type.</td></tr><tr><td>min</td><td>Minimum resource set size.</td></tr><tr><td>max</td><td>Maximum resource set size.</td></tr><tr><td>size</td><td>Current resource set size.</td></tr><tr><td>used</td><td>The measure of how much of the resource set is currently is use. This is calculated as the percentage utilization of the resource set multiplied by its size. If resource set has been reconfigured during last sampling interval, this value might be not reported (-).</td></tr><tr><td>load</td><td>The absolute representation of the load that is put on the resource set. For the definition of this property see libpool(3LIB).</td></tr></table>	id	Pool ID.	pool	Pool name.	rid	Resource set id.	rset	Resource set name.	type	Resource set type.	min	Minimum resource set size.	max	Maximum resource set size.	size	Current resource set size.	used	The measure of how much of the resource set is currently is use. This is calculated as the percentage utilization of the resource set multiplied by its size. If resource set has been reconfigured during last sampling interval, this value might be not reported (-).	load	The absolute representation of the load that is put on the resource set. For the definition of this property see libpool(3LIB) .
id	Pool ID.																				
pool	Pool name.																				
rid	Resource set id.																				
rset	Resource set name.																				
type	Resource set type.																				
min	Minimum resource set size.																				
max	Maximum resource set size.																				
size	Current resource set size.																				
used	The measure of how much of the resource set is currently is use. This is calculated as the percentage utilization of the resource set multiplied by its size. If resource set has been reconfigured during last sampling interval, this value might be not reported (-).																				
load	The absolute representation of the load that is put on the resource set. For the definition of this property see libpool(3LIB) .																				
选项	The following options are supported: <table><tr><td>-o format</td><td>Report statistics according to the format specification given in format. See DISPLAY FORMATS.</td></tr></table> The -o option accepts lists as arguments. Items in a list can be either separated by commas or enclosed in quotes and separated by commas or spaces. You can specify multiple -o options. The format specification is interpreted as the whitespace separated concatenation of all the format option arguments.	-o format	Report statistics according to the format specification given in format. See DISPLAY FORMATS.																		
-o format	Report statistics according to the format specification given in format. See DISPLAY FORMATS.																				

The -o option must be used in conjunction with the -r option.

-p *pool-list* Report only pools whose names are in the given list. If the -r option is also used, this option selects only resource sets which belong to pools in the given list. Statistics for pools or resource sets are reported in the same order in which pool names are listed on the pool-list. Pool can be specified by name or by ID.

The -p option accepts lists as arguments. Items in a *pool-list* can only be separated by spaces.

-r *rset-list* Report resource set statistics. If the *rset-list* argument is “all”, then all possible resource set types are selected.

The -r option accepts lists as arguments. Items in a list can be either separated by commas or enclosed in quotes and separated by commas or spaces.

The following resource set types are supported:

all All resource set types

pset Processor set

-T u | d Display a time stamp.

Specify u for a printed representation of the internal representation of time. See [time\(2\)](#). Specify d for standard date format. See [date\(1\)](#).

操作数

The following operands are supported:

count The number of times that the statistics are repeated. By default, poolstat reports statistics only once.

If neither interval nor count are specified, statistics are reported once. If interval is specified and count is not, statistics are reported indefinitely.

interval The sampling interval in seconds.

If neither interval nor count are specified, statistics are reported once. If interval is specified and count is not, statistics are reported indefinitely.

示例

示例 1 Using poolstat

The following example shows the default output from the poolstat utility:

```
% poolstat
           pset
id pool      size used load
0 pool_default 4  3.6  6.2
1 pool_admin   4  3.3  8.4
```

示例2 Reporting Resource Set Statistics

The following example reports resource set statistics.

```
% poolstat -r pset
  id pool          type rid rset          min  max size used load
  0 pool_default   pset  -1 pset_default   1  65K   2  1.2  8.3
  1 pool_admin     pset   1 pset_admin    1    1    1  0.4  5.2
  2 pool_other     pset  -1 pset_default   1  65K   2  1.2  8.3
```

Resource sets attached to multiple pools, as pset_default in the example above, are listed multiple times, once for each pool.

示例3 Restricting the Output to the List of Pools

The following example restricts the output to the list of pools

```
% poolstat -p pool_default
      pset
  id pool          size used load
  0 pool_default   8  5.3 10.3

% poolstat -p 'pool_admin pool_default'
      pset
  id pool          size used load
  1 pool_admin     6  4.3  5.3
  0 pool_default   2  1.9  2.0

% poolstat -r all -p 'pool_admin pool_default'
  id pool          type rid rset          min  max size used load
  1 pool_admin     pset   1 pset_admin    1    1    1  0.9  2.3
  2 pool_default   pset  -1 pset_default   1  65K   2  2.0  2.0
```

示例4 Customizing Output

The following example customizes output:

```
% poolstat -r -o pool,rset,size,load
  pool          rset          size load
  pool_default   pset_default   4  4.5
  pool_admin     pset_admin     4  2.1
```

退出状态

The following exit values are returned:

- 0 Successful completion.
- 1 An error occurred.
- 2 Invalid command line options were specified.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/resource-mgmt/resource-pools
Interface Stability	Committed

另请参见 [libpool\(3LIB\)](#), [attributes\(5\)](#)

《[Oracle Solaris 管理：Oracle Solaris Zones、Oracle Solaris 10 Zones 和资源管理](#)》

附注 The system ids associated with resources can change after the system reboots or the resource configuration is altered.

引用名	ports – creates /dev entries and inittab entries for serial lines
用法概要	<code>/usr/sbin/ports [-r rootdir]</code>
描述	devfsadm(1M) is now the preferred command for /dev and /devices and should be used instead of ports. The ports command creates symbolic links in the /dev/term and /dev/cua directories to the serial-port character device files in /devices and adds new entries in /etc/inittab for non-system ports found. System-board ports are given single lower-case letters for names (such as a and b) while other ports are named numerically. ports searches the kernel device tree to find the serial devices attached to the system. It also checks /dev/term and /dev/cua to see what symbolic links to serial devices already exist. ports then performs the following: 1. Assigns new numbers (or letters for system-board ports) to ports that are attached to the system but do not have /dev/term and /dev/cua entries. The numbers or letters assigned are the lowest-unused numbers or letters. 2. Removes dangling links: links from /dev/term and /dev/cua pointing to no-longer-existing ports. 3. Creates new /dev/term and /dev/cua links for new serial devices. If the configuration has not changed, ports exits without doing anything.
Notice to Driver Writers	ports considers devices with a node type of DDI_NT_SERIAL, DDI_NT_SERIAL_MB, DDI_NT_SERIAL_DO, or DDI_NT_SERIAL_MB_DO to be serial port devices. Devices with one of these node types must create minor device names that obey the following conventions when calling ddi_create_minor_node(9F). ■ The minor name for non-system port devices (DDI_NT_SERIAL) consists of an ASCII numeric string, where the first port on the device is named 0, the second named 1, the third named 2, up to the number of ports provided by the device. ■ The minor name for non-system dialout devices (DDI_NT_SERIAL_DO) is the ASCII numeric port name, concatenated with ,cu. For example, the minor name for the first dialout port on the serial board is 0,cu. ■ The minor name for system-board port devices (DDI_NT_SERIAL_MB) consists of a string containing a single ASCII lowercase character, where the first port on the device is named a, the second is named b, the third is named c, for all ports on the device (or up through port z). ■ The minor name for system-board dialout devices (DDI_NT_SERIAL_MB_DO) consists of the lowercase character port name, concatenated with ,cu. For example, the minor name for the first dialout port on the on-board serial device is a,cu. To prevent disks from attempting to automatically generate links for a device, drivers must specify a private node type and refrain from using one of the above node types when calling ddi_create_minor_node(9F).

- 选项**
- The following options are supported:
- `-r rootdir` Causes ports to presume that the `/dev/term`, `/dev/cua`, and `/devices` directories are found under *rootdir*, not directly under `/`.
- 示例**
- 示例 1 Creating the Serial and Dialout Minor Device Nodes
- The following example creates the serial and dialout minor device nodes from the `xkserial` driver's `attach(9E)` function:

```
/*
 * Create the minor number by combining the instance number
 * with the port number.
 */
#define XKNUMPORTS      8
#define XKMINORNUM(i, p) ((i) <= 4 | (p))
#define XKMINORNUM_DO(i, p) ((i) <= 4 | (p) | 0x80)
int
xkserialattach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    int instance, portnum;
    char name[8];
    /* other stuff in attach... */
    instance = ddi_get_instance(dip);
    for (portnum = 0; portnum < XKNUMPORTS; portnum++) {
        /*
         * create the serial port device
         */
        sprintf(name, "%d", portnum);
        ddi_create_minor_node(dip, name, S_IFCHR,
            XKMINORNUM(instance, portnum), DDI_NT_SERIAL, 0);

        /*
         * create the dialout device
         */
        sprintf(name, "%d,cu", portnum);
        ddi_create_minor_node(dip, name, S_IFCHR,
            XKMINORNUM_DO(instance, portnum), DDI_NT_SERIAL_DO, 0);
    }
}
```

示例 2 Installing the `xkserial` Port Driver on a Sun Fire 4800

The following example installs the `xkserial` port driver on a Sun Fire 4800 (with the driver controlling the fictional `XKSerial` 8 port serial board), with these special files in `/devices`:

```
# ls -l /devices/ssm@0,0/pci@18,700000/pci@1/xkserial@f,800000/
crw-r----- 1 root sys  32,  16 Aug 29 00:02 xkserial@2000:0
crw-r----- 1 root sys  32, 144 Aug 29 00:02 xkserial@2000:0,cu
crw-r----- 1 root sys  32,  17 Aug 29 00:02 xkserial@2000:1
```


示例 2 Installing the xkserial Port Driver on a Sun Fire 4800 (续)

```
crw-r----- 1 root sys 32, 145 Aug 29 00:02 xkserial@2000:1,cu
crw-r----- 1 root sys 32, 18 Aug 29 00:02 xkserial@2000:2
crw-r----- 1 root sys 32, 146 Aug 29 00:02 xkserial@2000:2,cu
crw-r----- 1 root sys 32, 19 Aug 29 00:02 xkserial@2000:3
crw-r----- 1 root sys 32, 147 Aug 29 00:02 xkserial@2000:3,cu
crw-r----- 1 root sys 32, 20 Aug 29 00:02 xkserial@2000:4
crw-r----- 1 root sys 32, 148 Aug 29 00:02 xkserial@2000:4,cu
crw-r----- 1 root sys 32, 21 Aug 29 00:02 xkserial@2000:5
crw-r----- 1 root sys 32, 149 Aug 29 00:02 xkserial@2000:5,cu
crw-r----- 1 root sys 32, 22 Aug 29 00:02 xkserial@2000:6
crw-r----- 1 root sys 32, 150 Aug 29 00:02 xkserial@2000:6,cu
crw-r----- 1 root sys 32, 23 Aug 29 00:02 xkserial@2000:7
crw-r----- 1 root sys 32, 151 Aug 29 00:02 xkserial@2000:7,cu
```

/dev/term contain symbolic links to the serial port device nodes in /devices

```
# ls -l /dev/term
/dev/term/0 -> ../../devices/[...]/xkserial@2000:0
/dev/term/1 -> ../../devices/[...]/xkserial@2000:1
/dev/term/2 -> ../../devices/[...]/xkserial@2000:2
/dev/term/3 -> ../../devices/[...]/xkserial@2000:3
/dev/term/4 -> ../../devices/[...]/xkserial@2000:4
/dev/term/5 -> ../../devices/[...]/xkserial@2000:5
/dev/term/6 -> ../../devices/[...]/xkserial@2000:6
/dev/term/7 -> ../../devices/[...]/xkserial@2000:7
```

and /dev/cua contain symbolic links to the dialout port device nodes in /devices

```
# ls -l /dev/cua
/dev/cua/0 -> ../../devices/[...]/xkserial@2000:0,cu
/dev/cua/1 -> ../../devices/[...]/xkserial@2000:1,cu
/dev/cua/2 -> ../../devices/[...]/xkserial@2000:2,cu
/dev/cua/3 -> ../../devices/[...]/xkserial@2000:3,cu
/dev/cua/4 -> ../../devices/[...]/xkserial@2000:4,cu
/dev/cua/5 -> ../../devices/[...]/xkserial@2000:5,cu
/dev/cua/6 -> ../../devices/[...]/xkserial@2000:6,cu
/dev/cua/7 -> ../../devices/[...]/xkserial@2000:7,cu
```

文件

- /dev/term/*n* Logical serial port devices
- /dev/cua/*n* Logical dialout port devices
- /etc/inittab Controls dispatching by init.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[add_drv\(1M\)](#), [devfsadm\(1M\)](#), [drvconfig\(1M\)](#), [attributes\(5\)](#), [devfs\(7FS\)](#), [attach\(9E\)](#), [ddi_create_minor_node\(9F\)](#)

[《Writing Device Drivers》](#)

引用名	poweradm – manage power management properties
用法概要	<pre>poweradm [-v] get [-a all smf platform current] <i>property</i> ... poweradm [-v] set <i>property=value</i> ... poweradm [-v] list poweradm -?</pre>
描述	The poweradm program is used to display and manage the Power Management settings within a Solaris instance. The Power Management properties and their values are stored in the service management facility (see smf(5)). All users can run the list and get forms given in the SYNOPSIS. These commands allow all users to view the current Power Management settings within a Solaris instance. Only users and roles that belong to the Maintenance and Repair RBAC profile can execute the set form of poweradm given in the SYNOPSIS. See also the NOTES section below. poweradm supports the following service property: <pre>administrative-authority</pre> The value for this property is specified as a string and can have the values: ▪ smf ▪ platform ▪ none ...with platform as the default value. The significance of this property is the source of administrative control for power management within the Solaris kernel. That is, this property indicates the origin of time-to-full-capacity and time-to-minimum-responsiveness settings within the Solaris kernel. Only a Solaris user with appropriate privileges can set this property. When the administrative-authority is set to platform the values of time-to-full-capacity and time-to-minimum-responsiveness will be taken from the platform code. Setting these values in SMF using the poweradm command will have no effect upon the values in the kernel, while administrative-authority is set to platform. Under this condition, the poweradm command will indicate that the kernel is currently using the platform values. As these values are changed by the platform administrator, so they will be modified in kernel. Also, commands to enable or disable the power management within the Solaris kernel will come from the platform code. When the administrative-authority is set to smf, the values of time-to-full-capacity and time-to-minimum-responsiveness will be taken from SMF. As these values in SMF are changed by the Solaris administrator, so they will be applied to kernel. Setting these values in the platform will have no effect upon the values in the kernel while administrative-authority is set to smf. Under this condition, he poweradm

command will indicate that the kernel is currently using the values from SMF. As these values are changed by the Solaris administrator, so they will be modified in kernel.

When the `administrative-authority` is set to `none`, power management within the Solaris kernel will be turned off, although the power service will continue to run. Any power management instructions from the platform will be ignored, as will the settings in SMF. Only when `administrative-authority` is set to one of the other values will power management within the Solaris kernel restart, using the settings from the specified source.

`time-to-full-capacity`

Specified in microseconds.

This parameter constrains the dynamic capacity adjustment allowed while the system is in an active state.

This parameter defines the maximum time the system is allowed to reach (re-provision and make available) its full capacity, returning from any lower-capacity/less-responsive state, while it has been using any or all of the PM features falling within this bound.

By default, this value is taken from the platform (for example, `i86pc`), because the default setting for `administrative-authority` is set to `platform`.

Alternatively, if `administrative-authority` is set to `smf`, this value is taken from the definition provided by the power service (that is, SMF). At install time this value is set to be undefined. If the Solaris administrator chooses to modify this property, a value appropriate to the needs of the workload or applications must be picked.

`time-to-minimum-responsiveness`

Specified in milliseconds.

This parameter constrains the dynamic capacity adjustment allowed while the system is in an inactive state.

This parameter defines how long the system is allowed to return to its active state—that is, to provide the minimum capacity required to meet the above `time-to-full-capacity` constraint.

Moderate values (seconds) allow hardware components or subsystems on the platform to be placed in slower-response inactive states; larger values still (for example, 30 seconds to minutes) allow for such as whole system suspension, using techniques such as `suspend-to-RAM`.

By default, this value is taken from the platform (for example, `i86pc`), because the default setting for `administrative-authority` is set to `platform`.

Alternatively, if `administrative-authority` is set to `smf`, this value is taken from the definition provided by the power service (that is, SMF). At install time this value is set to be undefined. If the Solaris administrator chooses to modify this property, a value appropriate to the needs of the workload or applications must be picked.

suspend-enable

By default, only a small number of machines are allowed to suspend. Setting this property to true will permit other machines to attempt to suspend. By default this property is set to false, unless a machine has S3 according the the ACPI kstat module and is on the white list. If these two conditions are met, the default is true. The value of administrative-authority has no effect upon this property.

选项

The following options are supported.

-?

Display a synopsis of available subcommands and options.

-v

Provide verbose output. Can be used with any of the subcommands listed below.

子命令

The following subcommands are supported:

get [-a all|smf|platform|current]

Retrieves the current value of the named property. The -a option can be used to indicate the origin of the value either: SMF (smf), the platform (platform), the current value used by the kernel (current) or all of the preceeding (all). By default, if no origin is specified then current is assumed.

The administrative-authority and suspend-enable properties do not have a platform value

set *property=value...*

Changes the named *property* to the given *value*. The properties administrative-authority and suspend-enable are automatically synchronized to the new value in the kernel. The properties time-to-full-capacity and time-to-minimum-responsiveness are synchronized to the kernel if and only if administrative-authority is set to smf. If the -v option is used and the kernel cannot be updated immediately because administrative-authority is not set to smf, a warning message will be issued.

Only users and roles that belong to the Maintenance and Repair RBAC profile can execute the set subcommand.

list

Lists all the available Power Management properties values and indicate whether power management is active.

退出状态

0

Successful completion.

1

An error occurred.

2

Invalid command line options were specified.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/kernel/power
Interface Stability	Committed

另请参见 [attributes\(5\)](#), [smf\(5\)](#), [smf_security\(5\)](#)

附注 The power service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/system/power:default
```

The properties that can be set by poweradm are defined in that service as:

- `active_control/administrative-authority`, which is described under `administrative-authority` above
- `active_config/time-to-full-capacity`, which is described under `time-to-full-capacity` above
- `active_config/time-to-minimum-responsiveness`, which is described under `time-to-minimum-responsiveness` above
- `suspend/suspend-enable`, which is described under `suspend-enable` above

If the service is disabled, no power management settings will be uploaded to the kernel in the future. Existing settings will not be undone until the next reboot. Disabling all power management is best accomplished by setting `administrative-authority` to `none`.

If `administrative-authority` is set to the value `smf` before both `time-to-full-capacity` and `time-to-minimum-responsiveness` have been set, the service will go into maintenance mode. In such a situation, set `administrative-authority` to the value `none` then set both `time-to-full-capacity` and `time-to-minimum-responsiveness` to the values you want, clear the service and then set `administrative-authority` to `smf`.

To set properties in the `active_config` and `suspend` property groups, the `solaris.smf.value.power_config` authorization is required. To set properties in the `active_control` property group requires the `solaris.smf.value.power_control` authorization is required. Both of these authorizations are part of the Maintenance and Repair profile.

引用名	powertop – 报告和分析影响电源管理的事件
用法概要	<code>powertop [-c <i>processor_id</i>] [-d <i>count</i>] [-t <i>interval</i>] [-v] [-h]</code>
描述	PowerTOP 是一个监测工具，可显示系统利用 CPU 电源管理功能的有效程度。通过在空闲系统上运行此工具，用户可以查看在不同电源状态下 CPU 可以运行多久。理想状态下，未使用的（空闲）系统将始终以最低电源状态运行，但是由于后台用户和内核活动（随机软件定期转醒以轮询状态），空闲系统的实际功耗会高于理想状态。 此工具定期分析系统活动，并显示处理器在每个支持的电源状态下运行的时间长度的汇总。它还会显示最容易导致 CPU 转醒而耗能更多的活动。通过此报告，用户可以发现和诊断系统中存在问题的区域，并优化其电源效率。 PowerTOP 将计算使 CPU 无法进入较低功耗状态的活动量的均值，并将其显示在 "Wakeups-from-idle per second"（每秒从空闲状态转醒的次数）字段中。该值等于转醒总次数除以当前间隔。请注意，屏幕上不会始终显示所有事件。 在执行过程中，用户可以按 R 键强制刷新分析结果。间隔时间将恢复为缺省值或指定值。要退出此应用程序，用户必须按 Q 键。 PowerTOP 允许您冻结其三个子窗口中的每个子窗口，以便更好地分析每个子窗口的内容。同时运行三个不同的 DTrace 脚本时，冻结子窗口可降低 CPU 消耗。在该工具运行时按 i、f 或 e 键可以激活此功能。按上述各键将分别冻结空闲、频率或事件报告。再一次按对应的键可取消冻结。在要分析的应用程序处于运行状态时调用此功能最为有用。 PowerTOP 在某些虚拟域中运行。但是，空闲状态转换报告可能不准确，因为物理 CPU 可由不同的虚拟 CPU 共享。转醒计数和事件报告均显示有关当前虚拟化环境的信息。
选项	所支持的选项如下： -c [<i>processor_id</i>] 指定该工具应监测的 CPU。 -d [<i>count</i>] 将系统活动的 <i>count</i> 分析结果转储到屏幕。 -h 显示命令的用法。 -t [<i>interval</i>] 指定该工具分析系统的间隔时间（以秒为单位）。可能的值在 1 到 30 范围内；缺省值为 5 秒。 -v 切换到详细模式，包括事件报告中内核循环子系统的触发说明。

示例

示例1 设置时间间隔
以下命令将间隔设置为两秒。

`% powertop -t 2`

示例2 分析和转储系统活动
以下命令四次分析系统活动并将其转储至标准输出。

`% powertop -d 4`

示例3 报告循环子系统活动
以下命令报告循环子系统活动。

`% powertop -v`

示例4 分析特定处理器上的活动
以下命令运行 PowerTOP 并仅显示 CPU 3 的数据：

`% powertop -c 3`

退出状态

- 0
操作成功。
- 1
出现错误。
- 2
使用方法不正确。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
体系结构	x86、SPARC
可用性	diagnostic/powertop
接口稳定性	Volatile（可变）

另请参见

[kstat\(1M\)](#)、[psrinfo\(1M\)](#)、[uadmin\(2\)](#)、[libdevinfo\(3LIB\)](#)、[attributes\(5\)](#)、[cpr\(7\)](#)、[pm\(7D\)](#)、[pm-com](#)
非 SunOS 手册页中的 [xscreensaver\(1\)](#) 和 [dtpower\(1M\)](#)。

用法

您必须具有 DTrace 特权才能运行 PowerTOP，并且必须具有根用户（超级用户）特权或具有主管理员角色才能通过该工具获得改进系统的建议。

引用名	pppd – point to point protocol daemon	
用法概要	pppd [<i>tty_name</i>] [<i>speed</i>] [<i>options</i>]	
描述	The point-to-point protocol (PPP) provides a method for transmitting datagrams over serial point-to-point links. PPP is composed of three components: a facility for encapsulating datagrams over serial links, an extensible link control protocol (LCP), and a family of network control protocols (NCP) for establishing and configuring different network-layer protocols. The encapsulation scheme is provided by driver code in the kernel. pppd provides the basic LCP authentication support and several NCPs for establishing and configuring the Internet Protocol (referred to as the IP Control Protocol or “IPCP”) and IPv6 (IPV6CP).	
选项	The following sections discuss the pppd options:	
Options Files	Options are taken from files and the command line. pppd reads options from the files <code>/etc/ppp/options</code>, <code>\$HOME/.ppprc</code> and <code>/etc/ppp/options.<i>ttyname</i></code> (in that order) before processing the options on the command line. (Command-line options are scanned for the terminal name before the <code>options.<i>ttyname</i></code> file is read.) To form the name of the <code>options.<i>ttyname</i></code> file, the initial <code>/dev/</code> is removed from the terminal name, and any remaining forward slash characters (<code>/</code>) are replaced with dots. For example, with serial device <code>/dev/cua/a</code>, option file <code>/etc/ppp/options.cua.a</code> is read. An options file is parsed into a series of words that are delimited by whitespace. Whitespace can be included in a word by enclosing the word in double-quotes (<code>"</code>). A backslash (<code>\</code>) quotes the succeeding character. A hash (<code>#</code>) starts a comment, which continues until the end of the line. There is no restriction on using the <code>file</code> or <code>call</code> options within an options file.	
Frequently Used Options	<code><tty_name></code>	Communicate over the named device. The string <code>/dev/</code> is prepended if necessary. If no device name is given, or if the name of the terminal connected to the standard input is given, pppd uses that terminal and does not fork to put itself in the background. A value for this option from a privileged source cannot be overridden by a non-privileged user.
	<code><speed></code>	Set the baud rate to <code><speed></code> (a decimal number). The default is to leave the baud rate unchanged. This option is normally needed for dial-out only.
	<code>asyncmap <map></code>	Set the async character map to <code><map></code> . The map describes which control characters cannot be successfully received over the serial line. pppd asks the peer to send these characters as a 2-byte escape sequence. The argument is a 32 bit hex number, with each bit representing a character to escape. Bit 0 (00000001) represents the character 0x00; bit 31 (80000000) represents the character 0x1f or <code>^_</code> . If multiple <code>asyncmap</code> options are given, the values are ORed together. If no <code>asyncmap</code> option is given, pppd attempts to negotiate a value of 0. If the

	peer agrees, this disables escaping of the standard control characters. Use the <code>default-asyncmap</code> option to disable negotiation and escape all control characters.
<code>auth</code>	Require the peer to authenticate itself before allowing network packets to be sent or received. This option is the default if the system has a default route. If the <code>auth</code> or the <code>noauth</code> option is not specified, <code>pppd</code> allows the peer to use only those IP addresses to which the system does not already have a route.
<code>call name</code>	Read options from the file <code>/etc/ppp/peers/name</code> . This file may contain privileged options, including <code>noauth</code> , even if <code>pppd</code> is not being run by root. The <i>name</i> string may not begin with a slash ("/") or include consecutive periods ("..") as a pathname component.
<code>callback number</code>	Request a callback to the given telephone number using Microsoft CBCP.
<code>connect script</code>	Use the executable or shell command specified by <i>script</i> to set up the serial line. This script would typically use the chat(1M) program to dial the modem and start the remote PPP session. A value for this option originating from a privileged source cannot be overridden by a non-privileged user.
<code>crtcts</code>	Use hardware flow control, that is, RTS/CTS, to control the flow of data on the serial port. If the <code>crtcts</code> , <code>nocrtcts</code> , <code>cdtrcts</code> or <code>nocdtrcts</code> option is not provided, the hardware flow control setting for the serial port is left unchanged. Some serial ports lack a true RTS output and use this mode to implement unidirectional flow control. The serial port suspends transmission when requested by the modem by means of CTS but cannot request the modem to stop sending to the computer. This mode allows the use of DTR as a modem control line.
<code>defaultroute</code>	Add a default route to the system routing tables when IPCP negotiation successfully completes, using the peer as the gateway. This entry is removed when the PPP connection is broken. This option is privileged if the <code>nodefaultroute</code> option is specified.
<code>disconnect script</code>	Run the executable or shell command specified by <i>script</i> after <code>pppd</code> terminates the link. Typically, this script is used to command the modem to hang up if hardware modem control signals are not available. <code>disconnect</code> is not run if the modem has already hung up. A value for this option originating from a privileged source cannot be overridden by a non-privileged user.
<code>escape xx,yy,...</code>	Specifies that certain characters be escaped on transmission regardless of whether the peer requests them to be escaped with its <code>async</code> control

character map. The characters to be escaped are specified as a list of hex numbers separated by commas. Note that almost any character can be specified for the escape option, unlike the `asynctest` option which allows only control characters to be specified. Characters that cannot be escaped are those containing hex values 0x20 through 0x3f and 0x5e.

<code>file name</code>	Read options from file <i>name</i> . If this option is used on the command line or in <code>\$HOME/.ppprc</code> , the file must be readable by the user invoking <code>pppd</code> . See Options Files for a list of files that <code>pppd</code> always reads, regardless of the use of this option.
<code>init script</code>	Run the executable or shell command specified by <i>script</i> to initialize the serial line. This script would typically use the chat(1M) program to configure the modem to enable auto-answer. A value for this option from a privileged source cannot be overridden by a non-privileged user.
<code>lock</code>	Directs <code>pppd</code> to create a UUCP-style lock file for the serial device to ensure exclusive access to the device.
<code>mru n</code>	Set the Maximum Receive Unit (MRU) value to <i>n</i> . <code>pppd</code> asks the peer to send packets of no more than <i>n</i> bytes. Minimum MRU value is 128. Default MRU value is 1500. A value of 296 is recommended for slow links (40 bytes for TCP/IP header + 256 bytes of data). For IPv6, MRU must be at least 1280.
<code>mtu n</code>	Set the Maximum Transmit Unit (MTU) value to <i>n</i> . Unless the peer requests a smaller value via MRU negotiation, <code>pppd</code> requests the kernel networking code to send data packets of no more than <i>n</i> bytes through the PPP network interface. For IPv6, MTU must be at least 1280.
<code>passive</code>	Enables the "passive" option in the LCP. With this option, <code>pppd</code> attempts to initiate a connection; if no reply is received from the peer, <code>pppd</code> waits passively for a valid LCP packet instead of exiting, as it would without this option.

Options

`<local_IP_address>:<remote_IP_address>`

Set the local and/or remote interface IP addresses. Either one may be omitted, but the colon is required. The IP addresses are specified with a host name or in decimal dot notation, for example: `:10.1.2.3`. The default local address is the first IP address of the system unless the `noipdefault` option is provided. The remote address is obtained from the peer if not specified in any option. Thus, in simple cases, this option is not required. If a local and/or remote IP address is specified with this option, `pppd` will not accept a different value from the peer in the IPCP negotiation unless the `ipcp-accept-local` and/or `ipcp-accept-remote` options are given, respectively.

allow-fcs *fcs-type*

Set allowable FCS type(s) for data sent to the peer. The *fcs-type* is a comma-separated list of "crc16", "crc32", "null", or integers. By default, all known types are allowed. If this option is specified and the peer requests a type not listed, a LCP Configure-Nak is sent to request only the listed types.

allow-ip *address(es)*

Allow peers to use the given IP address or subnet without authenticating themselves. The parameter is parsed in the same manner as each element of the list of allowed IP addresses is parsed in the secrets files. See the [Authentication](#) section for more details.

bsdcomp *nr,nt*

Request that the peer compress packets that it sends using the BSD-Compress scheme, with a maximum code size of *nr* bits, and agree to compress packets sent to the peer with a maximum code size of *nt* bits. If *nt* is not specified, it defaults to the value given for *nr*. Values in the range 9 to 15 may be used for *nr* and *nt*; larger values provide better compression but consume more kernel memory for compression dictionaries. Alternatively, a value of 0 for *nr* or *nt* disables compression in the corresponding direction. Use `nobsdcomp` or `bsdcomp 0` to disable BSD-Compress compression entirely. If this option is read from a privileged source, a nonprivileged user may not specify a code size larger than the value from the privileged source.

cdtrcts

Use a non-standard hardware flow control such as DTR/CTS to control the flow of data on the serial port. If the `crtcts`, `nocrtcts`, `cdtrcts` or `nocdtrcts` option is not specified, the hardware flow control setting for the serial port is left unchanged. Some serial ports lack a true RTS output. Such serial ports use this mode to implement true bi-directional flow control. Note that this flow control mode does not permit using DTR as a modem control line.

chap-interval *n*

If this option is given, `pppd` will rechallenge the peer every *n* seconds.

chap-max-challenge *n*

Set the maximum number of CHAP challenge transmissions to *n* (default 10).

chap-restart *n*

Set the CHAP restart interval (retransmission timeout for challenges) to *n* seconds. The default is 3.

connect-delay *n*

Wait for up to *n* milliseconds after the connect script finishes for a valid PPP packet from the peer. When the wait period elapses or when a valid PPP packet is received from the peer, `pppd` begins negotiation by sending its first LCP packet. The default value is 1000 (1 second). A wait period applies only if the `connect` or `pty` option is used.

datarate *n*

Set maximum data rate to *n* (in bytes per second) when using the `pty`, `notty`, `record`, or `socket` options.

debug

Enables connection debugging facilities. If this option is given, pppd logs the contents of all control packets sent or received in a readable form. The packets are logged through syslog with facility daemon and level debug. This information can be directed to a file by configuring `/etc/syslog.conf` appropriately.

default-asynmap

Disable asynmap negotiation, forcing all control characters to be escaped for both the transmit and the receive direction.

default-fcs

Disable FCS Alternatives negotiation entirely. By default, no FCS Alternatives option is sent to the peer, but the option is accepted. If this option is specified by the peer, then LCP Configure-Reject is sent.

default-mru

Disable MRU [Maximum Receive Unit] negotiation. With this option, pppd uses the default MRU value of 1500 bytes for the transmit and receive directions.

deflate *nr,nt,e*

Request that the peer compress packets that it sends, using the deflate scheme, with a maximum window size of 2^{nr} bytes, and agree to compress packets sent to the peer with a maximum window size of 2^{nt} bytes and effort level of *e* (1 to 9). If *nt* is not specified, it defaults to the value given for *nr*. If *e* is not specified, it defaults to 6. Values in the range 9 to 15 may be used for *nr* and *nt*; larger values provide better compression but consume more kernel memory for compression dictionaries. (Value 8 is not permitted due to a zlib bug.) Alternatively, a value of 0 for *nr* or *nt* disables compression in the corresponding direction. Use `nodeflate` or `deflate 0` to disable deflate compression entirely. (Note: pppd requests deflate compression in preference to BSD-Compress if the peer can do either.) If this option is read from a privileged source, a nonprivileged user may not specify a code size larger than the value from the privileged source.

demand

Initiate the link only on demand, that is, when data traffic is present. With this option, the remote IP address must be specified by the user on the command line or in an options file. pppd initially configures and enables the interface for IP traffic without connecting to the peer. When traffic is available, pppd connects to the peer and performs negotiation, authentication and other actions. When completed, pppd passes data packets across the link. The demand option implies the persist option. If this behavior is not desired, use the nopersist option after the demand option. The idle and holdoff options can be used in conjunction with the demand option.

domain *d*

Append the domain name *d* to the local host name for authentication purposes. For example, if `gethostname()` returns the name `porsche`, but the fully qualified domain name is `porsche.Quotron.COM`, you could specify `domain Quotron.COM`. With this configuration, pppd uses the name `porsche.Quotron.COM` for accessing secrets in the secrets file and as the default name when authenticating to the peer. This option is privileged.

endpoint endpoint-value

Set the endpoint discriminator (normally used for RFC 1990 Multilink PPP operation).

The *endpoint-value* consists of a class identifier and a class-dependent value. The class identifier is one of "null," "local," "IP," "MAC," "magic," "phone," or a decimal integer. If present, the class-dependent value is separated from the identifier by a colon (":") or period ("."). This value may be a standard dotted-decimal IP address for class "IP," an optionally colon-or-dot separated hex Ethernet address for class "MAC" (must have 6 numbers), or an arbitrary string of bytes specified in hex with optional colon or dot separators between bytes. Although this option is available, this implementation does not support multilink.

fcs fcs-type

Set FCS type(s) desired for data sent by the peer. The *fcs-type* is a comma-separated list of `crc16`, `crc32`, `null`, or integers. By default, an FCS Alternatives option is not specified, and the medium-dependent FCS type is used. If this option is specified and the peer sends an LCP Configure-Nak, only the listed types are used. If none are in common, the FCS Alternatives option is omitted from the next LCP Configure-Request to drop back to the default.

hide-password

When logging the contents of PAP packets, this option causes `pppd` to exclude the password string from the log. This is the default.

holdoff n

Specifies how many seconds to wait before re-initiating the link after it terminates. This option is effective only if the `persist` or `demand` option is used. The holdoff period is not applied if the link is terminated because it was idle.

ident string

Set the LCP Identification string. The default value is a version string similar to that displayed by the `--version` option.

idle n

Specifies that `pppd` must disconnect if the link is idle for *n* seconds. The link is idle when no data packets (i.e. IP packets) are being sent or received. Do not use this option with the `persist` option but without the `demand` option.

ipcp-accept-local

With this option, `pppd` accepts the peer's idea of the local IP address, even if the local IP address is specified in an option.

ipcp-accept-remote

With this option, `pppd` accepts the peer's idea of its remote IP address, even if the remote IP address is specified in an option.

ipcp-max-configure n

Set the maximum number of IPCP Configure-Request transmissions to *n* (default 10).

`ipcp-max-failure n`

Set the maximum number of IPCP Configure-NAKs sent before sending Configure-Rejects instead to *n* (default 10).

`ipcp-max-terminate n`

Set the maximum number of IPCP terminate-request transmissions to *n* (default 3).

`ipcp-restart n`

Set the IPCP restart interval (retransmission timeout) to *n* seconds (default 3).

`ipparam string`

Provides an extra parameter to the ip-up and ip-down scripts. When this option is given, the *string* supplied is given as the sixth parameter to those scripts. See the [Scripts](#) section.

`ipv6 <local_interface_identifier>, <remote_interface_identifier>`

Set the local and/or remote 64-bit interface identifier. Either one may be omitted. The identifier must be specified in standard ASCII notation of IPv6 addresses (for example: `::dead:beef`). If the `ipv6cp-use-ipaddr` option is given, the local and remote identifiers are derived from the respective IPv4 addresses (see above). The `ipv6cp-use-persistent` option can be used instead of the `ipv6 <local>, <remote>` option.

`ipv6cp-accept-local`

Accept peer's interface identifier for the local link identifier.

`ipv6cp-max-configure n`

Set the maximum number of IPv6CP Configure-Request transmissions to *n* (default 10).

`ipv6cp-max-failure n`

Set the maximum number of IPv6CP Configure-NAKs sent before sending Configure-Rejects instead to *n* (default 10).

`ipv6cp-max-terminate n`

Set the maximum number of IPv6CP terminate-request transmissions to *n* (default 3).

`ipv6cp-restart n`

Set the IPv6CP restart interval (retransmission timeout) to *n* seconds (default 3).

`ipv6cp-use-ipaddr`

If either the local or remote IPv6 address is unspecified, use the corresponding configured IPv4 address as a default interface identifier. (This option uses the configured addresses, not the negotiated addresses. Do not use it with `ipcp-accept-local` if the local IPv6 identifier is unspecified or with `ipcp-accept-remote` if the remote IPv6 identifier is unspecified.)

`ipv6cp-use-persistent`

Use uniquely-available persistent value for link local address.

`kdebug n`

Enable debugging code in the kernel-level PPP driver. Argument *n* is the sum of the following values: 1 to enable general debug messages, 2 to request that contents of received

packets be printed, and 4 to request contents of transmitted packets be printed. Messages printed by the kernel are logged by `syslogd(1M)` to a file directed in the `/etc/syslog.conf` configuration file. Do not use the `kdebug` option to debug failed links. Use the `debug` option instead.

`lcp-echo-failure n`

If this option is given, `pppd` presumes the peer to be dead if *n* LCP Echo-Requests are sent without receiving a valid LCP Echo-Reply. If this happens, `pppd` terminates the connection. This option requires a non-zero value for the `lcp-echo-interval` parameter. This option enables `pppd` to terminate after the physical connection is broken (for example, if the modem has hung up) in situations where no hardware modem control lines are available.

`lcp-echo-interval n`

If this option is given, `pppd` sends an LCP Echo-Request frame to the peer every *n* seconds. Normally the peer responds to the Echo-Request by sending an Echo-Reply. This option can be used with the `lcp-echo-failure` option to detect that the peer is no longer connected.

`lcp-max-configure n`

Set the maximum number of LCP Configure-Request transmissions to *n* (default 10).

`lcp-max-failure n`

Set the maximum number of LCP Configure-NAKs sent before starting to send Configure-Rejects instead to *n* (default 10).

`lcp-max-terminate n`

Set the maximum number of LCP Terminate-Request transmissions to *n* (default 3).

`lcp-restart n`

Set the LCP restart interval (retransmission timeout) to *n* seconds (default 3).

`linkname name`

Sets the logical name of the link to *name*. `pppd` creates a file named `ppp-name.pid` in `/var/run` containing its process ID. This is useful in determining which instance of `pppd` is responsible for the link to a given peer system. This is a privileged option.

`local`

Do not use modem control lines. With this option, `pppd` ignores the state of the CD (Carrier Detect) signal from the modem and does not change the state of the DTR (Data Terminal Ready) signal.

`logfd n`

Send log messages to file descriptor *n*. `pppd` sends log messages to (at most) one file or file descriptor (as well as sending the log messages to `syslog`), so this option and the `logfile` option are mutually exclusive. By default `pppd` sends log messages to `stdout` (file descriptor 1) unless the serial port is open on `stdout`.

`logfile filename`

Append log messages to the file *filename* (and send the log messages to `syslog`). The file is opened in append mode with the privileges of the user who invoked `pppd`.

login

Use the system password database for authenticating the peer using PAP, and record the user in the system wtmp file. Note that the peer must have an entry in the `/etc/ppp/pap-secrets` file and the system password database to be allowed access.

maxconnect *n*

Terminate the connection after it has been available for network traffic for *n* seconds (that is, *n* seconds after the first network control protocol starts). An LCP Time-Remaining message is sent when the first NCP starts, and again when 5, 2, and 0.5 minutes are remaining.

maxfail *n*

Terminate after *n* consecutive failed connection attempts. A value of 0 means no limit. The default value is 10.

modem

Use the modem control lines. This option is the default. With this option, pppd waits for the CD (Carrier Detect) signal from the modem to be asserted when opening the serial device (unless a connect script is specified), and drops the DTR (Data Terminal Ready) signal briefly when the connection is terminated and before executing the connect script.

ms-dns <*addr*>

If pppd is acting as a server for Microsoft Windows clients, this option allows pppd to supply one or two DNS (Domain Name Server) addresses to the clients. The first instance of this option specifies the primary DNS address; the second instance (if given) specifies the secondary DNS address. If the first instance specifies a name that resolves to multiple IP addresses, then the first two addresses are used. (This option is present in some older versions of pppd under the name `dns-addr`.)

ms-lanman

If pppd connects as a client to a Microsoft server and uses MS-CHAPv1 for authentication, this option selects the LAN Manager password style instead of Microsoft NT.

ms-wins <*addr*>

If pppd acts as a server for Microsoft Windows or Samba clients, this option allows pppd to supply one or two WINS (Windows Internet Name Services) server addresses to the clients. The first instance of this option specifies the primary WINS address; the second instance (if given) specifies the secondary WINS address. As with `ms-dns`, if the name specified resolves to multiple IP addresses, then the first two will be taken as primary and secondary.

name *name*

Set the name of the local system for authentication purposes to *name*. This is a privileged option. With this option, pppd uses lines in the secrets files that have *name* as the second field to look for a secret to use in authenticating the peer. In addition, unless overridden with the `user` option, *name* is used as the name to send to the peer when authenticating the local system. (Note that pppd does not append the domain name to *name*.)

no-accm-test

Disable use of asyncmap (ACCM) checking using LCP Echo-Request messages. If the `lcp-echo-failure` is used on an asynchronous line, `pppd` includes all control characters in the first *n* LCP Echo-Request messages. If the asyncmap is set incorrectly, the link drops rather than continue operation with random failures. This option disables that feature.

noaccomp

Disable HDLC Address/Control compression in both directions (send and receive).

noauth

Do not require the peer to authenticate itself. This option is privileged.

nobsdcomp

Disables BSD-Compress compression; `pppd` will not request or agree to compress packets using the BSD-Compress scheme. This option is not necessary if `noccp` is specified.

noccp

Disable CCP (Compression Control Protocol) negotiation. This option should only be required if the peer has bugs or becomes confused by requests from `pppd` for CCP negotiation. If CCP is disabled, then BSD and deflate compression do not need to be separately disabled.

nocrtscts

Disable hardware flow control (i.e. RTS/CTS) on the serial port. If the `crtscts`, `nocrtscts`, `cdtrcts` or `nocdtrcts` options are not given, the hardware flow control setting for the serial port is left unchanged.

nocdtrcts

This option is a synonym for `nocrtscts`. Either option will disable both forms of hardware flow control.

nodefaultroute

Disable the default route option. You can prevent non-root users from creating default routes with `pppd` by placing this option in the `/etc/ppp/options` file.

nodeflate

Disables deflate compression; `pppd` will not request or agree to compress packets using the deflate scheme. This option is not necessary if `noccp` is specified.

nodeflatedraft

Do not use Internet Draft (incorrectly assigned) algorithm number for deflate compression. This option is not necessary if `noccp` is specified.

nodetach

Do not detach from the controlling terminal. Without this option, `pppd` forks to become a background process if a serial device other than the terminal on the standard input is specified.

noendpoint

Do not send or accept the Multilink Endpoint Discriminator option.

noident

Disable use of LCP Identification. LCP Identification messages will not be sent to the peer, but received messages will be logged. (Specify this option twice to completely disable LCP Identification. In this case, pppd sends LCP Code-Reject in response to received LCP Identification messages.)

noip

Disable IPCP negotiation and IP communication. Use this option only if the peer has bugs or becomes confused by requests from pppd for IPCP negotiation.

noipv6

Disable IPv6CP negotiation and IPv6 communication. IPv6 is not enabled by default.

noipdefault

Disables the default behavior when no local IP address is specified, which is to determine (if possible) the local IP address from the hostname. With this option, the peer must supply the local IP address during IPCP negotiation (unless it specified explicitly on the command line or in an options file).

nolog

Do not send log messages to a file or file descriptor. This option cancels the `logfd` and `logfile` options. `nologfd` acts as an alias for this option.

nomagic

Disable magic number negotiation. With this option, pppd cannot detect a looped-back line. Use this option only if the peer has bugs. Do not use this option to work around the “Serial line is looped back” error message.

nopam

This privileged option disables use of pluggable authentication modules. If this option is specified, pppd reverts to standard authentication mechanisms. The default is not to use PAM.

nopcomp

Disable protocol field compression negotiation in the receive and the transmit direction.

nopersist

Exit once a connection has been made and terminated. This is the default unless the `persist` or `demand` option is specified.

noplink

Cause pppd to use `I_LINK` instead of `I_PLINK`. This is the default. When `I_LINK` is used, the system cleans up terminated interfaces (even when `SIGKILL` is used) but does not allow [ifconfig\(1M\)](#) to unplumb PPP streams or insert or remove modules dynamically. Use the `plink` option if [ifconfig\(1M\)](#) `modinsert`, `modremove` or `unplumb` support is needed.

nopredictor1

Do not accept or agree to Predictor-1 compression. (This option is accepted for compatibility. The implementation does not support Predictor-1 compression.)

noproxyarp

Disable the proxyarp option. If you want to prevent users from creating proxy ARP entries with pppd, place this option in the `/etc/ppp/options` file.

notty

Normally, pppd requires a terminal device. With this option, pppd allocates itself a pseudo-tty master/slave pair and uses the slave as its terminal device. pppd creates a child process to act as a character shunt to transfer characters between the pseudo-tty master and its standard input and output. Thus, pppd transmits characters on its standard output and receives characters on its standard input even if they are not terminal devices. This option increases the latency and CPU overhead of transferring data over the ppp interface as all of the characters sent and received must flow through the character shunt process. An explicit device name may not be given if this option is used.

novj

Disable Van Jacobson style TCP/IP header compression in both the transmit and the receive direction.

novjccomp

Disable the connection-ID compression option in Van Jacobson style TCP/IP header compression. With this option, pppd does not omit the connection-ID byte from Van Jacobson compressed TCP/IP headers, nor does it ask the peer to do so. This option is unnecessary if novj is specified.

pam

This privileged option enables use of PAM. If this is specified, pppd uses the [pam\(3PAM\)](#) framework for user authentication with a service name of "ppp" if the login option and PAP authentication are used. The default is not to use PAM.

papcrypt

Indicates that pppd should not accept a password which, before encryption, is identical to the secret from the `/etc/ppp/pap-secrets` file. Use this option if the secrets in the `pap-secrets` file are in [crypt\(3C\)](#) format.

pap-max-authreq *n*

Set the maximum number of PAP authenticate-request transmissions to *n* (default 10).

pap-restart *n*

Set the PAP restart interval (retransmission timeout) to *n* seconds (default 3).

pap-timeout *n*

Set the maximum time that pppd waits for the peer to authenticate itself with PAP to *n* seconds (0= no limit). The default is 30 seconds.

password *string*

Password string for authentication to the peer.

persist

Do not exit after a connection is terminated; instead try to reopen the connection.

plink

Cause pppd to use I_PLINK instead of I_LINK. The default is to use I_LINK, which cleans up terminated interface (even if SIGKILL is used), but does not allow **ifconfig(1M)** to unplumb PPP streams or insert or remove modules dynamically. Use this option if **ifconfig(1M)** modinsert/modremove/unplumb support is needed. See also the **plumbed** option.

plugin filename

Load the shared library object file *filename* as a plugin. This is a privileged option. Unless the filename specifies an explicit path, */etc/ppp/plugins* and */usr/lib/inet/ppp* will be searched for the object to load in that order.

plumbed

This option indicates that pppd should find a plumbed interface and use that for the session. If IPv4 addresses or IPv6 interface IDs or link MTU are otherwise unspecified, they are copied from the interface selected. This mode mimics some of the functionality of the older aspppd implementation and may be helpful when pppd is used with external applications that use **ifconfig(1M)**.

pppmux timer

Enable PPP Multiplexing option negotiation and set transmit multiplexing timeout to *timer* microseconds.

privgroup group-name

Allows members of group *group-name* to use privileged options. This is a privileged option. Because there is no guarantee that members of *group-name* cannot use pppd to become root themselves, you should be careful using this option. Consider it equivalent to putting the members of *group-name* in the root or sys group.

proxyarp

Add an entry to the system's Address Resolution Protocol (ARP) table with the IP address of the peer and the Ethernet address of this system. When you use this option, the peer appears to other systems to be on the local Ethernet. The remote address on the PPP link must be in the same subnet as assigned to an Ethernet interface.

pty script

Specifies that the command *script*, and not a specific terminal device is used for serial communication. pppd allocates itself a pseudo-tty master/slave pair and uses the slave as its terminal device. *script* runs in a child process with the pseudo-tty master as its standard input and output. An explicit device name may not be given if this option is used. (Note: if the record option is used in conjunction with the pty option, the child process will have pipes on its standard input and output.)

receive-all

With this option, pppd accepts all control characters from the peer, including those marked in the receive asyncmap. Without this option, pppd discards those characters as specified in RFC 1662. This option should be used only if the peer has bugs, as is often found with dial-back implementations.

record *filename*

Directs pppd to record all characters sent and received to a file named *filename*. *filename* is opened in append mode, using the user's user-ID and permissions. Because this option uses a pseudo-tty and a process to transfer characters between the pseudo-tty and the real serial device, it increases the latency and CPU overhead of transferring data over the PPP interface. Characters are stored in a tagged format with timestamps that can be displayed in readable form using the pppdump(1M) program. This option is generally used when debugging the kernel portion of pppd (especially CCP compression algorithms) and not for debugging link configuration problems. See the debug option.

remotename *name*

Set the assumed name of the remote system for authentication purposes to *name*. Microsoft WindowsNT does not provide a system name in its CHAP Challenge messages, and this option is often used to work around this problem.

refuse-chap

With this option, pppd will not agree to authenticate itself to the peer using standard Challenge Handshake Authentication Protocol (CHAP). (MS-CHAP is not affected.)

refuse-mschap

Do not agree to authenticate to peer with MS-CHAPv1. If this option is specified, requests for MS-CHAPv1 authentication from the peer are declined with LCP Configure-Nak. That option does not disable any other form of CHAP.

refuse-mschapv2

Do not agree to authenticate to peer with MS-CHAPv2. If specified, this option requests that MS-CHAPv2 authentication from the peer be declined with LCP Configure-Nak. That option does not disable any other form of CHAP.

refuse-pap

With this option, pppd will not agree to authenticate itself to the peer using Password Authentication Protocol (PAP).

require-chap

Require the peer to authenticate itself using standard CHAP authentication. MS-CHAP is not affected.

require-mschap

Require the peer to authenticate itself using MS-CHAPv1 authentication.

require-mschapv2

Require the peer to authenticate itself using MS-CHAPv2 authentication.

require-pap

Require the peer to authenticate itself using PAP authentication.

show-password

When logging contents of PAP packets, this option causes pppd to show the password string in the log message.

silent

With this option, `pppd` will not transmit LCP packets to initiate a connection until a valid LCP packet is received from the peer. This is like the “passive” option with older versions of `pppd` and is retained for compatibility, but the current `passive` option is preferred.

small-accm-test

When checking the `asynctest` (ACCM) setting, `pppd` uses all 256 possible values by default. See `no-accm-test`. This option restricts the test so that only the 32 values affected by standard ACCM negotiation are tested. This option is useful on very slow links.

socket *host:port*

Connect to given host and port using TCP and run PPP over this connection.

sync

Use synchronous HDLC serial encoding instead of asynchronous. The device used by `pppd` with this option must have sync support. Currently supports `zs`, `se`, and `hsi` drivers.

unit *n*

Set PPP interface unit number to *n*, if possible.

updetach

With this option, `pppd` detaches from its controlling terminal after establishing the PPP connection. When this is specified, messages sent to `stderr` by the connect script, usually `chat(1M)`, and debugging messages from the `debug` option are directed to `pppd`'s standard output.

usehostname

Enforce the use of the hostname with domain name appended, if given, as the name of the local system for authentication purposes. This overrides the `name` option. Because the `name` option is privileged, this option is normally not needed.

usepeerdns

Ask the peer for up to two DNS server addresses. Addresses supplied by the peer, if any, are passed to the `/etc/ppp/ip-up` script in the environment variables `DNS1` and `DNS2`. In addition, `pppd` creates an `/etc/ppp/resolv.conf` file containing one or two nameserver lines with the address(es) supplied by the peer.

user *name*

Sets the name used for authenticating the local system to the peer to *name*.

vj-max-slots *n*

Sets the number of connection slots to be used by the Van Jacobson TCP/IP header compression and decompression code to *n*, which must be between 2 and 16 (inclusive).

welcome *script*

Run the executable or shell command specified by *script* before initiating PPP negotiation, after the connect script, if any, has completed. A value for this option from a privileged source cannot be overridden by a non-privileged user.

`xonxoff`
Use software flow control, that is, XON/XOFF, to control the flow of data on the serial port.

Obsolete Options	The following options are obsolete:
<code>+ua name</code>	Read a PAP user name and password from the file <i>name</i> . This file must have two lines for name and password. Name and password are sent to the peer when the peer requests PAP authentication.
<code>+ipv6</code>	Enable IPv6 and IPv6CP without specifying interface identifiers.
<code>--version</code>	Show version number and exit.
<code>--help</code>	Show brief help message and exit.

扩展描述 The following sections discuss miscellaneous features of pppd:

Security pppd allows system administrators to provide legitimate users with PPP access to a server machine without fear of compromising the security of the server or the network it runs on. Access control is provided by restricting IP addresses the peer may use based on its authenticated identity (if any), and through restrictions on options a non-privileged user may use. Options that permit potentially insecure configurations are privileged. Privileged options are accepted only in files that are under the control of the system administrator or when pppd is being run by root.

By default, pppd allows an unauthenticated peer to use a given IP address only if the system does not already have a route to that IP address. For example, a system with a permanent connection to the wider Internet will normally have a default route, meaning all peers must authenticate themselves to set up a connection. On such a system, the `auth` option is the default. Conversely, a system with a PPP link that comprises the only connection to the Internet probably does not possess a default route, so the peer can use virtually any IP address without authenticating itself.

Security-sensitive options are privileged and cannot be accessed by a non-privileged user running pppd, either on the command line, in the user's `$HOME/.ppprc` file, or in an options file read using the `file` option. Privileged options may be used in `/etc/ppp/options` file or in an options file read using the `call` option. If pppd is run by the root user, privileged options can be used without restriction. If the `/etc/ppp/options` file does not exist, then only root may invoke pppd. The `/etc/ppp/options` file must be created (but may be empty) to allow ordinary non-root users to access pppd.

When opening the device, pppd uses the invoking user's user ID or the root UID (that is, 0), depending if the device name was specified by the user or the system administrator. If the device name comes from a privileged source, that is, `/etc/ppp/options` or an options file read using the `call` option, pppd uses full root privileges when opening the device. Thus, by creating an appropriate file under `/etc/ppp/peers`, the system administrator can allow users

to establish a PPP connection via a device that they would not normally have access to. Otherwise pppd uses the invoking user's real UID when opening the device.

Authentication

During the authentication process, one peer convinces the other of its identity by sending its name and some secret information to the other. During authentication, the first peer becomes the "client" and the second becomes the "server." Authentication names can (but are not required to) correspond to the peer's Internet hostnames.

pppd supports four authentication protocols: the Password Authentication Protocol (PAP) and three forms of the Challenge Handshake Authentication Protocol (CHAP). With the PAP protocol, the client sends its name and a cleartext password to the server to authenticate itself. With CHAP, the server initiates the authentication exchange by sending a challenge to the client who must respond with its name and a hash value derived from the shared secret and the challenge.

The PPP protocol is symmetrical, meaning that each peer may be required to authenticate itself to the other. Different authentication protocols and names can be used for each exchange.

By default, pppd authenticates if requested and does not require authentication from the peer. However, pppd does not authenticate itself with a specific protocol if it has no secrets that can do so.

pppd stores authentication secrets in the `/etc/ppp/pap-secrets` (for PAP), and `/etc/ppp/chap-secrets` (for CHAP) files. Both files use the same format. pppd uses secrets files to authenticate itself to other systems and to authenticate other systems to itself.

Secrets files contain one secret per line. Secrets are specific to a particular combination of client and server and can only be used by that client to authenticate itself to that server. Each line in a secrets file has a minimum of three fields that contain the client and server names followed by the secret. Often, these three fields are followed by IP addresses that are used by clients to connect to a server.

A secrets file is parsed into words, with client name, server name and secrets fields allocated one word each. Embedded spaces or other special characters within a word must be quoted or escaped. Case is significant in all three fields.

A secret beginning with an at sign ("@") is followed by the name of a file containing the secret. An asterisk (*) as the client or server name matches any name. When choosing a match, pppd selects the one with the fewest wildcards. Succeeding words on a line are interpreted by pppd as acceptable IP addresses for that client. IP Addresses are disallowed if they appear in lines that contain only three words or lines whose first word begins with a hyphen ("-"). To allow any address, use "*". An address starting with an exclamation point ("!") indicates that the specified address is not acceptable. An address may be followed by "/" and a number *n* to indicate a whole subnet (all addresses that have the same value in the most significant *n* bits).

In this form, the address may be followed by a plus sign ("+") to indicate that one address from the subnet is authorized, based on the ppp network interface unit number in use. In this case, the host part of the address is set to the unit number, plus one.

When authenticating the peer, pppd chooses a secret with the peer's name in the first field of the secrets file and the name of the local system in the second field. The local system name defaults to the hostname, with the domain name appended if the `domain` option is used. The default can be overridden with the `name` option unless the `usehostname` option is used.

When authenticating to the peer, pppd first determines the name it will use to identify itself to the peer. This name is specified with the `user` option. If the `user` option is not used, the name defaults to the host name of the local system. pppd then selects a secret from the secrets file by searching for an entry with a local name in the first field and the peer's name in the second field. pppd will know the name of the peer if standard CHAP authentication is used because the peer will have sent it in the Challenge packet. However, if MS-CHAP or PAP is being used, pppd must determine the peer's name from the options specified by the user. The user can specify the peer's name directly with the `remotename` option. Otherwise, if the remote IP address was specified by a name, rather than in numeric form, that name will be used as the peer's name. If that fails, pppd uses the null string as the peer's name.

When authenticating the peer with PAP, the supplied password is compared with data in the secrets file. If the password and secret do not match, the password is encrypted using `crypt()` and checked against the secret again. If the `papcrypt` option is given, the first unencrypted comparison is omitted for better security, and entries must thus be in encrypted `crypt(3C)` form.

If the `login` option is specified, the username and password are also checked against the system password database. This allows you to set up the `pap-secrets` file to enable PPP access only to certain users, and to restrict the set of IP addresses available to users. Typically, when using the `login` option, the secret in `/etc/ppp/pap-secrets` would be "", which matches any password supplied by the peer. This makes having the same secret in two places unnecessary. When `login` is used, the `pam` option enables access control through `pam(3PAM)`.

Authentication must be completed before IPCP (or other network protocol) can be started. If the peer is required to authenticate itself and fails, pppd closes LCP and terminates the link. If IPCP negotiates an unacceptable IP address for the remote host, IPCP is closed. IP packets are sent or received only when IPCP is open.

To allow hosts that cannot authenticate themselves to connect and use one of a restricted set of IP addresses, add a line to the `pap-secrets` file specifying the empty string for the client name and secret.

Additional pppd options for a given peer may be specified by placing them at the end of the secrets entry, separated by two dashes (--). For example

```
peername servername secret ip-address -- novj
```

Routing

When IPCP negotiation is complete, pppd informs the kernel of the local and remote IP addresses for the PPP interface and creates a host route to the remote end of the link that enables peers to exchange IP packets. Communication with other machines generally requires further modification to routing tables and/or Address Resolution Protocol (ARP) tables. In most cases the default route and/or proxyarp options are sufficient for this, but further intervention may be necessary. If further intervention is required, use the `/etc/ppp/ip-up` script or a routing protocol daemon.

To add a default route through the remote host, use the `default route` option. This option is typically used for “client” systems; that is, end-nodes that use the PPP link for access to the general Internet.

In some cases it is desirable to use proxy ARP, for example on a server machine connected to a LAN, to allow other hosts to communicate with the remote host. `proxyarp` instructs pppd to look for a network interface on the same subnet as the remote host. That is, an interface supporting broadcast and ARP that is not a point-to-point or loopback interface and that is currently up. If found, pppd creates a permanent, published ARP entry with the IP address of the remote host and the hardware address of the network interface.

When the `demand` option is used, the interface IP addresses are already set at the time when IPCP comes up. If pppd cannot negotiate the same addresses it used to configure the interface, it changes the interface IP addresses to the negotiated addresses. This may disrupt existing connections. Using demand dialing with peers that perform dynamic IP address assignment is not recommended.

Scripts

pppd invokes scripts at various stages during processing that are used to perform site-specific ancillary processing. These scripts may be shell scripts or executable programs. pppd does not wait for the scripts to finish. The scripts are executed as root (with the real and effective user-id set to 0), enabling them to update routing tables, run privileged daemons, or perform other tasks. Be sure that the contents of these scripts do not compromise your system's security. pppd runs the scripts with standard input, output and error redirected to `/dev/null`, and with an environment that is empty except for some environment variables that give information about the link. The pppd environment variables are:

DEVICE	Name of the serial tty device.
IFNAME	Name of the network interface.
IPLOCAL	IP address for the link's local end. This is set only when IPCP has started.
IPREMOTE	IP address for the link's remote end. This is set only when IPCP has started.
PEERNAME	Authenticated name of the peer. This is set only if the peer authenticates itself.
SPEED	Baud rate of the tty device.
ORIG_UID	Real user-id of user who invoked pppd.

PPPLLOGNAME Username of the real user-id who invoked pppd. This is always set.

pppd also sets the following variables for the ip-down and auth-down scripts:

CONNECT_TIME Number of seconds between the start of PPP negotiation and connection termination.

BYTES_SENT Number of bytes sent at the level of the serial port during the connection.

BYTES_RCV Number of bytes received at the level of the serial port during the connection.

LINKNAME Logical name of the link, set with the linkname option.

If they exist, pppd invokes the following scripts. It is not an error if they do not exist.

- /etc/ppp/auth-up Program or script executed after the remote system successfully authenticates itself. It is executed with five command-line arguments: interface-name peer-name user-name tty-device speed. Note that this script is not executed if the peer does not authenticate itself, for example, when the noauth option is used.
- /etc/ppp/auth-down Program or script executed when the link goes down if /etc/ppp/auth-up was previously executed. It is executed in the same manner with the same parameters as /etc/ppp/auth-up.
- /etc/ppp/ip-up A program or script that is executed when the link is available for sending and receiving IP packets (that is, IPCP has come up). It is executed with six command-line arguments: interface-name tty-device speed local-IP-address remote-IP-address ipparam.
- /etc/ppp/ip-down A program or script which is executed when the link is no longer available for sending and receiving IP packets. This script can be used for undoing the effects of the /etc/ppp/ip-up script. It is invoked in the same manner and with the same parameters as the ip-up script.
- /etc/ppp/ipv6-up Similar to /etc/ppp/ip-up, except that it is executed when the link is available for sending and receiving IPv6 packets. Executed with six command-line arguments: interface-name tty-device speed local-link-local-address remote-link-local-address ipparam.
- /etc/ppp/ipv6-down Similar to /etc/ppp/ip-down, but executed when IPv6 packets can no longer be transmitted on the link. Executed with the same parameters as the ipv6-up script.

示例

示例 1 Using the auth Option

The following examples assume that the /etc/ppp/options file contains the auth option.

示例 1 Using the auth Option (续)

pppd is commonly used to dial out to an ISP. You can do this using the “pppd call isp” command where the /etc/ppp/peers/isp file is set up to contain a line similar to the following:

```
cua/a 19200 crtscts connect '/usr/bin/chat -f /etc/ppp/chat-isp' noauth
```

For this example, [chat\(1M\)](#) is used to dial the ISP's modem and process any login sequence required. The /etc/ppp/chat-isp file is used by chat and could contain the following:

```
ABORT "NO CARRIER"
ABORT "NO DIALTONE"
ABORT "ERROR"
ABORT "NO ANSWER"
ABORT "BUSY"
ABORT "Username/Password Incorrect"
"" "at"
OK "at&f&d2&c1"
OK "atdt2468135"
"name:" "^Umyuserid"
"word:" "\qmypassword"
"ispts" "\q^Uppp"
"--^Uppp--"
```

See the [chat\(1M\)](#) man page for details of chat scripts.

示例 2 Using pppd with proxyarp

pppd can also provide a dial-in ppp service for users. If the users already have login accounts, the simplest way to set up the ppp service is to let the users log in to their accounts and run pppd as shown in the following example:

```
example% pppd proxyarp
```

示例 3 Providing a User with Access to PPP Facilities

To provide a user with access to the PPP facilities, allocate an IP address for the user's machine, create an entry in /etc/ppp/pap-secrets or /etc/ppp/chap-secrets. This enables the user's machine to authenticate itself. For example, to enable user “Joe” using machine “joespc” to dial in to machine “server” and use the IP address “joespc.my.net,” add the following entry to the /etc/ppp/pap-secrets or /etc/ppp/chap-secrets files:

```
joespc      server      "joe's secret"      joespc.my.net
```

Alternatively, you can create another username, for example “ppp,” whose login shell is /usr/bin/pppd and whose home directory is /etc/ppp. If you run pppd this way, add the options to the /etc/ppp/.ppprc file.

示例 3 Providing a User with Access to PPP Facilities (续)

If your serial connection is complex, it may be useful to escape such control characters as XON (^Q) and XOFF (^S), using `asynmap a0000`. If the path includes a `telnet`, escape `^`] (`asynmap 200a0000`). If the path includes a `rlogin` command, add `escape ff` option to the options, because `rlogin` removes the window-size-change sequence [0xff, 0xff, 0x73, 0x73, followed by any 8 bytes] from the stream.

退出状态

The `pppd` exit status indicates errors or specifies why a link was terminated. Exit status values are:

- 0 `pppd` has detached or the connection was successfully established and terminated at the peer's request.
- 1 An immediately fatal error occurred. For example, an essential system call failed.
- 2 An error was detected in the options given. For example, two mutually exclusive options were used, or `/etc/ppp/options` is missing and the user is not root.
- 3 `pppd` is not `setuid - root` and the invoking user is not root.
- 4 The kernel does not support PPP. For example, the PPP kernel driver is not included or cannot be loaded.
- 5 `pppd` terminated because it was sent a SIGINT, SIGTERM or SIGHUP signal.
- 6 The serial port could not be locked.
- 7 The serial port could not be opened.
- 8 The connect script failed and returned a non-zero exit status.
- 9 The command specified as the argument to the `pty` option could not be run.
- 10 The PPP negotiation failed because no network protocols were able to run.
- 11 The peer system failed or refused to authenticate itself.
- 12 The link was established successfully, but terminated because it was idle.
- 13 The link was established successfully, but terminated because the connect time limit was reached.
- 14 Callback was negotiated and an incoming call should arrive shortly.
- 15 The link was terminated because the peer is not responding to echo requests.
- 16 The link was terminated by the modem hanging up.
- 17 The PPP negotiation failed because serial loopback was detected.
- 18 The init script failed because a non-zero exit status was returned.
- 19 Authentication to the peer failed.

文件	<code>/var/run/spppn.pid</code>	Process-ID for pppd process on PPP interface unit <i>n</i> .
	<code>/var/run/ppp-name.pid</code>	Process-ID for pppd process for logical link name (see the <code>linkname</code> option).
	<code>/etc/ppp/pap-secrets</code>	Username, passwords and IP addresses for PAP authentication. This file should be owned by root and not readable or writable by any other user, otherwise pppd will log a warning.
	<code>/etc/ppp/chap-secrets</code>	Names, secrets and IP addresses for all forms of CHAP authentication. The <code>/etc/ppp/pap-secrets</code> file should be owned by root should not readable or writable by any other user, otherwise, pppd will log a warning.
	<code>/etc/ppp/options</code>	System default options for pppd, read before user default options or command-line options.
	<code>\$HOME/.ppprc</code>	User default options, read before <code>/etc/ppp/options.ttyname</code> .
	<code>/etc/ppp/options.ttyname</code>	System default options for the serial port in use; read after <code>\$HOME/.ppprc</code> . The <code>ttyname</code> component of this filename is formed when the initial <code>/dev/</code> is stripped from the port name (if present), and slashes (if any) are converted to dots.
	<code>/etc/ppp/peers</code>	Directory with options files that may contain privileged options, even if pppd was invoked by a user other than root. The system administrator can create options files in this directory to permit non-privileged users to dial out without requiring the peer to authenticate, but only to certain trusted peers.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/ppp
Interface Stability	Committed

另请参见 [chat\(1M\)](#), [ifconfig\(1M\)](#), [crypt\(3C\)](#), [pam\(3PAM\)](#), [attributes\(5\)](#)

Haskin, D., Allen, E. RFC 2472 – IP Version 6 Over PPP. Network Working Group. December 1998.

Jacobson, V. RFC 1144, Compressing TCP/IP Headers for Low-Speed Serial Links. Network Working Group. February, 1990

Lloyd, B., Simpson, W. RFC 1334, PPP Authentication Protocols. Network Working Group. October 1992.

McGregor, G. RFC 1332, The PPP Internet Protocol Control Protocol (IPCP). Network Working Group. May 1992.

Rivest, R. RFC 1321, The MD5 Message-Digest Algorithm. Network Working Group. April 1992

Simpson, W. RFC 1661, The Point-to-Point Protocol (PPP). Network Working Group. July 1994.

Simpson, W. RFC 1662, HDLC-like Framing . Network Working Group. July 1994.

附注

These signals affect pppd behavior:

SIGINT , SIGTERM	Terminate the link, restore the serial device settings and exit.
SIGHUP	Terminate the link, restore the serial device settings and close the serial device. If the <code>persist</code> or <code>demand</code> option is specified, pppd attempts to reopen the serial device and start another connection after the holdoff period. Otherwise pppd exits. If received during the holdoff period, SIGHUP causes pppd to end the holdoff period immediately.
SIGUSR1	Toggles the state of the <code>debug</code> option and prints link status information to the log.
SIGUSR2	Causes pppd to renegotiate compression. This is useful to re-enable compression after it has been disabled as a result of a fatal decompression error. (Fatal decompression errors generally indicate a bug in an implementation.)

诊断

Messages are sent to the syslog daemon using facility `LOG_DAEMON`. To see error and debug messages, edit the `/etc/syslog.conf` file to direct the messages to the desired output device or file, or use the `updetach` or `logfile` options.

The `debug` option causes the contents of all LCP, PAP, CHAP or IPCP control packets sent or received to be logged. This is useful if PPP negotiation does not succeed or if authentication fails.

Debugging can also be enabled or disabled by sending a `SIGUSR1` signal, which acts as a toggle to the pppd process.

引用名	pppoe – PPPoE chat utility
用法概要	<pre>pppoe [-omillsecs] [-smillsecs] [-v] device [service [[except]server... [only]]]</pre> <pre>pppoe [-omillsecs] [-v] -i [device]</pre>
描述	The pppoe utility implements the client-side negotiation of PPPoE. It is intended to be used with the pppd(1M) connect option, in the same manner as the chat(1M) utility is used for asynchronous dial-up PPP. When given with the <code>-i</code> flag, pppoe sends out a broadcast query on the given interface named by the <i>device</i> parameter. You can specify no other arguments in this mode. All responding PPPoE servers and the offered services are displayed on standard output. Otherwise, when given without the <code>-i</code> flag, pppoe does the full PPPoE client-side negotiation. The <i>device</i> parameter is the intended Ethernet interface, and must already be plumbed with sppptun(1M). The optional <i>service</i> parameter specifies a particular service desired; other offered services will be ignored. The optional <i>server</i> parameter specifies a specific server desired. You can specify <i>server</i> as an Ethernet address in the usual x:x:x:x:x:x format (with "*" in any of the six byte positions interpreted to mean "any"), or as a symbolic name resolved through <code>/etc/ethers</code> (or <code>NIS</code>), or as a PPPoE access concentrator name. The sense of the match (true or false) can be inverted by specifying the keyword <code>except</code> before this string. This parameter can be specified more than once, and the first match is taken. If you specify the <i>server</i> parameter, then the selected servers become "preferred." If no preferred server responds, then the first responding server is used instead. To exclude non-matching servers entirely, append the keyword <code>only</code>.
选项	The following options are supported: <code>-i</code> Sends out broadcast query over interface specified by <i>device</i>. <code>-o</code> Sets the initial wait time in milliseconds for PADO from the server before PADI is retried. The default is 500 milliseconds for normal operation, or 3000 milliseconds (3 seconds) for inquiry (<code>-i</code>) mode. <code>-s</code> Sets the initial wait time in milliseconds for PADS from the server before PADR is retried. The default is 2000 milliseconds (2 seconds). <code>-v</code> Displays verbose progress messages, including all PPPoE messages sent, and all state machine transitions. You normally do not need to adjust the parameters set with <code>-o</code> and <code>-s</code>. They are provided for coping with unusually slow servers.
操作数	The following operands are supported: <i>device</i> plumbed Ethernet interface

server preferred server or, if you specify only, the specified server
service desired service; other available services are ignored

示例

示例 1 Connecting to Any Service on hme0

The following command enables you to connect to any PPPoE service on hme0:

```
# /usr/bin/pppd sppptun plugin pppoe.so \  
connect "/usr/lib/inet/pppoe hme0" debug
```

Often, a command such as the preceding is specified in an `/etc/ppp/peers` file instead. For example, enter the following in `/etc/ppp/peers/myisp`:

```
sppptun  
plugin pppoe.so  
connect "/usr/lib/inet/pppoe hme0"  
debug
```

To invoke the PPP connection described in the file, enter:

```
% /usr/bin/pppd call myisp
```

Note that, because the `/etc/ppp/peers` files are considered privileged by `pppd`, you need not be root to invoke the preceding command.

示例 2 Connecting to a Particular Service

A more complex example: on hme0, connect to only the internet service offered by PPPoE servers with access concentrator name `isp`, but not to any Ethernet addresses starting with `40:0:1a`.

```
# /usr/lib/inet/pppoe hme0 internet except 40:0:1a:*:*:* isp only
```

Note that the `except 40:0:1a:*:*:*` filter must come before `isp`, because the filters are first-match.

退出状态

The following exit values are returned:

```
0 Successful completion.  
>0 An error occurred.
```

文件

<code>/usr/lib/inet/pppoe</code>	executable command
<code>/dev/sppptun</code>	Solaris PPP tunneling device driver.
<code>/etc/ppp/connect-errors</code>	usual location of error output (see DIAGNOSTICS , below)

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/ppp/tunnel

另请参见 [pppd\(1M\)](#), [sppptun\(1M\)](#), [pppoed\(1M\)](#), [sppptun\(7M\)](#)

RFC 2516, Method for Transmitting PPP Over Ethernet (PPPoE), Mamakos et al, February 1999

诊断

Error messages are written to standard error, which is normally redirected by pppd to /etc/ppp/connect-errors. The errors can also be redirected to pppd's standard output by using the updetach option.

If you specify the -v, verbose progress messages are displayed, including all PPPoE messages sent, and all state machine transitions. Specifying the updetach or nodetach pppd option is helpful when using verbose mode.

引用名 pppoed – PPPoE server daemon

用法概要 ppoed [*options*]

描述 The pppoed daemon implements the server-side negotiation of PPPoE. When a client requests service from this daemon, a copy of [pppd\(1M\)](#) is invoked to handle the actual PPP communication.

At startup, options are read from the command line and the `/etc/ppp/pppoe` file. After these options have been read, options in the per-device `/etc/ppp/pppoe.device` files are read, using the device names specified on the command line or in `/etc/ppp/pppoe`. Device names are not permitted in the per-device files. It is not an error if any of these files are absent; missing files are ignored.

Options are reread in the same order on SIGHUP. Except for the possibility of short delays due to the processing time, SIGHUP does not interfere with any client operations. Current status, including options read, is dumped to `/tmp/pppoed.pid` on SIGINT.

The options are used to set up a list of services to be offered to PPPoE clients on the broadcast domains (Ethernet subnets) specified by the named devices. Option parsing is always in one of two modes, either global mode or service mode. The initial mode at the beginning of each file (and the command line) is global mode. Options specified in global mode serve as default values for subsequently defined services. Service mode is entered by the `service name` option. In this mode, the named option is defined. Options that appear in this mode override any global mode definitions for the current service.

The option parsing follows standard shell tokenizing rules, using whitespace to delimit tokens, quotes to enclose strings that can contain whitespace, and escape sequences for special characters. Environment variables are substituted using familiar `$VAR` and `${VAR}` syntax and set using `NEWVAR=string`. Variables are both usable in subsequent options and provided to the [pppd\(1M\)](#) processes spawned for each client, but they are interpreted as they are encountered during option processing. Thus, all set variables are seen by all processes spawned; position in the configuration files has no effect on this.

选项 The pppoed daemon supports the following options:

`client [except] client-list` This option restricts the clients that may receive the service. If the `except` keyword is given, then the clients on the list cannot access the service, but others can. If this keyword is not given, then only the listed clients can access the service.

This option can be specified more than once for a given service. For a given client, first match among all listed options encountered specifies the handling. If it matches an option with `except` specified, then access is denied. Otherwise, it is granted. The `client` list within a service is prepended to any list specified in the global context.

If no `client` options are given or if all options are specified with `except`, then all clients are permitted by default. If any `client` options without `except` are specified, then no clients are permitted by default.

The *client-list* is a comma-separated list of client identifiers. The match is made if any client on the list matches; thus, these are logically "ORed" together. Each client identifier can be either a symbolic name (resolved through `/etc/ethers` or NIS, as defined by `/etc/nsswitch.conf`) or a hexadecimal Ethernet address in the format `x:x:x:x:x:x`. In the latter case, any byte of the address can be "*", which matches any value in that position. For example, `40:0:1a:*:*:*` matches Ethernet adapters from the manufacturer assigned block `40:0:1a`.

<code>debug</code>	Increase debug logging detail level by one. The detail levels are 0 (no logging), 1 (errors only; the default), 2 (warnings), 3 (informational messages), and 4 (debug messages). Log messages are written by default to <code>syslog(3C)</code> using facility <i>daemon</i> (see the <code>log</code> option below). When specified on the command line or in the global context of the <code>/etc/ppp/pppoe</code> file, this option also sets the daemon's default (non-service-related) detail level.
<code>device device-list</code>	Specify the devices on which the service is available. The <i>device-list</i> is a comma-separated list of logical device names (without the leading <code>/dev/</code>), such as <code>hme0</code> . This option is ignored if encountered in the per-device <code>/etc/ppp/pppoe.device</code> files.
<code>extra string</code>	Specifies extra options to <code>pppd(1M)</code> . It defaults to " <code>plugin pppoe.so directtty</code> " and usually does not need to be overridden.
<code>file path</code>	Suspends parsing of the current file, returns to global mode, and reads options from <i>path</i> . This file must be present and readable; if it is not, an error is logged. When the end of that file is reached, processing returns to the current file and the mode is reset to global again.
	The global mode options specified in files read by this command use the options set in the current file's global mode; this condition extends to any file included by those files. All files read are parsed as though the command line had specified this option, and thus inherit the command line's global modes.

	This option can be used to revert to global mode at any point in an option file by specifying <code>file /dev/null</code> .
<code>group <i>name</i></code>	Specifies the group ID (symbolic or numeric) under which <code>pppd</code> is executed. If <code>pppoed</code> is not run as root, this option is ignored.
<code>log <i>path</i></code>	Specifies an alternate debug logging file. Debug messages are sent to this file instead of <code>syslog</code> . The special name <code>syslog</code> is recognized to switch logging back to <code>syslog</code> . When specified on the command line or in the global context of the <code>/etc/ppp/pppoe</code> file, this option also sets the daemon's default (non-service-related) log file.
<code>nodebug</code>	Set debug logging detail level to 0 (no logging). When specified on the command line or in the global context of the <code>/etc/ppp/pppoe</code> file, this option also sets the daemon's default (non-service-related) detail level.
<code>nowildcard</code>	Specifies that the current service should not be included in response to clients requesting "any" service. The client must ask for this service by name. When specified on the command line or in the global context of the <code>/etc/ppp/pppoe</code> file, this option causes <code>pppoed</code> to ignore all wildcard service requests.
<code>path <i>path</i></code>	Specifies the path to the <code>pppd</code> executable. Defaults to <code>/usr/bin/pppd</code> .
<code>pppd <i>string</i></code>	Passes command-line arguments to <code>pppd</code> . It can be used to set the IP addresses or configure security for the session. The default value is the empty string.
<code>server <i>string</i></code>	Specifies the PPPoE Access Concentrator name to be sent to the client. It defaults to "Solaris PPPoE".
<code>service <i>name</i></code>	Closes any service being defined and begins definition of a new service. The same service name can be used without conflict on multiple devices. If the same service name is used on a single device, then the last definition encountered during parsing overrides all previous definitions.
<code>user <i>name</i></code>	Specifies the user ID, symbolic or numeric, under which <code>pppd</code> is executed. If <code>pppoed</code> is not run as root, this option is ignored.
<code>wildcard</code>	Specifies that the service should be included in responses to client queries that request "any" service, which is done by requesting a service name of length zero. When specified on the command line or in the global context of the

/etc/ppp/pppoe file, this option causes pppoed to ignore all wildcard service requests. This is the default.

示例

示例1 Configuring for Particular Services

In the /etc/ppp/pppoe file:

```
service internet
    device $DEV
    pppd "proxyarp 192.168.1.1:"
service debugging
    device hme0,$DEV
    pppd "debug proxyarp 192.168.1.1:"
```

You then invoke the daemon with:

```
example% /usr/lib/inet/pppoed DEV=eri0
```

The lines in /etc/ppp/pppoe and the preceding command result in offering services "internet" and "debugging" (and responding to wildcard queries) on interface eri0, and offering only service "debugging" on interface hme0.

Signals

The pppoed daemon responds to the following signals:

- SIGHUP Causes pppoed to reparse the original command line and all configuration files, and close and reopen any log files.
- SIGINT Causes a snapshot of the state of the pppoed daemon to be written to /tmp/pppoed.*pid* (where *pid* is the decimal process ID of the daemon).

文件

/usr/lib/inet/pppoed	executable command
/dev/sppptun	Solaris PPP tunneling device driver
/etc/ppp/pppoe	main configuration option file
/etc/ppp/pppoe. <i>device</i>	per-device configuration option file
/etc/ppp/pppoe-errors	location of output from pppd's stderr
/etc/ppp/pppoe.if	list of Ethernet interfaces to be plumbed at boot time
/tmp/pppoed. <i>pid</i>	ASCII text file containing dumped pppoed state information

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/ppp/tunnel

另请参见

[pppd\(1M\)](#), [pppoe\(1M\)](#), [sppptun\(1M\)](#), [sppptun\(7M\)](#)

Mamakos, L., et al. RFC 2516, A Method for Transmitting PPP Over Ethernet (PPPoE). Network Working Group. February 1999

附注

Because `pppd` is installed setuid root, this daemon need not be run as root. However, if it is not run as root, the `user` and `group` options are ignored.

The Ethernet interfaces to be used must be plumbed for PPPoE using the [sppptun\(1M\)](#) utility before services can be offered.

The daemon operate runs even if there are no services to offer. If you want to modify a configuration, it is not necessary to terminate the daemon. Simply use `kill -HUP pppoe` after updating the configuration files.

The PPPoE protocol is far from perfect. Because it runs directly over Ethernet, there is no possibility of security and the MTU is limited to 1492 (violating RFC 1661's default value of 1500). It is also not possible to run the client and the server of a given session on a single machine with a single Ethernet interface for testing purposes. The client and server portions of a single session must be run on separate Ethernet interfaces with different MAC addresses.

引用名	pppstats – print PPP statistics														
用法概要	pppstats [-a] [-v] [-r] [-z] [-c <count>] [-w <secs>] [interface]														
描述	The pppstats utility reports PPP-related statistics at regular intervals for the specified PPP interface. If the interface is unspecified, pppstats defaults to sppp0. The display is split horizontally into input and output sections containing columns of statistics describing the properties and volume of packets received and transmitted by the interface.														
选项	The pppstats options are: -a Display absolute values rather than deltas. With this option, all reports show statistics for the time elapsed since the link was initiated. Without this option, the second and subsequent reports show statistics for the time since the last report. -c count Repeat the display count times. If this option is not specified, the default repeat count is 1 if the -w option is not specified, otherwise infinity. -r Display additional statistics summarizing the compression ratio achieved by the packet compression algorithm in use. -v Display additional statistics relating to the performance of the Van Jacobson TCP header compression algorithm. -w wait Pause wait seconds between each display. If this option is not specified, the default interval is five seconds. -z Instead of the standard display, show statistics indicating the performance of the packet compression algorithm in use.														
扩展描述	The following fields are printed on the input side when the -z option is not used: <table><tr><td>IN</td><td>Total number of bytes received by this interface.</td></tr><tr><td>PACK</td><td>Total number of packets received by this interface.</td></tr><tr><td>VJCOMP</td><td>Number of header-compressed TCP packets received by this interface.</td></tr><tr><td>VJUNC</td><td>Number of header-uncompressed TCP packets received by this interface. Not reported when the -r option is specified.</td></tr><tr><td>VJERR</td><td>Number of corrupted or bogus header-compressed TCP packets received by this interface. Not reported when the -r option is specified.</td></tr><tr><td>VJT0SS</td><td>Number of VJ header-compressed TCP packets dropped on reception by this interface because of preceding errors. Only reported when the -v option is specified.</td></tr><tr><td>NON-VJ</td><td>Total number of non-TCP packets received by this interface. Only reported when the -v option is specified.</td></tr></table>	IN	Total number of bytes received by this interface.	PACK	Total number of packets received by this interface.	VJCOMP	Number of header-compressed TCP packets received by this interface.	VJUNC	Number of header-uncompressed TCP packets received by this interface. Not reported when the -r option is specified.	VJERR	Number of corrupted or bogus header-compressed TCP packets received by this interface. Not reported when the -r option is specified.	VJT0SS	Number of VJ header-compressed TCP packets dropped on reception by this interface because of preceding errors. Only reported when the -v option is specified.	NON-VJ	Total number of non-TCP packets received by this interface. Only reported when the -v option is specified.
IN	Total number of bytes received by this interface.														
PACK	Total number of packets received by this interface.														
VJCOMP	Number of header-compressed TCP packets received by this interface.														
VJUNC	Number of header-uncompressed TCP packets received by this interface. Not reported when the -r option is specified.														
VJERR	Number of corrupted or bogus header-compressed TCP packets received by this interface. Not reported when the -r option is specified.														
VJT0SS	Number of VJ header-compressed TCP packets dropped on reception by this interface because of preceding errors. Only reported when the -v option is specified.														
NON-VJ	Total number of non-TCP packets received by this interface. Only reported when the -v option is specified.														

RATIO	Compression ratio achieved for received packets by the packet compression scheme in use, defined as the uncompressed size divided by the compressed size. Only reported when the - r option is specified.
UBYTE	Total number of bytes received, after decompression of compressed packets. Only reported when the - r option is specified.

The following fields are printed on the output side:

OUT	Total number of bytes transmitted from this interface.
PACK	Total number of packets transmitted from this interface.
VJCOMP	Number of TCP packets transmitted from this interface with VJ-compressed TCP headers.
VJUNC	Number of TCP packets transmitted from this interface with VJ-uncompressed TCP headers. Not reported when the - r option is specified.
NON-VJ	Total number of non-TCP packets transmitted from this interface. Not reported when the - r option is specified.
VJSRCH	Number of searches for the cached header entry for a VJ header compressed TCP packet. Only reported when the - v option is specified.
VJMISS	Number of failed searches for the cached header entry for a VJ header compressed TCP packet. Only reported when the - v option is specified.
RATIO	Compression ratio achieved for transmitted packets by the packet compression scheme in use, defined as the size before compression divided by the compressed size. Only reported when the - r option is specified.
UBYTE	Total number of bytes to be transmitted before packet compression is applied. Only reported when the - r option is specified.

When the - z option is specified, pppstats displays the following fields relating to the packet compression algorithm currently in use. If packet compression is not in use, these fields display zeroes. The fields displayed on the input side are:

COMPRESSED BYTE	Number of bytes of compressed packets received.
COMPRESSED PACK	Number of compressed packets received.
INCOMPRESSIBLE BYTE	Number of bytes of incompressible packets (that is, those which were transmitted in uncompressed form) received.
INCOMPRESSIBLE PACK	Number of incompressible packets received.
COMP RATIO	Recent compression ratio for incoming packets, defined as the uncompressed size divided by the compressed size (including both compressible and incompressible packets).

The fields displayed on the output side are:

- COMPRESSED BYTE Number of bytes of compressed packets transmitted.
- COMPRESSED PACK Number of compressed packets transmitted.
- INCOMPRESSIBLE BYTE Number of bytes of incompressible packets received; that is, those that were transmitted by the peer in uncompressed form.
- INCOMPRESSIBLE PACK Number of incompressible packets transmitted.
- COMP RATIO Recent compression ratio for outgoing packets.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/ppp
Interface Stability	Committed

另请参见

[pppd\(1M\)](#), [attributes\(5\)](#)

引用名	<code>praudit</code> – print contents of an audit trail file
用法概要	<code>praudit [-lrsx] [-ddel] [filename]...</code>
描述	<code>praudit</code> reads the listed <i>filenames</i> (or standard input, if no <i>filename</i> is specified) and interprets the data as audit trail records as defined in audit.log(4) . By default, times, user and group IDs (UIDs and GIDs, respectively) are converted to their ASCII representation. Record type and event fields are converted to their ASCII representation. A maximum of 100 audit files can be specified on the command line.
选项	The following options are supported: <code>-ddel</code> Use <i>del</i> as the field delimiter instead of the default delimiter, which is the comma. If <i>del</i> has special meaning for the shell, it must be quoted. The maximum size of a delimiter is three characters. The delimiter is not meaningful and is not used when the <code>-x</code> option is specified. <code>-l</code> Print one line per record. <code>-r</code> Print records in their raw form. Times, UIDs, GIDs, record types, and events are displayed as integers. This option is useful when naming services are offline. The <code>-r</code> option and the <code>-s</code> option are exclusive. If both are used, a format usage error message is output. <code>-s</code> Display records in their short form. Numeric fields' ASCII equivalents are looked up by means of the sources specified in the <code>/etc/nsswitch.conf</code> file (see nsswitch.conf(4)). All numeric fields are converted to ASCII and then displayed. The short ASCII representations for the record type and event fields are used. This option and the <code>-r</code> option are exclusive. If both are used, a format usage error message is output. <code>-x</code> Print records in XML form. Tags are included in the output to identify tokens and fields within tokens. Output begins with a valid XML prolog, which includes identification of the DTD which can be used to parse the XML.
文件	<code>/etc/security/audit_event</code> Audit event definition and class mappings. <code>/etc/security/audit_class</code> Audit class definitions. <code>/usr/share/lib/xml/dtd</code> Directory containing the versioned DTD file referenced in XML output, for example, <code>adt_record.dtd.1</code>. <code>/usr/share/lib/xml/style</code> Directory containing the versioned XSL file referenced in XML output, for example, <code>adt_record.xsl.1</code>.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	See below.

The command stability is evolving. The output format is unstable.

另请参见 [getent\(1M\)](#), [getpwuid\(3C\)](#), [gethostbyaddr\(3NSL\)](#), [ethers\(3SOCKET\)](#),
[getipnodebyaddr\(3SOCKET\)](#), [audit.log\(4\)](#), [audit_class\(4\)](#), [audit_event\(4\)](#), [group\(4\)](#),
[nsswitch.conf\(4\)](#), [passwd\(4\)](#), [attributes\(5\)](#)

See the section on Solaris Auditing in [《Oracle Solaris 管理：安全服务》](#).

引用名	projadd – administer a new project on the system
用法概要	<pre>projadd [-n] [-f <i>filename</i>] [-p <i>projid</i> [-o]] [-c <i>comment</i>] [-U <i>user</i> [,<i>user</i>]...] [-G <i>group</i> [,<i>group</i>]...] [[-K <i>name</i> [=value [,value]...]...]] <i>project</i></pre>
描述	projadd adds a new project entry to the /etc/project file. If the files backend is being used for the project database, the new project is available for use immediately upon the completion of the projadd command.
选项	The following options are supported: -c <i>comment</i> Add a project comment. Comments are stored in the project's entry in the /etc/project file. Generally, comments contain a short description of the project and are used as the field for the project's full name. Specify <i>comment</i> as a text string. <i>comment</i> cannot contain a colon (:) or NEWLINE. -f <i>filename</i> Specify the project file to modify. If no <i>filename</i> is specified, the system project file, /etc/project, is modified. -G <i>group</i>[,<i>group</i>...] Specify a group list for the project. -K <i>name</i>[=<i>value</i>[,<i>value</i>...] Specify an attribute list for the project. Multiple -K options can be specified to set values on multiple keys, such as: -K <i>key1</i>=<i>value1</i> -K "<i>key2</i>=(<i>value2a</i>), (<i>value2b</i>)" - Resource control attributes use parentheses to specify values for a key. Because many user shells interpret parentheses as special characters, it is best to enclose an argument to -K that contains parentheses with double quotes, as shown above and in EXAMPLES, below. See resource_controls(5) for a description of the resource controls you can specify for a project. -n Syntax check. Check the format of the existing system project file and modifications only. The contents of the existing project file, such as user names, group names, and resources that are specified in the project attributes are not checked. -o This option allows the project ID specified by the -p option to be non-unique within the project file. -p <i>projid</i> Set the project ID of the new project. Specify <i>projid</i> as a non-negative decimal integer below UID_MAX as defined in <code>limits.h</code>. <i>projid</i> defaults to the next available

unique number above the highest number currently assigned. For example, if *projids* 100, 105, and 200 are assigned, the next default *projid* is 201. *projids* between 0-99 are reserved by SunOS.

`-U user[,user...]` Specify a user list for the project.

操作数

The following operands are supported:

project The name of the project to create. The *project* operand is a string consisting of characters from the set of alphabetic characters, numeric characters, underline (`_`), and hyphen (`-`). The period (`.`) is reserved for projects with special meaning to the operating system. The first character of the project name must be a letter. An error message is displayed if these restrictions are not met.

示例

示例 1 Adding a Project

The following command creates the project `salesaudit` and sets the resource controls specified as arguments to the `-K` option.

```
projadd -p 111 -G sales,finance -c "Auditing Project" \
-K "rcap.max-rss=10GB" \
-K "process.max-file-size=(priv,50MB,deny)" \
-K "task.max-lwps=(priv,100,deny)" salesaudit
```

This command would produce the following entry in `/etc/project`:

```
salesaudit:111:Auditing Project::sales,finance: \
process.max-file-size=(priv,52428800,deny); \
rcap.max-rss=10737418240;task.max-lwps=(priv,100,deny)
```

Note that the preceding would appear as one line in `/etc/project`.

Comparing the `projadd` command and resulting output in `/etc/project`, note the effect of the scaling factor in the resource cap (`rcap.max-rss=10GB`) and the resource control (`process.max-file-size=(priv,50MB,deny)`). Modifiers, such as B, KB, and MB, and scaling factors are specified in [resource_controls\(5\)](#).

退出状态

The following exit values are returned:

- 0 Successful completion.
- 2 The command syntax was invalid. A usage message for `projadd` is displayed.
- 3 An invalid argument was provided to an option.
- 4 The *projid* given with the `-p` option is already in use.
- 5 The project files contain an error. See [project\(4\)](#).
- 6 The project to be added, group, user, or resource does not exist.

文件
属性

- 9 The project is already in use.
- 10 Cannot update the /etc/project file.
- /etc/project System project file
- See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	See below.

另请参见

Invocation is evolving. Human readable output is unstable.

[projects\(1\)](#), [groupadd\(1M\)](#), [groupdel\(1M\)](#), [groupmod\(1M\)](#), [grpck\(1M\)](#), [projdel\(1M\)](#), [projmod\(1M\)](#), [useradd\(1M\)](#), [userdel\(1M\)](#), [usermod\(1M\)](#), [project\(4\)](#), [attributes\(5\)](#), [resource_controls\(5\)](#)

附注

In case of an error, projadd prints an error message and exits with a non-zero status.

projadd adds a project definition only on the local system. If a network name service such as NIS or LDAP is being used to supplement the local /etc/project file with additional entries, projadd cannot change information supplied by the network name service.

引用名	projdel – delete a project from the system
用法概要	projdel [-f <i>filename</i>] <i>project</i>
描述	The projdel utility deletes a project from the system and makes the appropriate changes to the system file.
选项	The following options are supported: -f <i>filename</i> Specify the project file to modify. If no <i>filename</i> is specified, the system project file, /etc/project, is modified.
操作数	The following operands are supported: <i>project</i> The name of the project to be deleted.
退出状态	The following exit values are returned: 0 Successful completion. 2 The command syntax was invalid. A usage message for projdel is displayed. 3 An invalid argument was provided to an option. 4 The <i>projid</i> given with the -p option is already in use. 5 The project files contain an error. See project(4). 6 The project to be modified, group, user, or resource does not exist. 9 The project is already in use. 10 Cannot update the /etc/project file.
文件	/etc/project System project file
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface stability	See below.

Invocation is evolving. Human readable output is unstable.

另请参见 [projects\(1\)](#), [groupadd\(1M\)](#), [groupdel\(1M\)](#), [groupmod\(1M\)](#), [grpck\(1M\)](#), [logins\(1M\)](#), [projadd\(1M\)](#), [projmod\(1M\)](#), [useradd\(1M\)](#), [userdel\(1M\)](#), [usermod\(1M\)](#), [project\(4\)](#), [attributes\(5\)](#)

诊断

In case of an error, `projdel` prints an error message and exits with a non-zero status.

附注

`projdel` deletes a project definition only on the local system. If a network name service such as NIS or LDAP is being used to supplement the local `/etc/project` file with additional entries, `projdel` cannot change information supplied by the network name service.

引用名	projmod – modify a project's information on the system												
用法概要	<pre>projmod [-n] [-A -f filename -] projmod [-n] [-A -f filename -] [-p projid [-o]] [-c comment] [-a -s -r] [-U user [,user]...] [-G group [,group]...] [[-K name [=value [,value]...]...]] [-l new_projectname] project</pre>												
描述	The projmod utility modifies a project's definition on the system. projmod changes the definition of the specified project and makes the appropriate project-related system file and file system changes.												
选项	The following options are supported: <table><tr><td>-A</td><td>Apply the project's resource controls, as defined in the system's project database, to the project if it is active.</td></tr><tr><td>-a</td><td>Specify that the users, groups, attributes, or attribute values specified by the -U, -G or -K options should be added to the project, rather than replacing the existing member or attribute list.</td></tr><tr><td>-c comment</td><td>Specify <i>comment</i> as a text string. Generally, <i>comment</i> contains a short description of the project. This information is stored in the project's /etc/project entry.</td></tr><tr><td>-f filename -</td><td>Specify the project file to modify or validate or specify input from stdin for validation. As noted under OPERANDS, if you do not specify a project in a projmod command line, projmod validates the argument to -f. If you do not use this option, the system project file, /etc/project, is modified.</td></tr><tr><td>-G group [,group...]</td><td>Specify a replacement list of member groups of the project. When used in conjunction with the -a or -r options, this option specifies a list of groups to be added or removed from the project.</td></tr><tr><td>-K name[=value[,value...]]</td><td>Specify a replacement list of project attributes for the project. When used in conjunction with the -a, -r, or -s options, this option specifies a list of attribute values to be added, removed, or replaced in the project. Attributes must be delimited by semicolons (;). Multiple -K options can be specified to set, add, remove, or substitute values on multiple keys, such as: -K key1=value1 -K "key2=(value2a) ,(value2b)"</td></tr></table> Resource control attributes use parentheses to specify values for a key. Because many user shells interpret parentheses as special	-A	Apply the project's resource controls, as defined in the system's project database, to the project if it is active.	-a	Specify that the users, groups, attributes, or attribute values specified by the -U, -G or -K options should be added to the project, rather than replacing the existing member or attribute list.	-c comment	Specify <i>comment</i> as a text string. Generally, <i>comment</i> contains a short description of the project. This information is stored in the project's /etc/project entry.	-f filename -	Specify the project file to modify or validate or specify input from stdin for validation. As noted under OPERANDS, if you do not specify a project in a projmod command line, projmod validates the argument to -f. If you do not use this option, the system project file, /etc/project, is modified.	-G group [,group...]	Specify a replacement list of member groups of the project. When used in conjunction with the -a or -r options, this option specifies a list of groups to be added or removed from the project.	-K name[=value[,value...]]	Specify a replacement list of project attributes for the project. When used in conjunction with the -a, -r, or -s options, this option specifies a list of attribute values to be added, removed, or replaced in the project. Attributes must be delimited by semicolons (;). Multiple -K options can be specified to set, add, remove, or substitute values on multiple keys, such as: -K key1=value1 -K "key2=(value2a) ,(value2b)"
-A	Apply the project's resource controls, as defined in the system's project database, to the project if it is active.												
-a	Specify that the users, groups, attributes, or attribute values specified by the -U, -G or -K options should be added to the project, rather than replacing the existing member or attribute list.												
-c comment	Specify <i>comment</i> as a text string. Generally, <i>comment</i> contains a short description of the project. This information is stored in the project's /etc/project entry.												
-f filename -	Specify the project file to modify or validate or specify input from stdin for validation. As noted under OPERANDS, if you do not specify a project in a projmod command line, projmod validates the argument to -f. If you do not use this option, the system project file, /etc/project, is modified.												
-G group [,group...]	Specify a replacement list of member groups of the project. When used in conjunction with the -a or -r options, this option specifies a list of groups to be added or removed from the project.												
-K name[=value[,value...]]	Specify a replacement list of project attributes for the project. When used in conjunction with the -a, -r, or -s options, this option specifies a list of attribute values to be added, removed, or replaced in the project. Attributes must be delimited by semicolons (;). Multiple -K options can be specified to set, add, remove, or substitute values on multiple keys, such as: -K key1=value1 -K "key2=(value2a) ,(value2b)"												

characters, it is best to enclose an argument to `-K` that contains parentheses with double quotes, as shown above and in EXAMPLES, below. See [resource_controls\(5\)](#) for a description of the resource controls you can specify for a project.

- `-l new_projectname` Specify the new project name for the project. The *new_projectname* argument is a string consisting of characters from the set of alphabetic characters, numeric characters, period (.), underline (_), and hyphen (-). The first character should be alphabetic. An error message is written if these restrictions are not met. The project name must also be unique within the project file.
- `-n` Syntax check. Check the format of the existing system project file and modifications only. The contents of the existing project file, such as user names, group names, and resources that are specified in the project attributes are not checked.
- `-o` This option allows the project ID specified by the `-p` option to be non-unique within the project file.
- `-p projid` Specify a new project ID for the project. It must be a non-negative decimal integer less than MAXUID as defined in `param.h`. This value must be unique within the project file if the `-o` option is not specified.
- `-r` Specify that the users, groups, attributes, or attribute values specified by the `-U`, `-G` or `-K` options should be removed from the project, rather than replacing the existing member or attribute list.
- `-s` Specify that the list of attributes specified by the `-K` option should have their values replaced. If the attributes do not exist, they are added as if the `a` option was used. This option has no effect the `-U` or `-G` options.
- `-U user [,user...]` Specify a replacement list of member users of the project. When used in conjunction with the `-a` or `-r` options, this option specifies a list of users to be added or removed from the project.

操作数

The following operands are supported:

- project* An existing project name to be modified or displayed.
- (*none*) If no operand is given, the project file is validated without modifying any project.

示例

示例 1 Using the -K Option for Addition of an Attribute Value

Consider the following `project(4)` entry:

```
salesaudit:111:Auditing Project::sales,finance: \
  process.max-file-size=(priv,52428800,deny); \
  task.max-lwps=(priv,100,deny)
```

The preceding would appear as one line in `/etc/project`. For this and the following examples, the focus is on the attributes field in the `project` entry. That is, the last field, the field following the last semicolon.

The attributes field for the project `salesaudit` lists the following resource control:

```
task.max-lwps=(priv,1000,signal=KILL)
```

The following `projmod` command adds an action clause to the preceding entry:

```
# projmod -a -K "task.max-lwps=(priv,100,deny)" salesaudit
```

...with the resulting attributes field in the entry for `salesaudit`:

```
task.max-lwps=(priv,100,deny),(priv,1000,signal=KILL)
```

示例 2 Using the -K Option for the Substitution of an Attribute Value

Assume an attributes field in a `project(4)` entry for the project `salesaudit` that lists the following resource control:

```
task.max-lwps=(priv,100,deny),(priv,1000,signal=KILL)
```

The following `projmod` command substitutes the action clause specified in the command for the action clauses in the preceding entry:

```
# projmod -s -K "task.max-lwps=(priv,500,signal=SIGSTOP)" salesaudit
```

...with the resulting attributes field in the entry for `salesaudit`:

```
task.max-lwps=(priv,500,signal=SIGSTOP)
```

示例 3 Using the -K Option for Removal of an Attribute Value

Assume an attributes field in a `project(4)` entry for a project `salesaudit` that lists the following resource control:

```
task.max-lwps=(priv,100,deny),(priv,1000,signal=KILL)
```

The following `projmod` command removes the first action clause from the preceding entry:

```
# projmod -r -K "task.max-lwps=(priv,100,deny)" salesaudit
```

...with the resulting attributes field in the entry for `salesaudit`:

示例 3 Using the -K Option for Removal of an Attribute Value (续)

```
task.max-lwps=(priv,1000,signal=KILL)
```

示例 4 Specifying Multiple Attribute Values

Suppose you want to achieve the following resource controls for the project salesaudit:

```
task.max-lwps=(priv,100,deny)
process.max-file-size=(priv,50MB,deny)
```

The following projmod command adds these resource controls for salesaudit:

```
# projmod -a -K "task.max-lwps=(priv,100,deny)" \
-K "process.max-file-size=(priv,50MB,deny)" salesaudit
```

...with the resulting attributes field in the entry for salesaudit:

```
task.max-lwps=(priv,100,deny);process.max-file-size=(priv,52428800,deny)
```

In this example, note the effect of the use of the modifier and scaling factor for the resource control process.max-file-size. The specification in projmod:

```
"process.max-file-size=(priv,50MB,deny)"
```

...becomes, in /etc/project:

```
process.max-file-size=(priv,52428800,deny)
```

That is, 50MB is expanded to 52428800. The modifiers, such as MB, and scaling factors you can use for resource controls are specified in [resource_controls\(5\)](#).

示例 5 Binding a Pool to a Project

The following command sets the project.pool attribute for the project sales.

```
# projmod -a -K project.pool=salespool sales
```

示例 6 Evaluating Input from stdin

The following command uses the -f option without a project name operand to evaluate the contents of an NIS projects map.

```
# ypcat project | projmod -f -
```

退出状态

In case of an error, projmod prints an error message and exits with one of the following values:

The following exit values are returned:

- 0 Successful completion.
- 2 The command syntax was invalid. A usage message for projmod is displayed.

- 3 An invalid argument was provided to an option.
- 4 The *projid* given with the -p option is already in use.
- 5 The project files contain an error. See [project\(4\)](#).
- 6 The project to be modified, group, user, or resource does not exist.
- 9 The project is already in use.
- 10 Cannot update the /etc/project file.

文件

- /etc/group System file containing group definitions
- /etc/project System project file
- /etc/passwd System password file
- /etc/shadow System file containing users' encrypted passwords and related information

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	See below.

Invocation is evolving. Human readable output is unstable.

另请参见

[groupadd\(1M\)](#), [groupdel\(1M\)](#), [groupmod\(1M\)](#), [projadd\(1M\)](#), [projdel\(1M\)](#), [useradd\(1M\)](#), [userdel\(1M\)](#), [usermod\(1M\)](#), [passwd\(4\)](#), [project\(4\)](#), [attributes\(5\)](#), [resource_controls\(5\)](#)

附注

The projmod utility modifies project definitions only in the local /etc/project file. If a network name service such as NIS or LDAP is being used to supplement the local files with additional entries, projmod cannot change information supplied by the network name service. However projmod verifies the uniqueness of project name and project ID against the external name service.

引用名

prstat – report active process statistics

用法概要

```
prstat [-achJLmRrtTv] [-d u | d] [-C psrsetlist] [-h lgrplist]
      [-j projlist] [-k tasklist] [-n ntop[,nbottom]]
      [-p pidlist] [-P cpulist] [-s key | -S key ]
      [-u euidlist] [-U uidlist] [-z zoneidlist] [-Z]
      [interval [count]]
```

描述

The `prstat` utility iteratively examines all active processes on the system and reports statistics based on the selected output mode and sort order. `prstat` provides options to examine only processes matching specified PIDs, UIDs, zone IDs, CPU IDs, and processor set IDs.

The `-j`, `-k`, `-C`, `-p`, `-P`, `-u`, `-U`, and `-z` options accept lists as arguments. Items in a list can be either separated by commas or enclosed in quotes and separated by commas or spaces.

If you do not specify an option, `prstat` examines all processes and reports statistics sorted by CPU usage.

选项

The following options are supported:

`-a`

Report information about processes and users. In this mode `prstat` displays separate reports about processes and users at the same time.

`-C`

Print new reports below previous reports instead of overprinting them.

`-C psrsetlist`

Report only processes or lwps that are bound to processor sets in the given list. Each processor set is identified by an integer as reported by [psrset\(1M\)](#). The load averages displayed are the sum of the load averages of the specified processor sets (see [pset_getloadavg\(3C\)](#)). Processes with one or more LWPs bound to processor sets in the given list are reported even when the `-L` option is not used.

`-d u | d`

Specify `u` for a printed representation of the internal representation of time. See [time\(2\)](#). Specify `d` for standard date format. See [date\(1\)](#).

`-h lgrplist`

Report only processes or lwps whose home *lgroup* is in the given list of *lgroups*. No processes or lwps will be listed for invalid *lgroups*.

`-H`

Report information about home *lgroup*. In this mode, `prstat` adds an extra column showing process or lwps home *lgroup* with the header LGRP.

`-j projlist`

Report only processes or lwps whose project ID is in the given list. Each project ID can be specified as either a project name or a numerical project ID. See [project\(4\)](#).

- J
Report information about processes and projects. In this mode `prstat` displays separate reports about processes and projects at the same time.
- k *tasklist*
Report only processes or lwps whose task ID is in *tasklist*.
- L
Report statistics for each light-weight process (LWP). By default, `prstat` reports only the number of LWPs for each process.
- m
Report microstate process accounting information. In addition to all fields listed in `-v` mode, this mode also includes the percentage of time the process has spent processing system traps, text page faults, data page faults, waiting for user locks and waiting for CPU (latency time).
- n *ntop*[,*nbottom*]
Restrict number of output lines. The *ntop* argument determines how many lines of process or lwp statistics are reported, and the *nbottom* argument determines how many lines of user, task, or projects statistics are reported if the `-a`, `-t`, `-T`, `-J` or `-Z` options are specified. By default, `prstat` displays as many lines of output that fit in a window or terminal. When you specify the `-c` option or direct the output to a file, the default values for *ntop* and *nbottom* are 15 and 5.
- p *pidlist*
Report only processes whose process ID is in the given list.
- P *cpulist*
Report only processes or lwps which have most recently executed on a CPU in the given list. Each CPU is identified by an integer as reported by `psrinfo(1M)`.
- R
Put `prstat` in the real time scheduling class. When this option is used, `prstat` is given priority over time-sharing and interactive processes. This option is available only for superuser.
- r
Disable lookups for user names and project names. (Note that this does not apply to lookups for the `-j`, `-u`, or `-U` options.)
- s *key*
Sort output lines (that is, processes, lwps, or users) by *key* in descending order. Only one *key* can be used as an argument.

There are five possible key values:

cpu
Sort by process CPU usage. This is the default.

pri
Sort by process priority.

rss
Sort by resident set size.

size
Sort by size of process image.

time
Sort by process execution time.

-S *key*
Sort output lines by *key* in ascending order. Possible *key* values are the same as for the **-s** option. See **-s**.

-t
Report total usage summary for each user. The summary includes the total number of processes or LWPs owned by the user, total size of process images, total resident set size, total cpu time, and percentages of recent cpu time and system memory.

-T
Report information about processes and tasks. In this mode `prstat` displays separate reports about processes and tasks at the same time.

-u *euidlist*
Report only processes whose effective user ID is in the given list. Each user ID may be specified as either a login name or a numerical user ID.

-U *uidlist*
Report only processes whose real user ID is in the given list. Each user ID may be specified as either a login name or a numerical user ID.

-v
Report verbose process usage. This output format includes the percentage of time the process has spent in user mode, in system mode, and sleeping. It also includes the number of voluntary and involuntary context switches, system calls and the number of signals received. Statistics that are not reported are marked with the **-** sign.

-z *zoneidlist*
Report only processes or LWPs whose zone ID is in the given list. Each zone ID can be specified as either a zone name or a numerical zone ID. See [zones\(5\)](#).

-Z
Report information about processes and zones. In this mode, `prstat` displays separate reports about processes and zones at the same time.

Output

The following list defines the column headings and the meanings of a `prstat` report:

PID
The process ID of the process.

USERNAME

The real user (login) name or real user ID.

SWAP

The sum of swap reservations of the associated processes for each user, project, task, or zone. This counts shared memory only once for each user, project, task, or zone. Swap is reserved when anonymous memory is allocated or files are mapped private. The value of swap is expressed in kilobytes (K), megabytes (M), or gigabytes (G).

RSS

The resident set size of the process (RSS), in kilobytes (K), megabytes (M), or gigabytes (G). The RSS value is an estimate provided by [proc\(4\)](#) that might underestimate the actual resident set size. Users who want to get more accurate usage information for capacity planning should use the -x option to [pmap\(1\)](#) instead.

STATE

The state of the process:

cpuN

Process is running on CPU *N*.

sleep

Sleeping: process is waiting for an event to complete.

wait

Waiting: process is waiting for CPU usage to drop to the CPU-caps enforced limits. See the description of CPU-caps in [resource_controls\(5\)](#).

run

Runnable: process in on run queue.

zombie

Zombie state: process terminated and parent not waiting.

stop

Process is stopped.

PRI

The priority of the process. Larger numbers mean higher priority.

NICE

Nice value used in priority computation. Only processes in certain scheduling classes have a nice value.

TIME

The cumulative execution time for the process.

CPU

The percentage of recent CPU time used by the process. If executing in a non-global zone and the pools facility is active, the percentage will be that of the processors in the processor set in use by the pool to which the zone is bound.

PROCESS

The name of the process (name of executed file).

LWPID

The lwp ID of the lwp being reported.

NLWP

The number of lwps in the process.

With the some options, in addition to a number of the column headings shown above, there are:

NPROC

Number of processes in a specified collection.

MEMORY

Percentage of memory used by a specified collection of processes.

The following columns are displayed when the `-v` or `-m` option is specified

USR

The percentage of time the process has spent in user mode.

SYS

The percentage of time the process has spent in system mode.

TRP

The percentage of time the process has spent in processing system traps.

TFL

The percentage of time the process has spent processing text page faults.

DFL

The percentage of time the process has spent processing data page faults.

LCK

The percentage of time the process has spent waiting for user locks.

SLP

The percentage of time the process has spent sleeping.

LAT

The percentage of time the process has spent waiting for CPU.

VCX

The number of voluntary context switches.

ICX

The number of involuntary context switches.

SCL

The number of system calls.

SIG
The number of signals received.

Under the -L option, one line is printed for each lwp in the process and some reporting fields show the values for the lwp, not the process.

The following column is displayed when the -H option is specified:

LGRP
The home *lgroup* of the process or lwp.

操作数 The following operands are supported:

count
Specifies the number of times that the statistics are repeated. By default, prstat reports statistics until a termination signal is received.

interval
Specifies the sampling interval in seconds; the default interval is 5 seconds.

示例 示例1 Reporting the Five Most Active Super-User Processes

The following command reports the five most active super-user processes running on CPU1 and CPU2:

example% prstat -u root -n 5 -P 1,2 1 1

PID	USERNAME	SWAP	RSS	STATE	PRI	NICE	TIME	CPU	PROCESS/LWP
306	root	3024K	1448K	sleep	58	0	0:00.00	0.3%	sendmail/1
102	root	1600K	592K	sleep	59	0	0:00.00	0.1%	in.rdisc/1
250	root	1000K	552K	sleep	58	0	0:00.00	0.0%	utmpd/1
288	root	1720K	1032K	sleep	58	0	0:00.00	0.0%	sac/1
1	root	744K	168K	sleep	58	0	0:00.00	0.0%	init/1
TOTAL:		25, load averages: 0.05, 0.08, 0.12							

示例2 Displaying Verbose Process Usage Information

The following command displays verbose process usage information about processes with lowest resident set sizes owned by users root and john.

example% prstat -S rss -n 5 -vc -u root,john

PID	USERNAME	USR	SYS	TRP	TFL	DFL	LCK	SLP	LAT	VCX	ICX	SCL	SIG	PROCESS/LWP
1	root	0.0	0.0	-	-	-	-	100	-	0	0	0	0	init/1
102	root	0.0	0.0	-	-	-	-	100	-	0	0	3	0	in.rdisc/1
250	root	0.0	0.0	-	-	-	-	100	-	0	0	0	0	utmpd/1
1185	john	0.0	0.0	-	-	-	-	100	-	0	0	0	0	csh/1
240	root	0.0	0.0	-	-	-	-	100	-	0	0	0	0	switchd/4
TOTAL:		71, load averages: 0.02, 0.04, 0.08												

引用名	prtconf – 打印系统配置
用法概要	<pre>/usr/sbin/prtconf [-V] [-F] [-x] [-bpv] [-acdDlPuv] [dev_path]</pre>
描述	prtconf 命令可打印系统配置信息。输出中包括内存总量以及以设备树格式显示的系统外围设备的配置。 如果在命令行上为可以获取设备路径的命令选项指定了设备路径，prtconf 将仅显示该设备节点的信息。
选项	所支持的选项如下： -a 显示命令行上指定设备直至设备树根节点的所有祖先设备节点。 -b 显示用于标识平台的固件设备树根属性。这些属性是 "name"、"compatible"、"banner-name" 和 "model"。 -c 显示以命令行上指定设备节点为根的设备子树，即，显示命令行上指定设备节点的所有子项。 -d 除了节点名称之外，还显示 PCI 和 PCI Express 设备的供应商 ID 和设备 ID。 -D 对于设备树中的每个系统外围设备，显示用于管理外围设备的设备驱动程序的名称。 -l 显示与设备节点关联的 /dev/chassis 位置。如果使用 -v，则暗指 -l。 -F 返回控制台帧缓冲区的设备路径名称（如果存在）。如果没有帧缓冲区，prtconf 将返回一个非零退出代码。此标志必须由其自己使用。它只返回控制台名称、帧缓冲区设备或非零退出代码。例如，如果 SUNW,Ultra-30 上的控制台帧缓冲区是 ffb，该命令将返回：/SUNW,ffb@1e,0:ffb0。此选项可用于为 /dev/fb 创建一个指向实际控制台设备的符号链接。 -p 显示由 SPARC 平台上的固件 (PROM) 或 x86 平台上的引导系统提供的设备树派生的信息。使用此选项显示的设备树信息是初始配置的快照，可能不会准确反映之后发生的重新配置事件。 -P 包含有关伪设备的信息。缺省情况下忽略有关伪设备的信息。 -u 与 -v 一起，显示供应商和管理列表中的设备列表属性的信息（如果有）。 -v 指定详细模式。 -V 显示（SPARC 平台上）取决于平台的 PROM 或（x86 平台上）引导系统的版本信息。此标志必须由其自己使用。输出的内容是字符串。字符串的格式是任意的，并与平台相关。 -x 报告此系统上的固件是否支持 64 位。某些现有平台可能需要进行固件升级以运行 64 位内核。如果操作不适用于此平台，或者固件支持 64 位，操作将以无提示方式退出，并返回代码零。如果操作适用于此平台，并且固件不支持 64

位，操作将在标准输出中显示一条描述性消息，同时退出并返回一个非零代码。硬件平台文档包含有关需要进行固件升级以运行 64 位内核的平台的更多信息。

此标志会覆盖所有其他标志，并且必须由其自己使用。

操作数 支持下列操作数：

dev_path 指向目标设备次要节点、设备子树节点或为其显示设备节点配置信息的设备链接的路径。

退出状态 将返回下列退出值：

0 没有出现错误。

非零 对于 -F 选项（仅限 SPARC），返回的非零值表示输出设备不是帧缓冲区。对于 -x 选项，返回的非零值表示固件不支持 64 位。对于所有其他情况，返回的非零值表示出现错误。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Uncommitted（未确定）

另请参见 [fuser\(1M\)](#)、[modinfo\(1M\)](#)、[sysdef\(1M\)](#)、[driver\(4\)](#)、[driver.conf\(4\)](#)、[attributes\(5\)](#)

《Sun 硬件平台指南》

仅限 SPARC [openprom\(7D\)](#)

附注 prtconf 命令的输出与系统中安装的 PROM 的版本高度相关。在所有情况下，输出都有可能受到潜在影响。

driver not attached（驱动程序未连接）消息表示当前没有驱动程序关联到设备的此实例。通常，会根据要求或在需要时装载和安装驱动程序（并关联到硬件实例），在不使用设备时卸载驱动程序。

在 x86 平台上，使用 prtconf -vp 将提供 prtconf -v 信息的子集。prtconf -vp 的整数属性值可能需要字节交换以进行正确解释。

引用名	prtdiag – 显示系统诊断信息
用法概要	/usr/sbin/prtdiag [-v] [-l]
描述	prtdiag 可显示 sun4u、sun4v 和 x86 系统上的系统配置和诊断信息。 诊断信息列出系统中所有失败的现场可更换单元 (Field Replaceable Unit, FRU)。 prtdiag 的目录分层结构中的接口、输出和位置均未确定，在将来的发行版中可能会更改。 在 Sun Enterprise 10000 服务器上执行时，prtdiag 不显示诊断信息和环境状态。请参见系统服务处理器 (SSP) 上的 /var/opt/SUNWssp/adm/\${SUNW_HOSTNAME}/messages 文件以获取此服务器的此类信息。
选项	所支持的选项如下： -l 记录输出。如果系统中存在故障或错误，则仅将此信息输出到 syslogd(1M)。 -v 详细模式。显示最近发生 AC 电源故障的时间、最近硬件致命错误信息以及（如果适用）环境状态。硬件致命错误信息可用于修复和创建 FRU 的详细诊断信息。
退出状态	将返回下列退出值： 0 在系统中没有检测到故障或错误。 1 在系统中检测到故障或错误。 2 发生了内部 prtdiag 错误，例如，内存不足。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	system/library/platform
接口稳定性	Uncommitted（未确定）*

*输出不稳定。

另请参见	modinfo(1M) 、 prtconf(1M) 、 psrinfo(1M) 、 sysdef(1M) 、 syslogd(1M) 、 attributes(5) 、 open
附注	并非所有的诊断和系统信息在每个 Solaris 平台上都可用，因此无法通过 prtdiag 显示。在这些平台上，可从系统控制器获取进一步的信息。

引用名	prtdscp – 显示 DSCP IP 地址
用法概要	<pre>prtdscp [-v] prtdscp [-v] -h prtdscp [-v] -d prtdscp [-v] -s</pre>
描述	prtdscp 可显示与服务处理器通信协议域 (Domain to Service Processor Communications Protocol, DSCP) 链接有关的 IP 地址。如果没有指定参数, prtdscp 将显示 DSCP 链接两端的 IP 地址。通过使用 -s 选项或 -d 选项, 可分别单独显示服务处理器或域端的 IP 地址。
选项	所支持的选项如下: -v 详细模式。将有关程序内部进度的附加详细信息打印到 stderr。 -h 帮助。打印程序使用方法的简要用法概要并退出。将忽略所有其他命令行参数。 -d 仅显示本地域的 IP 地址。 -s 仅显示远程服务处理器的 IP 地址。
示例	示例1 显示本地域和远程服务处理器的地址 以下示例显示本地域的 IP 地址和远程 SP 的 IP 地址: <pre># prtdscp Domain Address: 192.168.103.2 SP Address: 192.168.103.1</pre> 示例2 显示本地 IP 地址 以下示例显示本地域的 IP 地址: <pre># prtdscp -d 192.168.103.2</pre> 示例3 显示远程 IP 地址 以下示例显示远程 SP 的 IP 地址: <pre># prtdscp -s 192.168.103.1</pre>
属性	有关以下属性的说明, 请参见 attributes(5) :

属性类型	属性值
可用性	system/domain-service-processor-protocol/sparc-enterprise 、
接口稳定性	Committed（已确定）

另请参见

[attributes\(5\)](#)

引用名

prtfu – 打印有关系统或域中的 FRU 的 FRUID 特定信息

用法概要

```
/usr/sbin/prtfu [-d] | [-clx] [container]
```

描述

prtfu 实用程序用于从系统或域中获取 FRUID 数据。它的输出是一个树结构，回显 FRU（Field-Replaceable Unit, 现场可更换单元）树中指向每个容器的路径。当找到容器时，容器中的数据也会以类似树的结构打印出来。

不使用任何参数的 prtfu 将打印 FRU 分层结构和所有 FRUID 容器数据。prtfu 打印到标准输出，该标准输出可能会重定向到某个文件。

选项

所支持的选项如下：

- c 仅打印容器及其数据。此选项不打印 FRU 树分层结构。
- d 在标准输出中打印当前注册的 DTD。
- l 仅打印 FRU 树分层结构。此选项不打印容器数据。
- x 以 XML 格式打印，并带有 prtfureg.dtd 的系统标识符 (SYSTEM)。

可以结合使用选项 -c 和 -l 以获取容器列表。

操作数

支持下列操作数：

container FRU 分层结构中特定容器的名称，即，-l 选项中显示的容器的名称或路径/名称。

退出状态

将返回下列退出值：

- 0 成功找到并打印了所有信息。
- >0 出现错误。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/fru-id

另请参见 [fruadm\(1M\)](#)、[attributes\(5\)](#)

引用名	prtpicl – 打印 PICL 树
用法概要	<code>/usr/sbin/prtpicl [-c <i>picl_class</i>] [-v]</code>
描述	<code>prtpicl</code> 命令可打印 PICL 守护进程维护的 PICL 树。 <code>prtpicl</code> 的输出包括节点的名称和 PICL 类。
选项	所支持的选项如下： <code>-c <i>picl_class</i></code> 仅打印指定 PICL 类的节点。 <code>-v</code> 以详细模式打印。在详细模式下， <code>prtpicl</code> 将打印每个节点的属性和值的列表。缺省情况下详细模式是禁用的。
退出状态	将返回下列退出值： <code>0</code> 成功完成。 <code>非零</code> 出现错误。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	system/picl

另请参见 [picld\(1M\)](#)、[attributes\(5\)](#)

引用名	prvtoc – 报告有关磁盘几何参数和分区的信息
用法概要	prvtoc [-fhs] [-t vfstab] [-m mnttab] device
描述	prvtoc 命令允许查看标签内容。只有超级用户才能使用此命令。 device 名称可以是以 /dev/rdisk/c?t?s2 形式显示的原始设备的文件名，也可以是以 /dev/dsk/c?t?s2 形式显示的块设备的文件名。
选项	所支持的选项如下： -f 报告磁盘空闲空间，包括空闲空间的起始块地址、块数和未使用的分区。 -h 在常规输出中忽略标题。 -m mnttab 使用 mnttab 作为挂载的文件系统的列表，而不使用 /etc/mnttab。 -s 在常规输出中忽略除列标题之外的所有标题。 -t vfstab 使用 vfstab 作为文件系统缺省值的列表，而不使用 /etc/vfstab。

示例

示例1 使用 prvtoc 命令

以下示例在一个 424 MB 硬盘上使用 prvtoc 命令：

```
example# prvtoc /dev/rdisk/c0t3d0s2
* /dev/rdisk/c0t3d0s2 partition map
*
* Dimension:
*   512 bytes/sector
*   80 sectors/track
*   9 tracks/cylinder
*   720 sectors/cylinder
*   2500 cylinders
*   1151 accessible cylinders
*
* Flags:
*   1: unmountable
*   10: read-only
* *
* Partition  Tag  Flags  First  Sector  Last  Mount Directory
*   0         2    00      0    76320   76319  /
*   1         3    01    76320  132480  208799
*   2         5    00      0    828720  828719
*   5         6    00   208800  131760  340559  /opt
*   6         4    00   340560  447120  787679  /usr
*   7         8    00   787680  41040   828719  /export/home
example#
```

上面 Tag 列中的数据表示分区类型，如下所示：

示例1 使用 prvtoc 命令 （续）

<i>Name</i>	<i>Number</i>
UNASSIGNED	0x00
BOOT	0x01
ROOT	0x02
SWAP	0x03
USR	0x04
BACKUP	0x05
STAND	0x06
VAR	0x07
HOME	0x08
ALTSCTR	0x09
CACHE	0x0a
RESERVED	0x0b
SYSTEM	0x0c
BOOT	0x18

上面 **Flags** 列中的数据表示分区是如何挂载的，如下所示：

<i>Name</i>	<i>Number</i>
MOUNTABLE, READ AND WRITE	0x00
NOT MOUNTABLE	0x01
MOUNTABLE, READ ONLY	0x10

示例2 使用 prvtoc 命令和 -f 选项

以下示例在一个 424 MB 硬盘上使用 prvtoc 命令和 -f 选项：

```
example# prvtoc -f /dev/rdisk/c0t3d0s2
FREE_START=0 FREE_SIZE=0 FREE_COUNT=0 FREE_PART=34
```

示例3 在超过 1 TB 的磁盘上使用 prvtoc 命令

以下示例在一个超过 1 TB 的磁盘上使用 prvtoc 命令：

示例3 在超过 1 TB 的磁盘上使用 pvtvtoc 命令 (续)

```
example# pvtvtoc /dev/rdisk/clt1d0s2
* /dev/rdisk/clt1d0s2 partition map
*
* Dimensions:
*   512 bytes/sector
* 3187630080 sectors
* 3187630013 accessible sectors
*
* Flags:
*   1: unmountable
*  10: read-only
*
*
* Partition Tag  Flags      First      Sector      Last
* Partition Tag  Flags      Sector      Count      Sector  Mount Directory
0      2      00          34      262144      262177
1      3      01      262178      262144      524321
6      4      00      524322 3187089340 3187613661
8      11     00 3187613662      16384 318763004
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见 [devinfo\(1M\)](#)、[fmthard\(1M\)](#)、[format\(1M\)](#)、[mount\(1M\)](#)、[attributes\(5\)](#)

警告 mount 命令不检查 "not mountable"（不可挂载）位。

引用名	psradm – 更改处理器操作状态
用法概要	<pre>psradm -f -i -n -s [-v] [-F] <i>processor_id</i> psradm -a -f -i -n -s [-v] [-F]</pre>
描述	psradm 实用程序可更改处理器的操作状态。处理器的合法状态包括 on-line、off-line、spare、faulted 和 no-intr。 on-line 处理器可处理 LWP（lightweight processes，轻量级进程），并且可被系统中的 I/O 设备中断。 off-line 处理器不处理任何 LWP。通常情况下，系统中的 I/O 设备不能中断 off-line 处理器。在某些处理器上或在某些情况下，可能无法为 off-line 处理器禁用中断。因此，off-line 的实际效果可能会因计算机而异。 spare 处理器不处理任何 LWP。系统的特权用户或内核可以使备用处理器处于 on-line、off-line 或 no-intr 状态，以响应系统状态的更改。 faulted 处理器由内核标识，内核会定时监视处理器的行为。特权用户可以将 faulted 处理器的状态设置为 on-line、off-line、spare 或 no-intr，但是必须使用强制选项来执行此操作。 no-intr 处理器可处理 LWP，但不能被 I/O 设备中断。 如果有 LWP 绑定到处理器，则除非使用附加的 -F 选项，否则无法将处理器设为 off-line 或 spare 状态。-F 选项会在更改处理器的操作状态之前删除这类 LWP 的处理器绑定。在某些体系结构上，例如当系统依赖于处理器提供的某些资源时，可能无法将某些处理器设为 off-line 或 spare 状态。 系统中必须至少有一个处理器能够处理 LWP。还必须至少有一个处理器是可中断的。因为 off-line 或 spare 处理器是可中断的，所以可以实现这样一个操作系统，该操作系统中的一个处理器处于 no-intr 状态，所有其他处理器处于 off-line 或 spare 状态，但是有一个或多个处理器可接受中断。 如果任何指定的处理器被关闭，psradm 可能会启动一个或多个处理器。 只有具有 PRIV_SYS_RES_CONFIG 特权的用户才能使用 psradm 实用程序。
选项	所支持的选项如下： -a 在所有处理器或尽可能多的处理器上执行操作。 -f 将指定的处理器设为 off-line 状态。 -F 强制转换到附加指定状态。如果一个或多个指定处理器处于 faulted 状态，则需使用此选项。如果未指定其他转换选项，则将指定的处理器设置为 faulted 状态。只能强制转换到 faulted、spare 或 off-line 状态。建议管理员为 pbind(1M) 使用 -Q 选项，以找出强制执行处理器状态转换会影响哪些线程。 -i 将指定处理器设置为 no-intr。

- n 将指定处理器设为 on-line 状态。
- s 将指定处理器设为 spare 状态。
- v 输出一条消息，显示每个尝试操作的结果。

操作数

支持下列操作数：

processor_id 设置为 on-line 或 off-line、spare 或 no-intr 状态的处理器的处理器的 ID。

将 *processor_id* 指定为单个处理器编号（例如，3）、由空格分隔的多个处理器编号（例如，1 2 3）或者处理器编号的范围（例如，1-4）。也可以组合使用范围和单个或多个 *processor_id*（例如，1-3 5 7-8 9）。

示例

示例1 将处理器设置为 off-line

以下示例将处理器 2 和 3 设置为 off-line：

```
% psradm -f 2 3
```

示例2 将处理器设置为 no-intr

以下示例将处理器 1 和 2 设置为 no-intr：

```
% psradm -i 1 2
```

示例3 将处理器设置为 spare

以下示例将处理器 1 和 2 设置为处于 spare 状态（即使处理器 1 或 2 处于 faulted 状态）：

```
% psradm -F -s 1 2
```

示例4 将所有处理器设置为 on-line

```
% psradm -a -n
```

示例5 将处理器强制设置为 off-line

以下示例将处理器 1 和 2 设置为 offline 状态，并且撤销其与所绑定进程之间的绑定：

```
% psradm -F -f 1 2
```

退出状态

将返回下列退出值：

0 成功完成。

>0 出现错误。

文件 /etc/wtmpx 记录了处理器状态更改的记录

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见 [pbind\(1M\)](#)、[psrinfo\(1M\)](#)、[psrset\(1M\)](#)、[p_online\(2\)](#)、[processor_bind\(2\)](#)、[attributes\(5\)](#)

诊断

psradm: processor 4: Invalid argument
配置中不存在指定处理器。

psradm: processor 3: Device busy
无法将指定处理器设为 off-line 状态，因为它绑定有 LWP，或者是系统中最后一个 on-line 状态的处理器，或者系统需要该处理器提供某些基本服务。

psradm: processor 3: Device busy
无法将指定处理器设置为 no-intr 状态，因为它是系统中最后一个可中断的处理器，或者是系统中唯一可以提供系统所需的中断服务的处理器。

psradm: processor 3: Device busy
指定处理器已关闭，并且无法启动，因为某些特定于平台的资源不可用。

psradm: processor 0: Not owner
用户不具有更改处理器状态的权限。

psradm: processor 2: Operation not supported
指定处理器已关闭，并且平台不支持启动单个处理器。

引用名	psrinfo – 显示有关处理器的信息
用法概要	<pre>psrinfo [-p] [-v] [<i>processor_id</i>]... psrinfo [-p] -s <i>processor_id</i></pre>
描述	psrinfo 可显示有关处理器的信息。每个物理处理器可支持多个虚拟处理器。每个虚拟处理器都是一个实体，具有各自的中断 ID，并能够执行独立线程。 如果不带 <i>processor_id</i> 操作数，psrinfo 将为每个配置的处理器显示一行，表明该处理器是处于 on-line（联机）、no-intr（不可中断）、spare（备用）、off-line（脱机）、faulted（故障）状态还是已关闭，以及上一次状态发生变化的时间。使用 <i>processor_id</i> 操作数可显示特定处理器的信息。请参见“操作数”部分。
选项	所支持的选项如下： -s <i>processor_id</i> 无提示模式。如果指定的处理器完全联机，则显示 1。如果指定的处理器处于不可中断、备用、脱机、故障或关闭状态，则显示 0。 在 shell 脚本中使用 psrinfo 时，请使用无提示模式。 -p 显示系统中物理处理器的数目。 与 -v 选项结合使用时，报告每个物理处理器的附加信息。 -v 详细模式。显示指定处理器的附加信息，包括处理器类型、浮点单元类型和时钟速度。如果其中的任何信息无法判断，psrinfo 将显示 unknown（未知）。 与 -p 选项结合使用时，报告每个物理处理器的附加信息。
操作数	支持下列操作数： <i>processor_id</i> 要显示其信息的处理器的处理器 ID。 将 <i>processor_id</i> 指定为单个处理器编号（例如，3）、由空格分隔的多个处理器编号（例如，1 2 3）或者处理器编号的范围（例如，1-4）。也可以组合使用范围和单个或多个 <i>processor_id</i>（例如，1-3 5 7-8 9）。
示例	示例1 以详细模式显示所有配置的处理器的信息 以下示例以详细模式显示所有配置的处理器的信息。 <pre>psrinfo -v</pre> 示例2 确定处理器是否联机 以下示例在 shell 脚本中使用 psrinfo 判断处理器是否联机。

示例2 确定处理器是否联机 (续)

```
if [ "psrinfo -s 3 2> /dev/null" -eq 1 ]
then
    echo "processor 3 is up"
fi
```

示例3 显示关于系统中物理处理器的信息

没有附加参数时，-p 选项将显示一个单个整数，即系统中物理服务器的数目：

```
> psrinfo -p
8
```

psrinfo 也接受命令行参数（处理器 ID）：

```
> psrinfo -p 0 512    # IDs 0 and 512 exist on the
1                    # same physical processor

> psrinfo -p 0 1      # IDs 0 and 1 exist on different
2                    # physical processors
```

在本例中，虚拟处理器 0 和 512 存在于同一个物理处理器中。虚拟处理器 0 和 1 则非如此。这是此示例的特定情况，并非通用的规则。

退出状态

将返回下列退出值：

- 0 成功完成。
- >0 出现错误。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见

[psradm\(1M\)](#)、[p_online\(2\)](#)、[processor_info\(2\)](#)、[attributes\(5\)](#)

诊断

psrinfo: processor 9: Invalid argument 指定的处理器不存在。

引用名 psrset – creation and management of processor sets

用法概要

```
psrset -a [-F] processor_set_id processor_id...
psrset -b processor_set_id pid [/lwpid]...
psrset -c [-F] [processor_id]...
psrset -d processor_set_id...
psrset -e processor_set_id command [argument(s)]
psrset -f processor_set_id
psrset [-i] [processor_set_id]...
psrset -n processor_set_id
psrset -p [processor_id]...
psrset [-q] [pid [/lwpid]]...
psrset -Q [processor_set_id]...
psrset -r [-F] processor_id...
psrset -u pid [/lwpid]...
psrset -U [processor_set_id]...
```

描述

The psrset utility controls the management of processor sets. Processor sets allow the binding of processes or LWPs to groups of processors, rather than just a single processor. Processors assigned to processor sets can run only LWPs that have been bound to that processor set.

This command cannot be used to modify processor disposition when pools are enabled. Use [pooladm\(1M\)](#) and [poolcfg\(1M\)](#) to modify processor set configuration through the resource pools facility.

选项

The following options are supported:

- a Assign the specified processors to the specified processor set. With the additional -F option, all LWPs bound to the specified processors will be unbound prior to changing processor sets.

This option is restricted to users with the PRIV_SYS_RES_CONFIG privilege.

- b Bind all or a subset of the LWPs of the specified processes to the specified processor set.

LWPs bound to a processor set are restricted to run only on the processors in that set. Processes can only be bound to non-empty processor sets, that is, processor sets that have had processors assigned to them.

Bindings are inherited, so new LWPs and processes created by a bound LWP have the same binding. Binding an interactive shell to a processor, for example, binds all commands executed by the shell.

This option is restricted to users with the `PRIV_SYS_RES_CONFIG` privilege.

- c Create a new processor set and displays the new processor set ID. With the additional -F option, all LWPs bound to the specified processors will be unbound prior to assigning them to the processor set being created.

If a list of processors is given, it also attempts to assign those processors to the processor set. If this succeeds, the processors are idle until LWPs are bound to the processor set. This option is restricted to users with the `PRIV_SYS_RES_CONFIG` privilege.

Only a limited number of processor sets can be active (created and not destroyed) at a given time. This limit is always be greater than the number of processors in the system. If the -c option is used when the maximum number of processor sets is already active, the command fails.

The following format is used for the first line of output of the -c option when the `LC_MESSAGES` locale category specifies the “C” locale. In other locales, the strings `created`, `processor`, and `set` can be replaced with more appropriate strings corresponding to the locale.

`"created processor set %d\n" processor set ID`

- d Remove the specified processor set, releasing all processors and processes associated with it.

This option is restricted to users with the `PRIV_SYS_RES_CONFIG` privilege.

- e Execute a command (with optional arguments) in the specified processor set.

The command process and any child processes are executed only by processors in the processor set.

This option is restricted to users with the `PRIV_SYS_RES_CONFIG` privilege.

- f Disables interrupts for all processors within the specified processor set. See [psradm\(1M\)](#).

If some processors in the set cannot have their interrupts disabled, the other processors still have their interrupts disabled, and the command reports an error and return non-zero exit status.

This option is restricted to users with the `PRIV_SYS_RES_CONFIG` privilege.

- F Forces the specified processor set operation by unbinding all threads bound to the specified processor. Only the -a or the -r option can be used in combination with this option. Administrators are encouraged to use the -Q option for [pbind\(1M\)](#) to find out which threads will be affected by such operation.
- i Display a list of processors assigned to each named processor set. If no argument is given, a list of all processor sets and the processors assigned to them is displayed. This is also the default operation if the psrset command is not given an option.
- n Enable interrupts for all processors within the specified processor set. See [psradm\(1M\)](#).

This option is restricted to users with the PRIV_SYS_RES_CONFIG privilege.
- p Display the processor set assignments for the specified list of processors. If no argument is given, the processor set assignments for all processors in the system is given.
- q Display the processor set bindings of the specified processes or of all processes. If a process is composed of multiple LWPs which have different bindings and the LWPs are not explicitly specified, the bindings of only one of the bound LWPs is displayed. The bindings of a subset of LWPs can be displayed by appending “/lwpids” to the process IDs. Multiple LWPs may be selected using “-” and “,” delimiters. See EXAMPLES.
- Q Display the LWPs bound to the specified list of processor sets, or all LWPs with processor set bindings.
- r Remove a list of processors from their current processor sets. Processors that are removed return to the general pool of processors.

Processors with LWPs bound to them using [pbind\(1M\)](#) can be assigned to or removed from processor sets using the -F option.

This option is restricted to users with the PRIV_SYS_RES_CONFIG privilege.
- u Remove the processor set bindings of a subset or all the LWPs of the specified processes, allowing them to be executed on any on-line processor if they are not bound to individual processors through pbind.

Users with the PRIV_SYS_RES_CONFIG privilege can unbind any process or LWP from any active processor set. Other users can unbind processes and LWPs from processor sets that do not have the PSET_NOESCAPE attribute set. In addition, the user must have permission to control the affected processes; the real or effective user ID of the user must match the real or saved user ID of the target processes.
- U Removes the bindings of all LWPs bound to the specified list of processor sets, or to any processor set if no argument is specified.

- 操作数

The following operands are supported:

pid

Specify *pid* as a process ID.

lwpid

The set of LWPIDs of the specified process to be controlled or queried.
The syntax for selecting LWP IDs is as follows:

2,3,4-8

LWP IDs 2, 3, and 4 through 8

-4

LWPs whose IDs are 4 or below

4-

LWPs whose IDs are 4 or above

processor_id

Specify *processor_id* as an individual processor number (for example, 3), multiple processor numbers separated by spaces (for example, 1 2 3), or a range of processor numbers (for example, 1-4). It is also possible to combine ranges and (individual or multiple) *processor_ids* (for example, 1-3 5 7-8 9).

processor_set_id

Specify *processor_set_id* as a processor set ID.

退出状态

The following exit values are returned:

0

Successful completion.

non-0

An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[pbind\(1M\)](#), [pooladm\(1M\)](#), [poolcfg\(1M\)](#), [psradm\(1M\)](#), [psrinfo\(1M\)](#), [processor_bind\(2\)](#), [processor_info\(2\)](#), [pset_bind\(2\)](#), [pset_create\(2\)](#), [pset_info\(2\)](#), [sysconf\(3C\)](#), [libpool\(3LIB\)](#), [attributes\(5\)](#), [privileges\(5\)](#)

诊断

The following output indicates that the specified process did not exist or has exited:

psrset: cannot query pid 31: No such process

The following output indicates that the user does not have permission to bind the process:

psrset: cannot bind pid 31: Not owner

The following output indicates that the user does not have permission to assign the processor:

psrset: cannot assign processor 4: Not owner

The following output indicates that the specified processor is not on-line, or the specified processor does not exist.

psrset: cannot assign processor 8: Invalid argument

系统管理命令－第2单元

1805

The following output indicates that an LWP in the specified process is bound to a processor and cannot be bound to a processor set that does not include that processor:

```
psrset: cannot bind pid 67: Device busy
```

The following output indicates that the specified processor could not be added to the processor set. This can be due to bound LWPs on that processor, or because that processor cannot be combined in the same processor set with other processors in that set, or because the processor is the last one in its current processor set:

```
psrset: cannot assign processor 7: Device busy
```

The following output indicates that the specified processor set does not exist:

```
psrset: cannot execute in processor set 8: Invalid argument
```

The following output indicates that the maximum number of processor sets allowed in the system is already active:

```
psrset: cannot create processor set: Not enough space
```

The following output indicates that the pools facility is active.

```
psrset: cannot assign processor 7: Operation not supported
psrset: cannot bind pid 31: Operation not supported
psrset: cannot bind pid 31: Operation not supported
psrset: could not create processor set: Operation not supported
psrset: could not remove processor set 1: Operation not supported
psrset: cannot exec in processor set 1: Operation not supported
psrset: cannot remove processor 7: Operation not supported
psrset: cannot unbind pid 31: Operation not supported
```

引用名 pwck, grpck – password/group file checkers

用法概要 /usr/sbin/pwck [*filename*]
/usr/sbin/grpck [*filename*]

描述 pwck scans the password file and notes any inconsistencies. The checks include validation of the number of fields, login name, user ID, group ID, and whether the login directory and the program-to-use-as-shell exist. The default password file is /etc/passwd.

grpck verifies all entries in the group file. This verification includes a check of the number of fields, group name, group ID, whether any login names belong to more than NGROUPS_MAX groups, and that all login names appear in the password file. grpck also issues a warning if it finds an entry (a single line) in the group file longer than 2047 characters. Such an entry causes group maintenance commands, such as [groupdel\(1M\)](#) and [groupmod\(1M\)](#), to fail.

The default group file is /etc/group.

All messages regarding inconsistent entries are placed on the stderr stream.

文件

- /etc/group
- /etc/passwd

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [groupdel\(1M\)](#), [groupmod\(1M\)](#), [getpwent\(3C\)](#), [group\(4\)](#), [passwd\(4\)](#), [attributes\(5\)](#)

诊断 Group entries in /etc/group with no login names are flagged.

Group file '*filename*' is empty

The /etc/passwd or /etc/group file is an empty file.

cannot open file *filename*: No such file or directory

The /etc/passwd or /etc/group file does not exist.

附注 If no filename argument is given, grpck checks the local group file, /etc/group, and also makes sure that all login names encountered in the checked group file are known to the system [getpwent\(3C\)](#) routine. This means that the login names may be supplied by a network name service.

引用名

用法概要

描述

pwconv – installs and updates /etc/shadow with information from /etc/passwd

pwconv

The pwconv command creates and updates /etc/shadow with information from /etc/passwd.

pwconv relies on a special value of 'x' in the password field of /etc/passwd . This value of 'x' indicates that the password for the user is already in /etc/shadow and should not be modified.

If the /etc/shadow file does not exist, this command will create /etc/shadow with information from /etc/passwd. The command populates /etc/shadow with the user's login name, password, and password aging information. If password aging information does not exist in /etc/passwd for a given user, none will be added to /etc/shadow. However, the last changed information will always be updated.

If the /etc/shadow file does exist, the following tasks will be performed:

Entries that are in the /etc/passwd file and not in the /etc/shadow file will be added to the /etc/shadow file.

Entries that are in the /etc/shadow file and not in the /etc/passwd file will be removed from /etc/shadow.

Password attributes (for example, password and aging information) that exist in an /etc/passwd entry will be moved to the corresponding entry in /etc/shadow.

The pwconv command can only be used by the super-user.

- 文件
- /etc/opasswd
 - /etc/oshadow
 - /etc/passwd
 - /etc/shadow

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[passwd\(1\)](#), [usermod\(1M\)](#), [passwd\(4\)](#), [attributes\(5\)](#)

- 诊断
- pwconv exits with one of the following values:
- 0 SUCCESS.
 - 1 Permission denied.
 - 2 Invalid command syntax.
 - 3 Unexpected failure. Conversion not done.

- 4 Unexpected failure. Password file(s) missing.
- 5 Password file(s) busy. Try again later.
- 6 Bad entry in /etc/shadow file.

引用名	quot – summarize file system ownership
用法概要	<pre>quot [-acfhnv] filesystem...</pre> <pre>quot -a [-cfhnv]</pre>
描述	quot displays the number of blocks (1024 bytes) in the named <i>filesystem</i> (one or more) currently owned by each user. There is a limit of 2048 blocks. Files larger than this will be counted as a 2048 block file, but the total block count will be correct.
选项	The following options are supported: -a Generate a report for all mounted file systems. -c Display three columns giving a file size in blocks, the number of files of that size, and a cumulative total of blocks containing files of that size or a smaller size. -f Display three columns giving, for each user, the number of blocks owned, the count of number of files, and the user name. This option is incompatible with the -c and -v options. -h Estimate the number of blocks in the file. This does not account for files with holes in them. -n Attach names to the list of files read from standard input. quot -n cannot be used alone, because it expects data from standard input. For example, the pipeline ncheck myfilesystem sort +0n quot -n myfilesystem will produce a list of all files and their owners. This option is incompatible with all other options. -v In addition to the default output, display three columns containing the number of blocks not accessed in the last 30, 60, and 90 days.
操作数	<i>filesystem</i> mount-point of the filesystem(s) being checked
用法	See largefile(5) for the description of the behavior of quot when encountering files greater than or equal to 2 Gbyte (2 ³¹ bytes).
退出状态	0 Successful operation. 32 Error condition (bad or missing argument, bad path, or other error).
文件	/etc/mnttab Lists mounted file systems. /etc/passwd Used to obtain user names
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [du\(1\)](#), [mnttab\(4\)](#), [passwd\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

附注 This command can only be used by the super-user.

引用名 quota – display a user's ufs or zfs file system disk quota and usage

用法概要 quota [-v] [*username*]

描述 quota displays users' UFS or ZFS disk usage and limits. Only the super-user may use the optional *username* argument to view the limits of other users.

quota without options only display warnings about mounted file systems where usage is over quota. Remotely mounted file systems which do not have quotas turned on are ignored.

username can be the numeric UID of a user.

选项 -v Display user's quota on all mounted file systems where quotas exist.

用法 See [largefile\(5\)](#) for the description of the behavior of quota when encountering files greater than or equal to 2 Gbyte (2³¹ bytes).

文件 /etc/mnttab list of currently mounted filesystems

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [edquota\(1M\)](#), [quotaon\(1M\)](#), [quotacheck\(1M\)](#), [repquota\(1M\)](#), [rquotad\(1M\)](#), [attributes\(5\)](#), [largefile\(5\)](#), [zones\(5\)](#)

附注 quota displays quotas for NFS mounted UFS- or ZFS-based file systems if the rquotad daemon is running. See [rquotad\(1M\)](#). In a [zones\(5\)](#) environment, quota displays quotas only for the zone in which it is invoked.

quota can display entries for the same file system multiple times for multiple mount points. For example,

```
# quota -v user1
```

might display identical quota information for user1 at the mount points /home/user1, /home/user2, and /home/user, if all three mount points are mounted from the same file system with quotas turned on.

引用名	quotacheck – ufs file system quota consistency checker				
用法概要	<pre>quotacheck [-fp] [-v] filesystem... quotacheck -a [-fpv]</pre>				
描述	quotacheck examines each mounted ufs file system, builds a table of current disk usage, and compares this table against the information stored in the file system's disk quota file. If any inconsistencies are detected, both the quota file and the current system copy of the incorrect quotas are updated. <i>filesystem</i> is either a file system mount point or the block device on which the file system resides. quotacheck expects each file system to be checked to have a quota file named <i>quotas</i> in the root directory. If none is present, quotacheck will not check the file system. quotacheck accesses the character special device in calculating the actual disk usage for each user. Thus, the file systems that are checked should be quiescent while quotacheck is running.				
选项	The following options are supported: -a Check the file systems which <i>/etc/mnttab</i> indicates are ufs file systems. These file systems must be read-write mounted with disk quotas enabled, and have an <i>rq</i> entry in the <i>mntopts</i> field in <i>/etc/vfstab</i>. -f Force check on file systems with logging enabled. Use in combination with the <i>-p</i> option. -p Check quotas of file systems in parallel. For file systems with logging enabled, no check is performed unless the <i>-f</i> option is also specified. -v Indicate the calculated disk quotas for each user on a particular file system. quotacheck normally reports only those quotas modified.				
用法	See largefile(5) for the description of the behavior of quotacheck when encountering files greater than or equal to 2 Gbyte (2 ³¹ bytes).				
文件	<table><tr><td><i>/etc/mnttab</i></td><td>Mounted file systems</td></tr><tr><td><i>/etc/vfstab</i></td><td>List of default parameters for each file system</td></tr></table>	<i>/etc/mnttab</i>	Mounted file systems	<i>/etc/vfstab</i>	List of default parameters for each file system
<i>/etc/mnttab</i>	Mounted file systems				
<i>/etc/vfstab</i>	List of default parameters for each file system				
属性	See attributes(5) for descriptions of the following attributes:				

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[edquota\(1M\)](#), [quota\(1M\)](#), [quotaon\(1M\)](#), [repquota\(1M\)](#), [attributes\(5\)](#), [largefile\(5\)](#),
[quotactl\(7I\)](#), [mount_ufs\(1M\)](#)

引用名	quotaon, quotaoff – turn ufs file system quotas on and off
用法概要	<pre>quotaon [-v] filesystem... quotaon -a [-v] quotaoff [-v] filesystem... quotaoff -a [-v]</pre>
描述	quotaon turns on disk quotas for one or more ufs file systems. Before a file system may have quotas enabled, a file named <code>quotas</code>, owned by root, must exist in the root directory of the file system. See edquota(1M) for details on how to modify the contents of this file. quotaoff turns off disk quotas for one or more ufs file systems. The file systems specified must already be mounted. These commands update the <code>mntopts</code> field of the appropriate entries in <code>/etc/mnttab</code> to indicate when quotas are on or off for each file system. If quotas are on, the string <code>quota</code> will be added to <code>mntopts</code>; if quotas are off, the <code>quota</code> string is not present. <i>filesystem</i> must be either the mount point of a file system, or the block device on which the file system resides.
选项	
quotaon	-a This option is normally used at boot time to enable quotas. It applies only to those file systems in <code>/etc/vfstab</code> which have “<code>rq</code>” in the <code>mntopts</code> field, are currently mounted ‘<code>rw</code>’, and have a <code>quotas</code> file in the root directory. -v Display a message for each file system after quotas are turned on.
quotaoff	-a Force all file systems in <code>/etc/mnttab</code> to have their quotas disabled. -v Display a message for each file system affected.
用法	See largefile(5) for the description of the behavior of <code>quotaon</code> and <code>quotaoff</code> when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).
文件	<pre>/etc/mnttab mounted file systems /etc/vfstab list of default parameters for each file system</pre>
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	system/core-os

另请参见

[edquota\(1M\)](#), [quota\(1M\)](#), [quotacheck\(1M\)](#), [repquota\(1M\)](#), [mnttab\(4\)](#), [vfstab\(4\)](#),
[attributes\(5\)](#), [largefile\(5\)](#), [quotactl\(7I\)](#)

引用名	raidctl – RAID 硬件实用程序
用法概要	<pre> raidctl -C "disks" [-r <i>raid_level</i>] [-z <i>capacity</i>] [-s <i>stripe_size</i>] [-f] controller raidctl -d [-f] <i>volume</i> raidctl -F <i>filename</i> [-f] <i>controller...</i> raidctl -a {set unset} -g <i>disk</i> {<i>volume</i> <i>controller</i>} raidctl -p "<i>param=value</i>" [-f] <i>volume</i> raidctl -c [-f] [-r <i>raid_level</i>] <i>disk1 disk2 [disk3...]</i> raidctl -l -g <i>disk controller</i> raidctl -l <i>volume</i> raidctl -l <i>controller...</i> raidctl [-l] raidctl -S [<i>volume</i> <i>controller</i>] raidctl -S -g <i>disk controller</i> raidctl -h </pre>
描述	raidctl 实用程序是支持各种 RAID 控制器的硬件 RAID 配置工具，为最终用户提供了用于创建、删除或显示 RAID 卷的 CLI（Command Line Interface，命令行接口）。此实用程序还可用于设置卷的属性、将热备用 (hot-spare, HSP) 磁盘分配到卷或控制器，以及为 RAID 控制器更新固件/fcode/BIOS。 raidctl 实用程序要求具有受基础文件系统权限控制的特权。只有特权用户才能操控 RAID 系统配置。如果非特权用户尝试运行 raidctl，则此命令会失败，并返回退出状态 1。 如本手册页所述，raidctl 实用程序定义了一系列用于管理功能完备的 RAID 控制器的命令行选项。不过，对给定选项的支持依赖于以下两个元素： ■ 存在软件驱动程序 ■ RAID 设备的固件级别 之所以依赖于软件驱动程序，是 raidctl 的设计使然。此实用程序基于支持为各种驱动程序插入插件模块的公共库而构建。目前，Solaris 操作系统附带了用于 mpt 驱动程序的一个插件。该插件并不是支持所有 raidctl 选项。在给定的存储设备上，这些选项可能会进一步受设备的固件级别限制。 对各个 raidctl 选项的支持级别不能由 raidctl 来确定。用户必须依据 RAID 控制器或硬件平台的文档来确定。 当前，raidctl Currently, raidctl 对下列 RAID 控制器提供一定级别的支持： ■ LSI1020 SCSI HBA

- LSI1030 SCSI HBA
- LSI1064 SAS HBA
- LSI1068 SAS HBA

在 X86-32/64 和 SPARC 平台上，上述所有 HBA 都是由 mpt 驱动程序维护的。

选项

所支持的选项如下：

`-c "disks" [-r raid_level] [-z capacity] [-s stripe_size] [-f] controller`
使用指定的磁盘创建 RAID 卷。

使用此选项创建 RAID 卷时，会自动为新创建的卷生成标识，并且 `raidctl` 会将其报告给用户。

此选项指定的参数包含用来组成要创建的卷的元素。这些元素既可以是磁盘，也可以是子卷，其中磁盘用空格分隔，子卷是用括号分组的一组磁盘。所有磁盘都应该采用 `C.ID.L` 表示形式（例如，`0.1.2` 表示通道为 0，目标 ID 为 1，逻辑单元号为 2 的物理磁盘）。该参数必须与 `-r` 选项指定的 RAID 级别相匹配，即使会被忽略也应该匹配。这意味着该参数只能是：

对于 RAID 0
至少有 2 个磁盘

对于 RAID 1
只能有 2 个磁盘

对于 RAID 1E
至少有 3 个磁盘

对于 RAID 5
至少有 3 个磁盘

对于 RAID 10
至少有 2 个子卷，每个子卷必须由 2 个磁盘组成

对于 RAID 50
至少有 2 个子卷，每个子卷必须由至少 3 个磁盘组成，且每个子卷中的磁盘数量必须相同

例如，表达式 `"0.0.0 0.1.0"` 表示 2 个指定磁盘组成了一个 RAID 卷，该卷可以是 RAID 0 或 RAID 1 卷。`"(0.0.0 0.1.0)(0.2.0 0.3.0)"` 表示前 2 个磁盘和后 2 个磁盘分别组成了 2 个子卷，这 2 个子卷组成了 RAID 10 卷。有关更多样例，请参见示例部分。

`-r` 选项指定了要创建的卷的 RAID 级别。可能的级别包括 0、1、1E、5、10、50。如果省略了此选项，`raidctl` 在缺省情况下会创建 RAID 1 卷。

`-z` 选项指定了要创建的卷的容量。单位可以是兆字节、千兆字节或兆字节（例如 2t、10g、20m 等）。如果省略了此选项，`raidctl` 会计算指定的磁盘可以创建的最大卷容量，并使用该值来创建卷。

-s 选项指定要创建的卷的条带大小。可能的值包括 512、1k、2k、4k、8k、16k、32k、64k 或 128k。如果省略了此选项，raidctl 会为此卷选择合适的值（例如 64k）。

在某些情况下，创建 RAID 卷会导致指定磁盘上的数据丢失（例如，在 LSI1020、LSI1030、SAS1064 或 SAS1068 HBA 上就会出现这种情况），raidctl 会提示用户确认卷的创建。使用 -f 选项可在不提示用户确认的情况下强制创建卷。

controller（控制器）参数用于标识指定磁盘所属的 RAID 控制器。可使用 -l 选项列出控制器的 ID 号。

-d [-f] *volume*

删除指定为 *volume*（卷）的 RAID 卷。该卷是采用规范的格式（例如 `c0t0d0`）指定的。

卷被删除后，所有数据都会丢失。因此，除非指定了 -f 选项，否则 raidctl 会在删除卷之前提示用户确认。

从 LSI1020、LSI1030、SAS1064 或 SAS1068 HBA 中删除 RAID 1 卷时，主磁盘和辅助磁盘会“分裂”。如果卷原来处于 SYNCING（正在同步）状态，主磁盘会包含数据，辅助磁盘不会包含数据。如果卷状态原来为 OPTIMAL（最佳），则两种磁盘都会包含数据的完整映像。

-F *filename* [-f] *controller...*

更新在指定控制器上运行的固件。除非提供了 -f 选项，否则 raidctl 实用程序会提示用户确认此操作。

-a {set|unset} -g *disk* {*volume* | *controller*}

如果指定了卷，raidctl 会根据 -a 选项指定的值将磁盘设置为该卷专用的本地热备用磁盘，或取消将其设置为本地热备用磁盘。如果指定了控制器，raidctl 会将磁盘设置为全局热备用磁盘或取消将其设置为全局热备用磁盘。

-p "*param=value*" [-f] *volume*

更改给定 RAID 卷的属性值。可使用此选项来更改高速缓存写入策略或激活卷。当更改高速缓存写入策略时，*param* 应该为字符串 `wp` (SET_WR_POLICY)，*value* 可以是 `on` 或 `off`。当用来激活卷时，*param* 应该为 `state`，*value* 应该为 `activate`。

更改 RAID 卷的属性会影响 RAID 控制器的内部行为，因此，除非指定了 -f 选项，否则 raidctl 会在应用更改之前提示用户确认。

-c [-f] [-r *raid_level*] *disk1 disk2* [*disk3...*]

使用指定的磁盘创建卷。这是具有类似功能的 -c 选项的替代选项。此选项是为实现兼容性而保留的，但只能用于在 LSI1020、LSI1030、SAS1064 和 SAS1068 HBA 上创建 RAID 0、RAID 1 或 RAID 1E 卷。对于其他 HBA，用户只能使用 -c 选项。

可使用 -r 选项指定目标卷的 RAID 级别。如果省略了 -r 选项，raidctl 会创建 RAID 1 卷。

磁盘必须采用 Solaris 规范格式指定（例如 `c0t0d0`）。

使用此选项创建 RAID 1 卷会用磁盘 1 中的内容替换磁盘 2 中的内容。

当用户使用此选项创建 RAID 卷时，RAID 卷将采用磁盘 1 的标识。其他磁盘将成为不可见的，RAID 卷将显示为一个磁盘。

使用此选项创建卷在缺省情况下是交互式的。用户必须肯定的回答提示问题，才能创建卷。使用 `-f` 选项可在不提示用户确认的情况下强制创建卷。

`-l -g disk controller`

显示给定控制器的指定磁盘的相关信息。输出包括以下信息：

Disk

以 `C.ID.L` 表示形式显示磁盘。

Vendor

显示供应商 ID 字符串。

Product

显示产品 ID 字符串。

Capacity

显示磁盘的总容量。

Status

显示磁盘的当前状态。状态可能为 "GOOD"（良好）（正在正常运行）、"FAILED"（故障）（不能正常使用）或 "MISSING"（缺少）（磁盘不存在）。

HSP

指示磁盘已设置为全局热备用磁盘、本地热备用磁盘还是标准磁盘。如果是本地热备用磁盘，则会显示此磁盘被分配到的所有卷。

GUID

指定的磁盘的 GUID 字符串。这是附加的数据，在某些情况下可能不可用。

`-l volume`

显示指定卷的相关信息。输出包括以下信息：

卷号

以规范格式显示卷。

Sub

显示子卷，如果指定的卷为 RAID 10 或 RAID 50 卷。

Disk

显示组成指定卷的所有磁盘。

Stripe Size

显示卷的条带大小。

Status

显示指定卷或者组成指定卷的子卷或磁盘的状态。对于非活动卷，状态应当为 **INACTIVE**（非活动），其他情况下可能为 **OPTIMAL**（最佳）（正在以最佳状态运行）、**DEGRADED**（降级）（正在以降级的功能运行）、**FAILED**（故障）（不能正常使用）或 **SYNC**（同步）（磁盘正在同步）。对于磁盘，状态可能为 **GOOD**（良好）、**FAILED**（故障）或 **MISSING**（缺少）。

Cache

指示是否向 **I/O** 写入活动应用了高速缓存。高速缓存可能为 "ON" 或 "OFF"。

RAID level

显示 RAID 级别。RAID 级别可能为 0、1、1E、5、10 或 50。

-l controller ...

显示指定控制器的相关信息。输出包括以下信息：

控制器

显示 RAID 控制器的 ID 号。

Type

显示 RAID 控制器的产品类型。

fw_version

显示控制器的固件版本。

[-l]

列出 **raidctl** 实用程序可以操控的所有 RAID 相关对象，包括所有可用的 RAID 控制器、RAID 卷和物理磁盘。可以省略 **-l** 选项。

输出包括以下信息：

控制器

显示 RAID 控制器的 ID 号。

Volume

显示逻辑 RAID 卷名。

Disk

以 **C.ID.L** 表示形式显示 RAID 磁盘。

-S [volume | controller]

拍摄 RAID 配置信息（包括所有可用的 RAID 设备、RAID 控制器、卷和磁盘）的快照。

输出的每一行都指定了 RAID 设备及其相关信息，相互之间以空格分隔。所有卷和磁盘都属于最后指定的控制器。

输出列出以下信息：

Controller

显示控制器 ID 号以及用双引号引起来的控制器类型字符串。

Volume

显示 RAID 卷名、组件磁盘的数目、组件磁盘的 C.ID.L 表示形式、RAID 级别，以及状态。状态可能为 OPTIMAL（最佳）、DEGRADED（降级）、FAILED（故障）或 SYNCING（正在同步）。

Disk

显示磁盘的 C.ID.L 表示形式和状态。状态可能为 GOOD（良好）、FAILED（故障）或 HSP（磁盘已设置为备用磁盘）。

如果指定了卷或控制器，则只会拍摄指定的卷或控制器的信息快照。

-S -g *disk controller*

拍摄指定的磁盘的信息快照。

-h

输出用法字符串。

示例

示例1 创建 RAID 配置

以下命令在控制器 0 上创建 10G 的 RAID 0 卷，条带大小设置为 64k：

```
# raidctl -C "0.0.0 0.2.0" -r 0 -z 10g -s 64k 0
```

以下命令在控制器 2 上创建 RAID 1 卷：

```
# raidctl -C "0.0.0 1.1.0" -r 1 2
```

以下命令在控制器 2 上创建 RAID 5 卷：

```
# raidctl -C "0.0.0 0.1.0 0.2.0" -r 5 2
```

以下命令在控制器 0 上创建 RAID 10 卷：

```
# raidctl -C "(0.0.0 0.1.0)(0.2.0 0.3.0)" -r 10 0
```

以下命令在控制器 0 上创建 RAID 50 卷：

```
# raidctl -C "(0.0.0 0.1.0 0.2.0)(0.3.0 0.4.0 0.5.0)" -r 50 0
```

示例2 显示 RAID 配置

以下命令显示所有可用的控制器、卷和磁盘：

```
# raidctl -l
```

```
Controller: 0
```

```
Controller: 2
```

```
Volume: c2t0d0
```

```
Disk: 0.0.0
```

```
Disk: 0.1.0
```

```
Disk: 0.2.0
```

```
Disk: 0.3.0(HSP)
```

以下命令显示控制器 2 的相关信息：

示例2 显示 RAID 配置 (续)

```
# raidctl -l 2

Controller      Type              Fw_version
-----
c2              LSI 1030         1.03.39.00
```

以下命令显示指定卷的相关信息：

```
# raidctl -l c2t0d0

Volume      Sub      Size      Stripe  Status  Cache  RAID
              Disk      Size      Size
-----
c2t0d0              10240M  64K      OPTIMAL ON      RAID5
              0.0.0   5120M      GOOD
              0.1.0   5120M      GOOD
              0.2.0   5120M      GOOD
```

以下命令显示控制器 0 上的磁盘 0.0.0 的相关信息：

```
# raidctl -l -g 0.0.0 0

Disk  Vendor  Product              Firmware Capacity      Status  HSP
-----
0.0.0  HITACHI H101473SCSUN72G SQ02    68.3G      GOOD   N/A
GUID:2000000cca02536c
```

示例3 删除 RAID 配置

以下命令将删除卷：

```
# raidctl -d c0t0d0
```

示例4 更新控制器上的闪存映像

以下命令更新控制器 0 上的闪存映像：

```
# raidctl -F lsi_image.fw 0
```

示例5 设置或取消设置热备用磁盘

以下命令将控制器 2 上的磁盘 0.3.0 设置为全局热备用磁盘：

```
# raidctl -a set -g 0.3.0 2
```

以下命令将控制器 2 上的磁盘 0.3.0 设置为卷 c2t0d0 的本地热备用磁盘：

```
# raidctl -a set -g 0.3.0 c2t0d0
```

示例5 设置或取消设置热备用磁盘 (续)

以下命令将控制器 2 上的磁盘 0.3.0 从全局热备用磁盘转换为标准磁盘：

```
# raidctl -a unset -g 0.3.0 2
```

以下命令取消将磁盘 0.3.0 用作卷 c2t0d0 的本地热备用磁盘：

```
# raidctl -a unset -g 0.3.0 c2t0d0
```

示例6 设置卷的属性

以下命令将卷的写入策略设置为 "off"：

```
# raidctl -a set -p "wp=off" c0t0d0
```

示例7 使用 -c 选项创建卷

以下命令创建 RAID 1 卷：

```
# raidctl -c c0t0d0 c0t1d0
```

以下命令创建 RAID 0 卷：

```
# raidctl -c -r 0 c0t1d0 c0t2d0 c0t3d0
```

示例8 拍摄 RAID 配置的快照

以下命令拍摄所有 RAID 设备的快照：

```
# # raidctl -S
```

```
1 "LSI 1030"
c1t1d0 2 0.2.0 0.3.0 1 DEGRADED
0.2.0 GOOD
0.3.0 FAILED
```

以下命令拍摄关于卷 c1t0d0 的快照：

```
# raidctl -S c1t0d0
```

```
c1t0d0 2 0.0.0 0.1.0 1 OPTIMAL
```

以下命令拍摄关于控制器 1 上的磁盘 0.1.0 的快照：

```
# raidctl -S -g 0.1.0 1
```

```
0.1.0 GOOD
```

退出状态

将返回下列退出值：

```
0
```

成功完成。

- 1 命令行输入无效或权限被拒绝。
- 2 请求操作失败。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）

另请参见 [attributes\(5\)](#)、[mpt\(7D\)](#)
[《Oracle Solaris 管理：常见任务》](#)

警告 如果计划使用 Solaris 多路径 I/O 功能（也称为 MPxIO），请不要在内部 SAS 磁盘上创建 RAID 卷。在 Solaris 多路径下创建新的 RAID 卷会为根设备指定一个新的 GUID，这个新的 GUID 与现有设备的 GUID 不匹配。这会导致引导失败，因为 `/etc/vfstab` 中的根设备条目不匹配。

附注 在使用 `mpt` 驱动程序和 LSI RAID 控制器的系统上，`-z` 选项不受支持。

引用名	ramdiskadm – administer ramdisk pseudo device
用法概要	<pre>/usr/sbin/ramdiskadm -a <i>name size</i> [g m k b] /usr/sbin/ramdiskadm -d <i>name</i> /usr/sbin/ramdiskadm</pre>
描述	The <code>ramdiskadm</code> command administers ramdisk(7D), the ramdisk driver. Use <code>ramdiskadm</code> to create a new named ramdisk device, delete an existing named ramdisk, or list information about existing ramdisks. Ramdisks created using <code>ramdiskadm</code> are not persistent across reboots.
选项	The following options are supported: -a <i>name size</i> Create a ramdisk named <i>name</i> of size <i>size</i> and its corresponding block and character device nodes. <i>name</i> must be composed only of the characters a-z, A-Z, 0-9, _ (underbar), and - (hyphen), but it must not begin with a hyphen. It must be no more than 32 characters long. Ramdisk names must be unique. The size can be a decimal number, or, when prefixed with 0x, a hexadecimal number, and can specify the size in bytes (no suffix), 512-byte blocks (suffix b), kilobytes (suffix k), megabytes (suffix m) or gigabytes (suffix g). The size of the ramdisk actually created might be larger than that specified, depending on the hardware implementation. If the named ramdisk is successfully created, its block device path is printed on standard out. -d <i>name</i> Delete an existing ramdisk of the name <i>name</i>. This command succeeds only when the named ramdisk is not open. The associated memory is freed and the device nodes are removed. You can delete only ramdisks created using <code>ramdiskadm</code>. It is not possible to delete a ramdisk that was created during the boot process. Without options, <code>ramdiskadm</code> lists any existing ramdisks, their sizes (in decimal), and whether they can be removed by <code>ramdiskadm</code> (see the description of the -d option, above).
示例	示例 1 Creating a 2MB Ramdisk Named mydisk <pre># ramdiskadm -a mydisk 2m /dev/ramdisk/mydisk</pre>

示例 2 Listing All Ramdisks

```
# ramdiskadm
Block Device          Size  Removable
/dev/ramdisk/miniroot 134217728 No
/dev/ramdisk/certfs    1048576  No
/dev/ramdisk/mydisk    2097152  Yes
```

退出状态 ramdiskadm returns the following exit values:

- 0 Successful completion.
- >0 An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [attributes\(5\)](#), [ramdisk\(7D\)](#)

附注 The abilities of `ramdiskadm` and the privilege level of the person who uses the utility are controlled by the permissions of `/dev/ramdiskctl`. Read access allows query operations, for example, listing device information. Write access is required to do any state-changing operations, for example, creating or deleting ramdisks.

As shipped, `/dev/ramdiskctl` is owned by `root`, in group `sys`, and mode `0644`, so all users can do query operations but only `root` can perform state-changing operations. An administrator can give write access to non-privileged users, allowing them to add or delete ramdisks. However, granting such ability entails considerable risk; such privileges should be given only to a trusted group.

引用名 rcapadm – configure resource capping daemon

用法概要 rcapadm

```
rcapadm [ [-n] -E | -D]
          [-i interval=value,...,interval=value] [-c percent]
          [-z zonename -m maxvalue]
```

描述

The `rcapadm` command allows a user with the privileges described below to configure various attributes of the resource capping daemon. If used without arguments, `rcapadm` displays the current status of the resource capping daemon if it has been configured. See [rcapd\(1M\)](#) for more information.

In the current release of the Solaris operating environment, `rcapadm` is available to users with all privileges and to users who have the Process Management profile in their list of profiles. The System Administrator role includes the Process Management profile.

选项	-c percent	Set the minimum physical memory utilization for memory cap enforcement. Caps will not be enforced until the physical memory available to processes is low. The <i>percent</i> value should be in the range 0 to 100. The minimum (and default) value is 0, which means that memory caps are always enforced.
	-D	Disable the resource capping daemon so that it will not be started when the system is booted. Also stop the resource capping daemon now, if the -n option is not specified and it is currently running.
	-E	Enable the resource capping daemon so that it will be started each time the system is booted. Also start the resource capping daemon now, if the -n option is not specified and it is not currently running.
	-i interval=value,...,interval=value	Set intervals for various periodic operations performed by <code>rcapd</code> . All intervals are specified in seconds. You can set the following intervals:
	scan	The interval at which <code>rcapd</code> scans for new processes. The default scan interval is every 15 seconds. The minimum value is 1 second.
	sample	The interval of process resident set size sampling. The default sample interval is every 5 seconds. The minimum value is 1 second.

report The interval at which various paging statistics are updated by rcapd, in seconds. These statistics can be viewed by using `rcapstat(1SRM)`. The default reporting interval is every 5 seconds. When the interval is set to 0, statistics will not be updated.

注 – Paging refers to the act of relocating portions of memory, called pages, to or from physical memory. rcapd pages out the most infrequently used pages.

config The reconfiguration interval, in seconds. At each reconfiguration event, rcapd checks its configuration file for updates, and scans the project databases for new project caps. The default reconfiguration interval is every 60 seconds. The minimum interval is 0. When the interval is set to 0, no periodic reconfiguration occurs, although the running daemon can still be reconfigured by sending it `SIGHUP`.

-m maxvalue Used in conjunction with the **-z** option. Specifies a value for `rcap.max-rss`, a dynamically-set cap on the usage of physical memory for the zone specified by **-z**. You can apply a scale (K, M, G, T) to the value you specify. K means kilobyte; M, megabyte; G, gigabyte; and T, terabyte. For example, `100M` is 100 megabytes.

To remove an existing cap, specify `0M`.

-n Do not affect the running state of the resource capping daemon when enabling or disabling it.

-z zonename Used in conjunction with the **-m** option. Specifies the zone for which you are dynamically specifying a cap on physical memory usage (using **-m**).

注 – To set a persistent cap on memory usage within a zone, use [zonecfg\(1M\)](#).

示例

```
示例 1    Configuring the Resource Capping Daemon with Immediate Enforcement
# rcapadm -E -i scan=15,sample=5,report=5,config=60 -c 0
```

示例 2 Specifying a Resource Cap for a Zone

The command shown below specifies the maximum amount of memory that can be consumed by a specified zone. Note that this value lasts only until the next reboot. To set a persistent cap, use [zonecfg\(1M\)](#).

```
# rcapadm -z testzone -m 512M
```

退出状态

The following exit values are returned:

- 0 Successful completion. The modifications to the current configuration were valid and made successfully.
- 1 An error occurred. A fatal error occurred either in obtaining or modifying the resource capping configuration.
- 2 Invalid command-line options were specified.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/resource-mgmt/resource-caps
Interface Stability	Committed

The -z and -m options are committed interfaces.

另请参见

[rcapstat\(1\)](#), [rcapd\(1M\)](#), [zonecfg\(1M\)](#), [project\(4\)](#), [attributes\(5\)](#), [zones\(5\)](#)

“Physical Memory Control Using the Resource Capping Daemon” in 《Oracle Solaris 管理：Oracle Solaris Zones、Oracle Solaris 10 Zones 和资源管理》

引用名	rcapd – resource cap enforcement daemon
用法概要	rcapd [-d]
描述	The rcapd daemon enforces resource caps on collections of processes. Per-project and per-zone physical memory caps are supported. For information about projects, see project(4). For zones information, see zones(5) When the resident set size (RSS) of a collection of processes exceeds its cap, rcapd takes action and reduces the RSS of the collection. The virtual memory system divides physical memory into segments known as pages. To read data from a file into memory, the virtual memory system reads in individual pages. To reduce resource consumption, the daemon can page out, or relocate, infrequently used pages to an area outside of physical memory. In the project file, caps are defined for projects that have positive values for the following project attribute: <pre>rcap.max-rss The total amount of physical memory, in bytes, that is available to the project's member processes</pre> See project(4) for a description of project attributes. For a system with one or more zones, you can dynamically set the rcap.max-rss value for a zone with rcapadm(1M). To set a persistent cap on memory usage within a zone, you use zonecfg(1M). You configure rcapd through the use of rcapadm(1M). The daemon can be monitored with rcapstat(1). Configuration changes are incorporated into rcapd by sending it SIGHUP (see kill(1)), or according to the configuration interval (see rcapadm(1M)).
选项	The following option is supported: <pre>-d Enable debug mode. Messages are displayed on the invoking user's terminal.</pre>
示例	示例 1 Setting Resident Set Size Cap Attribute The following line in the /etc/project database sets an RSS cap of 1073741824 bytes for a project named foo. <pre>foo:100::foo,root::rcap.max-rss=10737418240</pre>
退出状态	The following exit values are returned: <pre>0 Successful completion. 1 An error occurred. 2 Invalid command-line options were specified.</pre>

文件 /etc/project Project database.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/resource-mgmt/resource-caps
Interface Stability	Committed

另请参见 [rcapstat\(1\)](#), [svcs\(1\)](#), [rcapadm\(1M\)](#), [zonecfg\(1M\)](#), [svcadm\(1M\)](#), [project\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [zones\(5\)](#)

“Physical Memory Control Using the Resource Capping Daemon” in 《[Oracle Solaris 管理：Oracle Solaris Zones、Oracle Solaris 10 Zones 和资源管理](#)》

附注 If killed with SIGKILL, rcapd can leave processes in a stopped state. Use SIGTERM to cause rcapd to terminate properly.

A collection's RSS can exceed its cap for some time before the cap is enforced, even if sufficient pageable memory is available. This period of time can be reduced by shortening the RSS sampling interval with rcapadm.

The rcapd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/system/rcap:default`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	<code>rctladm</code> – display or modify global state of system resource controls
用法概要	<code>rctladm [-lu] [-e <i>action</i>] [-d <i>action</i>] [<i>name...</i>]</code>
描述	The <code>rctladm</code> command allows the examination and modification of active resource controls on the running system. An instance of a resource control is referred to as an <i>rctl</i>. See setrctl(2) for a description of an <i>rctl</i>; see resource_controls(5) for a list of the <i>rctls</i> supported in the current release of the Solaris operating system. Logging of <i>rctl</i> violations can be activated or deactivated system-wide and active <i>rctls</i> (and their state) can be listed. An <code>rctladm</code> command without options is the equivalent of an <code>rctladm</code> with the <code>-l</code> option. See the description of <code>-l</code> below.
选项	The following options are supported: <code>-d <i>action</i></code> <code>-e <i>action</i></code> Disable (<code>-d</code>) or enable (<code>-e</code>) the global action on the specified <i>rctls</i>. If no <i>rctl</i> is specified, no action is taken and an error status is returned. You can use the special token <code>all</code> with the disable option to deactivate all global actions on a resource control. You can set the <code>syslog</code> action to a specific degree by assigning a severity level. To do this, specify <code>syslog=<i>level</i></code>, where <i>level</i> is one of the string tokens given as valid severity levels in syslog(3C). You can omit the common <code>LOG_</code> prefix on the severity level. Note that not all <i>rctls</i> support the <code>syslog</code> action. See resource_controls(5). <code>-l</code> List information about <i>rctls</i>. The name, global event actions and statuses, and global flags are displayed. If one or more name operands are specified, only those <i>rctls</i> matching the names are displayed. <code>-u</code> Configure resource controls based on the contents of <code>/etc/rctladm.conf</code>. Any name operands are ignored.
操作数	The following operands are supported: <i>name</i> The name of the <i>rctl</i> to operate on. Multiple <i>rctl</i> names can be specified. If no names are specified, and the list action has been specified, then all <i>rctls</i> are listed. If the enable or disable action is specified, one or more <i>rctl</i> names must be specified.
示例	示例 1 Activating System Logging for Specific Violations The following command activates system logging of all violations of <code>task.max-lwps</code>. <pre># rctladm -e syslog task.max-lwps #</pre>

示例 2 Examining the Current Status of a Specific Resource

The following command examines the current status of the `task.max-lwps` resource.

```
$ rctladm -l task.max-lwps
task.max-lwps          syslog=DEBUG
$
```

退出状态

The following exit values are returned:

0
Successful completion.

1
A fatal error occurred. A message is written to standard error to indicate each resource control for which the operation failed. The operation was successful for any other resource controls specified as operands.

2
Invalid command line options were specified.

文件

`/etc/rctladm.conf`
Each time `rctladm` is executed, it updates the contents of `rctladm.conf` with the current configuration.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[setrctl\(2\)](#), [getrctl\(2\)](#), [prctl\(1\)](#), [rctlblk_get_global_flags\(3C\)](#), [rctlblk_get_global_action\(3C\)](#), [attributes\(5\)](#), [resource_controls\(5\)](#)

附注

By default, there is no global logging of `rctl` violations.

引用名 rdate – set system date from a remote host

用法概要 rdate *hostname*

描述 rdate sets the local date and time from the *hostname* given as an argument. You must have the authorization `solaris.system.date` on the local system. Typically, rdate is used in a startup script.

rdate requests are responded to by the “time” service on the specified host. To enable the “time” service, use the following commands:

```
svcadm enable time:stream
svcadm enable time:dgram
```

用法 The rdate command is IPv6-enabled. See [ip6\(7P\)](#).

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-clients

另请参见 [inetd\(1M\)](#), [inetd.conf\(4\)](#), [attributes\(5\)](#), [ip6\(7P\)](#)

引用名 reboot – 重新启动操作系统

用法概要 /usr/sbin/reboot [-dlnq] [-f | -p] [boot_arguments]
 /usr/sbin/reboot [-f [-e environment] | -p] [-dlnq] [boot_arguments]

描述 reboot 实用程序可重新启动内核。PROM 监视器将内核装入内存，并将控制权转移给装入的内核。

在 x86 系统中，当指定了 -f 标志时，运行的内核会将下一个内核装入内存，然后将控制权转移给新装入的内核。此 reboot 形式如上面第二条用法概要所示。

虽然超级用户可以随时运行 reboot，但通常会先使用 shutdown(1M) 以警告登录用户服务即将丢失。有关详细信息，请参见 shutdown(1M)。

reboot 实用程序对磁盘执行 sync(1M) 操作，然后启动多用户重新引导。有关详细信息，请参见 init(1M)。在 x86 系统中，重新引导可能还会根据需要进行更新引导归档文件，以确保成功重新引导。

reboot 实用程序通常会将重新引导记录到系统日志守护进程 syslogd(1M) 中，并在登录记帐文件 /var/adm/wtmpx 中添加一条关机记录。如果提供了 -n 或 -q 选项，则会禁止这些操作。

通常，系统会在加电时或崩溃后进行重新引导。

选项 所支持的选项如下：

-d
 强制在重新引导前执行系统故障转储。有关配置系统故障转储的信息，请参见 dumpadm(1M)。

-e
 如果存在 -f 选项，将重新引导至指定的引导环境。
 次选项当前仅在 x86 系统中可用。

-f
 对于 x86 系统：
 快速重新引导，跳过固件和引导装载程序。运行的内核会将新内核装入内存，同时将控制权转移给新装入的内核。如果要指定磁盘或内核参数，则必须在其他引导参数之前指定它们。

 对于 SPARC 系统：

 通过跳过某些 POST 测试加快重新引导过程。

 缺省情况下，启用服务 svc:/system/boot-config:default。它需要 solaris.system.shutdown 作为 action_authorization 和 value_authorization。当 config/fastreboot_default 属性设置为 true 时，reboot 的行为相当于 reboot -f。此属性的值可以使用 svccfg(1M) 和 svcadm(1M) 更改，以控制缺省的重新引导行为。

有关详细信息，请参见**示例**。

-l

阻止向系统日志守护进程 **syslogd(1M)** 发送有关谁执行了 **reboot** 的消息。

-n

避免调用 **sync(2)**，并且不将重新引导记录到 **syslogd(1M)** 或 **/var/adm/wtmpx** 中。内核仍会在重新引导之前尝试同步文件系统，除非也使用了 **-d** 选项。如果将选项 **-d** 和 **-n** 一起使用，内核将不会尝试同步文件系统。

-p

重新引导至 **PROM**。此标志可用于通过固件重新引导系统，而不会更改由 **system/boot-config** 服务中的 **config/fastreboot_default** 属性设置指定的缺省重新引导行为。

-p 和 **-f** 选项互相排斥。

-q

快速。以快速、断然的方式重新引导，而不先关闭正在运行的进程。

操作数

支持下列操作数：

boot_arguments

可选的 *boot_arguments* 为 **uadmin(2)** 函数指定参数，这些参数在重新启动时将传递给引导程序和内核。参数的形式和列表在 **boot(1M)** 和 **kernel(1M)** 手册页中有述。如果指定参数，参数间的空白将替换为单空格，除非针对 **shell** 将空白放在引号中。如果 *boot_arguments* 以连字符开头，则必须在其前面加上 **--** 分界符（两个连字符），以表示 **reboot** 参数列表结束。

示例

示例 1 将 **-r** 和 **-v** 参数传递给 **boot**

在下面的示例中，必须使用分界符 **--**（两个连字符）将 **reboot** 的选项与 **boot(1M)** 的参数分隔开来。

```
example# reboot -dl -- -rv
```

示例 2 使用特定磁盘和内核重新引导

下面的示例使用特定的磁盘和内核重新引导。

```
example# reboot disk1 kernel.test/unix
```

示例 3 快速重新引导

下面的示例使用 **-f** 选项执行快速重新引导。

如果启用了 **svc:/system/boot-config:default** 服务，并且 **config/fastreboot_default** 属性设置为 **true**，则可省略 **-f** 选项。

在 **x86** 系统中，以下命令将重新引导至 **GRUB**（请参见 **grub(5)**）菜单文件 **menu.lst** 中的缺省条目。

示例3 快速重新引导 (续)

```
example# reboot -f
```

以下命令将重新引导至另一个 UFS 根磁盘。

```
example# reboot -f -- '/dev/dsk/c1d0s0'
```

以下命令将重新引导至另一个 ZFS 根池。

```
example# reboot -f -- 'rpool/R00T/root2'
```

以下命令使用 -k 选项重新引导至同一磁盘上的 mykernel。

```
example# reboot -f -- '/platform/i86pc/mykernel/amd64/unix -k'
```

以下命令将重新引导至挂载在 /mnt 上的另一个根磁盘上的 mykernel。

```
example# reboot -f -- '/mnt/platform/i86pc/mykernel/amd64/unix -k'
```

以下命令将重新引导至另一个名为 second_root 的引导环境中的 /platform/i86pc/kernel/\$ISADIR/unix。

```
example# reboot -f -e second_root
```

以下命令使用 -kv 选项重新引导至同一内核。

```
example# reboot -f -- '-kv'
```

以下命令将禁用缺省的快速重新引导行为。

```
example# svccfg -s "system/boot-config:default" \
setprop config/fastreboot_default=false
example# svcadm refresh svc:/system/boot-config:default
```

以下命令将重新启用缺省的快速重新引导行为。

```
example# svccfg -s "system/boot-config:default" \
setprop config/fastreboot_default=true
example# svcadm refresh svc:/system/boot-config:default
```

示例4 重新引导至特定 GRUB 菜单

以下命令将重新引导至 GRUB 菜单中的条目 2。

```
example# bootadm list-menu
the location for the active GRUB menu is: /rpool/boot/grub/menu.lst
default 0
timeout 10
0 zfsbe1
1 zfsbe1 failsafe
2 zfsbe2
3 zfsbe2 Solaris xVM
```

示例 4 重新引导至特定 GRUB 菜单 (续)

```
4 zfsbe2 failsafe
example# reboot 2
```

文件

/var/adm/wtmpx
登录记帐文件

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见

[mdb\(1\)](#)、[boot\(1M\)](#)、[dumpadm\(1M\)](#)、[fsck\(1M\)](#)、[halt\(1M\)](#)、[init\(1M\)](#)、[kernel\(1M\)](#)、[shutdown\(1M\)](#)

附注

`reboot` 实用程序不执行 `/etc/rcnum.d` 中的脚本，也不执行 [inittab\(4\)](#) 中的关机操作。为确保完全关闭系统服务，请使用 [shutdown\(1M\)](#) 或 [init\(1M\)](#) 重新引导 Solaris 系统。

引用名	rem_drv – remove a device driver from the system
用法概要	rem_drv [-b <i>basedir</i>] [-C] [-n] <i>device_driver</i>
描述	The <code>rem_drv</code> command informs the system that the device driver <i>device_driver</i> is no longer valid. If possible, <code>rem_drv</code> unloads <i>device_driver</i> from memory. <code>rem_drv</code> also updates the system driver configuration files. If <code>rem_drv</code> has been executed, the next time the system is rebooted it automatically performs a reconfiguration boot (see kernel(1M)).
选项	The following options are supported: -b <i>basedir</i> Sets the path to the root directory of the diskless client. Used on the server to execute <code>rem_drv</code> for a client. The client machine must be rebooted to unload the driver. - 注 – The root file system of any non-global zones must not be referenced with the -b option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See zones(5). -C Remove dangling device attribute nodes bound to the driver being removed. This causes any device ownership or permissions customizations made to any node not to be preserved if the driver is added back. Recommended for use when reprovisioning a machine from one configuration or use to another where past administrative customizations might not be desired. -n Do not try to detach and unload <i>device_driver</i>, just modify the system configuration files for that driver.
示例	示例 1 Removing the sd Driver The following example removes the sd driver from use: <pre>example% rem_drv sd</pre> 示例 2 Removing a Diskless Client The following example removes the driver from the sun1 diskless client. The driver is not uninstalled or unloaded until the client machine is rebooted. <pre>example% rem_drv -b /export/root/sun1 sd</pre> Note the caveat on the use of the -b option in the description of that option, above.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [add_drv\(1M\)](#), [kernel\(1M\)](#), [update_drv\(1M\)](#), [attributes\(5\)](#), [zones\(5\)](#), [devfs\(7FS\)](#)

引用名

remove_allocatable – remove entries from allocation databases

用法概要

```
/usr/sbin/remove_allocatable [-f] -n name
```

```
/usr/sbin/remove_allocatable [-f] [-d] -t dev-type
```

描述

remove_allocatable removes entries of user allocatable devices from the device allocation mechanism. remove_allocatable also removes entries of some non-allocatable devices, such as printers, whose label range is managed by the mechanism.

选项

The following options are supported:

-d

Removes system-supplied default attributes of the device type that is specified with -t.

-f

Force the removal of an entry. remove_allocatable exits with an error if this option is not specified when an entry with the specified device name no longer exists.

-n name

Removes the entry for the device name.

-t dev-type

Removes devices of type dev-type.

退出状态

When successful, remove_allocatable returns an exit status of 0 (true). remove_allocatable returns a nonzero exit status in the event of an error. The exit codes are as follows:

1

Invocation syntax error

2

Unknown system error

3

Device name or dev-type not found. This error occurs only when the -f option is not specified.

4

Permission denied. User does not have DAC or MAC access to database.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/trusted
Interface Stability	See below.

The invocation is Uncommitted. The options are Uncommitted. The output is Not-an-Interface.

另请参见

[allocate\(1\)](#), [deallocate\(1\)](#), [add_allocatable\(1M\)](#), [attributes\(5\)](#), [device_clean\(5\)](#)

附注

The functionality described on this manual page is available only if the system is configured with Trusted Extensions.

引用名	removef – remove a file from software database
用法概要	<pre>removef [[-M] -R <i>root_path</i>] [-V <i>fs_file</i>] <i>pkginst path</i>...</pre> <pre>removef [[-M] -R <i>root_path</i>] [-V <i>fs_file</i>] -f <i>pkginst</i></pre>
描述	removef informs the system that the user, or software, intends to remove a pathname. Output from removef is the list of input pathnames that may be safely removed (no other packages have a dependency on them).
选项	The following options are supported: -f After all files have been processed, removef should be invoked with the -f option to indicate that the removal phase is complete. -M Instruct removef not to use the <i>\$root_path/etc/vfstab</i> file for determining the client's mount points. This option assumes the mount points are correct on the server and it behaves consistently with Solaris 2.5 and earlier releases. -R <i>root_path</i> Define the full path name of a directory to use as the <i>root_path</i>. All files, including package system information files, are relocated to a directory tree starting in the specified <i>root_path</i>. The <i>root_path</i> may be specified when installing to a client from a server (for example, <i>/export/root/client1</i>). removef inherits the value of the <code>PKG_INSTALL_ROOT</code> environment variable. (See <code>ENVIRONMENT VARIABLES</code>, below.) If <code>PKG_INSTALL_ROOT</code> is set, such as when the -R option is used with pkgadd(1M) or pkgrm(1M) 注 – The root file system of any non-global zones must not be referenced with the -R option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See zones(5). -V <i>fs_file</i> Specify an alternative <i>fs_file</i> to map the client's file systems. For example, used in situations where the <i>\$root_path/etc/vfstab</i> file is non-existent or unreliable.
操作数	The following operands are supported: <i>path</i> The pathname to be removed. <i>pkginst</i> The package instance from which the pathname is being removed.

示例

```
示例1 Using removef
The following example uses the removef command in an optional pre-install script:

echo "The following files are no longer part of this package
    and are being removed."
removef $PKGINST /myapp/file1 /myapp/file2 |
while read pathname
do
    echo "$pathname"
    rm -f $pathname
done
removef -f $PKGINST || exit 2
```

环境变量

removef inherits the value of the following environment variable. This variable is set when [pkgadd\(1M\)](#) or [pkgrm\(1M\)](#)

PKG_INSTALL_ROOT
If present, defines the full path name of a directory to use as the system's PKG_INSTALL_ROOT path. All product and package information files are then looked for in the directory tree, starting with the specified PKG_INSTALL_ROOT path. If not present, the default system path of / is used.

退出状态

0
Successful completion.

>0
An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[pkginfo\(1\)](#), [pkgmk\(1\)](#), [pkgparam\(1\)](#), [pkgproto\(1\)](#), [pkgtrans\(1\)](#), [installf\(1M\)](#), [pkgadd\(1M\)](#), [pkgask\(1M\)](#), [pkgchk\(1M\)](#), [pkgrm\(1M\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

《Application Packaging Developer's Guide》

附注

Package commands are [largefile\(5\)](#)-aware. They handle files larger than 2 GB in the same way they handle smaller files. In their current implementations, [pkgadd\(1M\)](#), [pkgtrans\(1\)](#) and other package commands can process a datastream of up to 4 GB.

引用名	reparse – 重解析点守护进程
用法概要	/usr/lib/reparse/reparse
描述	reparse 守护进程处理对“重解析点”进行解释的内核调用，“重解析点”是在 Solaris SMB 和 NFS 文件服务器上支持 Microsoft DFS 引用和 NFS 引用的基础。只有超级用户或具有同等特权的用户可以运行此守护进程。此守护进程受 SMF 控制。
退出状态	0 守护进程启动成功。 >0 守护进程启动失败。 错误信息以 LOG_ERR 级别报告到 syslog 中。
文件	/usr/lib/reparse/*.so.1 reparse 的每服务插件。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	system/core-os
接口稳定性	Consolidation Private（整合专用）

另请参见 [nfsref\(1M\)](#)、[svcadm\(1M\)](#)、[syslogd\(1M\)](#)、[libreparse\(3LIB\)](#)、[attributes\(5\)](#)

附注 不要手动停止、启动或重新启动 reparse 守护进程。如果您需要更改此守护进程的状态，请使用以下命令：

```
# svcadm disable svc:/system/filesystem/reparse
# svcadm enable svc:/system/filesystem/reparse
# svcadm restart svc:/system/filesystem/reparse
```

有关其他信息，请参见 [svcadm\(1M\)](#)。

引用名	repquota – summarize quotas for a ufs file system				
用法概要	<pre>repquota [-v] filesystem... repquota -a [-v]</pre>				
描述	repquota prints a summary of the disk usage and quotas for the specified ufs file systems. The current number of files and amount of space (in kilobytes) is printed for each user along with any quotas created with edquota(1M). The <i>filesystem</i> must have the file quotas in its root directory. Only the super-user may view quotas which are not their own.				
选项	The following options are supported: -a Report on all mounted ufs file systems that have <i>rq</i> in the <i>mntopts</i> field of the <i>/etc/vfstab</i> file. -v Report quotas for all users, even those who do not consume resources.				
用法	See largefile(5) for the description of the behavior of repquota when encountering files greater than or equal to 2 Gbyte (2 ³¹ bytes).				
属性	See attributes(5) for descriptions of the following attributes:				
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>system/core-os</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	system/core-os
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	system/core-os				
另请参见	edquota(1M) , quota(1M) , quotacheck(1M) , quotaon(1M) , attributes(5) , largefile(5) , quotactl(7I)				

引用名

rmmount – removable media mounter for CD-ROM, Jaz drive, and others

用法概要

`/usr/sbin/rmmount [-D]`

描述

The rmmount utility is a removable media mounter that is executed by volume management whenever a removable medium, such as a CD-ROM is inserted. Removable media is managed by an application or a volume manager. rmmount can also be called by using [volrmmount\(1\)](#).

Upon insertion of a medium and following invocation of the [volcheck\(1\)](#) command, rmmount determines what type of file system (if any) is on that medium. If a file system is present, rmmount mounts the file system in one of the locations listed below.

For a CD-ROM or a DVD-ROM:

<code>/cdrom/cdrom0</code>	symbolic link to mounted CD-ROM in local CD-ROM drive
<code>/cdrom/CD-ROM_name</code>	mounted named CD-ROM
<code>/cdrom/CD-ROM_name/partition</code>	mounted named CD-ROM with partitioned file system
<code>/cdrom/unnamed_cdrom</code>	mounted unnamed CD-ROM

For a Zip drive:

<code>/rmdisk/zip0</code>	symbolic link to mounted Zip medium in local Zip drive
<code>/rmdisk/Zip_name</code>	mounted named Zip medium
<code>/rmdisk/Zip_name/partition</code>	mounted named Zip medium with partitioned file system
<code>/rmdisk/unnamed_zip</code>	mounted unnamed Zip medium

For a Jaz drive:

<code>/rmdisk/jaz0</code>	symbolic link to mounted Jaz medium in local Jaz drive
<code>/rmdisk/Jaz_name</code>	mounted named Jaz medium
<code>/rmdisk/Jaz_name/partition</code>	mounted named Jaz medium with partitioned file system
<code>/rmdisk/unnamed_Jaz</code>	mounted unnamed Jaz medium

For a generic “rmdisk” drive:

<code>/rmdisk/rmdisk0</code>	symbolic link to mounted removable medium in local removable medium drive
<code>/rmdisk/rmdisk_name</code>	mounted named removable medium
<code>/rmdisk/rmdisk_name/partition</code>	mounted named removable medium with partitioned file system

/rmdisk/unnamed_rmdisk mounted unnamed removable medium

If the media is read-only (for example, a CD-ROM), the file system is mounted read-only.

If a file system is not identified, rmmount does not mount a file system. See the 《Oracle Solaris 管理：常见任务》 for more information on the location of CD-ROM and other media without file systems.

If a file system type has been determined, it is then checked to see that it is “clean.” If the file system is “dirty,” fsck -p (see fsck(1M)) is run in an attempt to clean it. If fsck fails, the file system is mounted read-only.

After the mount is complete, “actions” associated with the media type are executed. These actions allow for the notification to other programs that new media are available.

Actions are executed in the order in which they appear in the configuration file. The action function can return either 1 or 0. If it returns 0, no further actions will be executed. This allows the function to control which applications are executed.

In order to execute an action, rmmount performs a dlopen(3C) on the shared object and calls the action function defined within it. The definition of the interface to actions can be found in /usr/include/rmmount.h.

File systems mounted by rmmount are always mounted with the nosuid flag set, thereby disabling setuid programs and access to block or character devices in that file system. Upon ejection, rmmount unmounts mounted file systems and executes actions associated with the media type. If a file system is “busy” (that is, it contains the current working directory of a live process), the ejection will fail.

选项
文件
属性

-D Turn on the debugging output from the rmmount dprintf calls.

/usr/lib/rmmount/*.so.1 shared objects used by rmmount.

See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/storage/media-volume-manager

另请参见

volcheck(1), volrmmount(1), fsck(1M), dlopen(3C), attributes(5)

《Oracle Solaris 管理：常见任务》

引用名

rmt – remote magtape protocol module

用法概要

/usr/sbin/rmt

描述

rmt is a program used by the remote dump and restore programs in manipulating a magnetic tape drive through an interprocess communication connection. rmt is normally started up with an [rexec\(3SOCKET\)](#) or [rcmd\(3SOCKET\)](#) call.

The rmt program accepts requests that are specific to the manipulation of magnetic tapes, performs the commands, then responds with a status indication. All responses are in ASCII and in one of two forms. Successful commands have responses of:

*A*number\n where *number* is an ASCII representation of a decimal number.

Unsuccessful commands are responded to with:

Error-number\n*error-message*\n where *error-number* is one of the possible error numbers described in [Intro\(3\)](#), and *error-message* is the corresponding error string as printed from a call to [perror\(3C\)](#).

The protocol consists of the following commands:

<i>S</i> \n	Return the status of the open device, as obtained with a MTIOCGET ioctl call. If the operation was successful, an “ack” is sent with the size of the status buffer, then the status buffer is sent (in binary).
<i>C</i> <i>device</i> \n	Close the currently open device. The <i>device</i> specified is ignored.
<i>I</i> <i>operation</i> \n <i>ncount</i> \n	Perform a MTIOCOP ioctl(2) command using the specified parameters. The parameters are interpreted as the ASCII representations of the decimal values to place in the <i>mt_op</i> and <i>mt_count</i> fields of the structure used in the ioctl call. When the operation is successful the return value is the <i>count</i> parameter.
<i>L</i> <i>offset</i> \n <i>nwhence</i> \n	Perform an lseek(2) operation using the specified parameters. The response value is returned from the lseek call.
<i>O</i> <i>device</i> \n <i>mode</i> \n	Open the specified <i>device</i> using the indicated <i>mode</i> . <i>device</i> is a full pathname, and <i>mode</i> is an ASCII representation of a decimal number suitable for passing to open(9E) . If a device is already open, it is closed before a new open is performed.
<i>R</i> <i>count</i> \n	Read <i>count</i> bytes of data from the open device. rmt performs the requested read(9E) and responds with <i>Acount-read</i> \n if the read was successful; otherwise an error in standard format is returned. If the read was successful, the data read is sent.
<i>W</i> <i>count</i> \n	Write data onto the open device. rmt reads <i>count</i> bytes from the connection, aborting if a premature EOF is encountered. The

response value is returned from the `write(9E)` call.

Any other command causes rmt to exit.

属性 See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-clients

另请参见 `ufsdump(1M)`, `ufsrestore(1M)`, `Intro(3)`, `ioctl(2)`, `lseek(2)`, `perror(3C)`, `rcmd(3SOCKET)`, `rexec(3SOCKET)`, `attributes(5)`, `mtio(7I)`, `open(9E)`, `read(9E)`, `write(9E)`

诊断 All responses are of the form described above.

已知问题 Do not use this for a remote file access protocol.

引用名 rmvolmgr – 能够识别 HAL 的可移除卷管理器

用法概要 /usr/sbin/rmvolmgr [-chns]

描述 rmvolmgr 命令是一个卷管理器，它能够自动挂载和卸载可移除的介质和可热插拔的存储器。缺省挂载点是 /media。

rmvolmgr 是 Solaris 操作系统附带的许多能够识别硬件抽象层 (Hardware Abstraction Layer, HAL) 的工具中的一种。请参见 [hald\(1M\)](#)。

可以同时运行 rmvolmgr 的多个实例。rmvolmgr 的系统实例在缺省情况下作为服务管理工具 (Service Management Facility, SMF) 服务运行（请参见 [smf\(5\)](#)）。其故障管理资源标识符 (fault management resource identifier, FMRI) 为：

```
svc:/system/filesystem/rmvolmgr
```

通过将 rmvolmgr 实例添加到 .xinitrc 文件或类似的会话启动脚本，您可以运行自己的 rmvolmgr 实例。在这种情况下，系统 rmvolmgr 实例将不会管理属于您（启动脚本的所有者）的卷。例如，如果一个用户登录到工作站控制台 (/dev/console) 并调用了自己的 rmvolmgr 实例，则该用户将成为在本地连接的设备的所有者，如 CD-ROM 驱动器和连接到本地 USB 或 FireWire 端口的设备。

除了在 /media 下挂载卷，rmvolmgr 还在 /cdrom 和 /rmdisk 下创建传统的符号链接。

通过在 /tmp/.removable 下维护通知文件，rmvolmgr 还提供与 CDE 可移除介质接口的向后兼容性。可以使用 -c 选项禁用该功能。

也可以将 -c 和 -n 选项指定为 SMF 属性。请参见“示例”。

选项 所支持的选项如下：

- c 禁用 CDE 兼容性。
- h 显示帮助信息并退出。
- n 不创建传统的挂载点符号链接。
- s 以系统实例模式进行调用。

示例 示例 1 使用 SMF 属性设置选项

下面的 [svccfg\(1M\)](#) 命令和子命令使用 SMF 属性来设置 -c 和 -n 选项。

```
example# svccfg
svc:> select rmvolmgr
svc:/system/filesystem/rmvolmgr> listprop rmvolmgr/*
rmvolmgr/legacy_mountpoints    boolean true
rmvolmgr/cde_compatible        boolean true
svc:/system/filesystem/rmvolmgr> setprop \
    rmvolmgr/legacy_mountpoints=false
svc:/system/filesystem/rmvolmgr> setprop rmvolmgr/cde_compatible=false
```


示例1 使用 SMF 属性设置选项 (续)

```
svc:/system/filesystem/rmvolmgr> exit
example#
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/storage/media-volume-manager
接口稳定性	Volatile (可变)

另请参见 [hald\(1M\)](#)、[svccfg\(1M\)](#)、[attributes\(5\)](#)、[hal\(5\)](#)、[smf\(5\)](#)

引用名 rndc – name server control utility

用法概要 rndc [-V] [-b *src-addr*] [-c *config-file*] [-k *key-file*] [-s *server*]
 [-p *port*] [-y *key_id*] *command*

描述 The rndc utility controls the operation of a name server. It supersedes the ndc utility that was provided in previous BIND releases. If rndc is invoked with no command line options or arguments, it prints a short summary of the supported commands and the available options and their arguments.

The rndc utility communicates with the name server over a TCP connection, sending commands authenticated with digital signatures. The only supported authentication algorithm in the current versions of rndc and [named\(1M\)](#) is HMAC-MD5, which uses a shared secret on each end of the connection. This algorithm provides TSIG-style authentication for the command request and the name server's response. All commands sent over the channel must be signed by a *key_id* known to the server.

The rndc utility reads a configuration file to determine how to contact the name server and decide what algorithm and key it should use.

选项 The following options are supported:

-b *source-address*

Use *source-address* as the source address for the connection to the server. Multiple instances are permitted to allow setting of both the IPv4 and IPv6 source addresses.

-c *config-file*

Use *config-file* as the configuration file instead of the default `/etc/rndc.conf`.

-k *key-file*

Use *key-file* as the key file instead of the default, `/etc/rndc.key`. The key in `/etc/rndc.key` is used to authenticate commands sent to the server if the *config-file* does not exist.

-s *server*

The *server* argument is the name or address of the server that matches a server statement in the configuration file for rndc. If no server is supplied on the command line, the host named by the default-server clause in the options statement of the rndc configuration file is used.

-p *port*

Send commands to TCP port *port* instead of BIND 9's default control channel port, 953.

-V

Enable verbose logging.

-y *key_id*

Use the key *key_id* from the configuration file. The *key_id* argument must be known by named with the same algorithm and secret string for control message validation to succeed. If no *key_id* is specified, rndc will first look for a key clause in the server statement of the

server being used, or if no server statement is present for that host, then the default-key clause of the options statement. The configuration file contains shared secrets that are used to send authenticated control commands to name servers. It should therefore not have general read or write access.

For the complete set of commands supported by rndc, see the BIND 9 Administrator Reference Manual or run rndc without arguments to see its help message.

Limitations The rndc utility does not support all the commands of the BIND 8 ndc utility.

There is no way to provide the shared secret for a *key_id* without using the configuration file.

Several error messages tend toward the cryptic.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	network/dns/bind
Interface Stability	Volatile

另请参见 [named\(1M\)](#), [rndc-confgen\(1M\)](#), [named.conf\(4\)](#), [rndc.conf\(4\)](#), [attributes\(5\)](#)

See the BIND 9 Administrator's Reference Manual. As of the date of publication of this man page, this document is available at <https://www.isc.org/software/bind/documentation>.

引用名	<code>rndc-confgen</code> – rndc key generation tool
用法概要	<code>rndc-confgen [-ah] [-b <i>keysize</i>] [-c <i>keyfile</i>] [-k <i>keyname</i>] [-p <i>port</i>] [-r <i>randomfile</i>] [-s <i>address</i>] [-t <i>chrootdir</i>] [-u <i>user</i>]</code>
描述	The <code>rndc-confgen</code> utility generates configuration files for rndc(1M) . This utility can be used as a convenient alternative to writing by hand the <code>rndc.conf(4)</code> file and the corresponding controls and key statements in <code>named.conf</code> . It can also be run with the <code>-a</code> option to set up a <code>rndc.key</code> file and avoid altogether the need for a <code>rndc.conf</code> file and a controls statement.
选项	The following options are supported: <code>-a</code> Perform automatic rndc configuration. This option creates a file <code>rndc.key</code> in <code>/etc</code> (or however <code>sysconfdir</code> was specified when BIND was built) that is read by both <code>rndc</code> and named(1M) on startup. The <code>rndc.key</code> file defines a default command channel and authentication key allowing <code>rndc</code> to communicate with <code>named</code> with no further configuration. Running <code>rndc-confgen</code> with <code>-a</code> specified allows BIND 9 and <code>rndc</code> to be used as drop-in replacements for BIND 8 and <code>ndc</code>, with no changes to the existing BIND 8 <code>named.conf</code> file. If a more elaborate configuration than that generated by <code>rndc-confgen -a</code> is required, for example if <code>rndc</code> is to be used remotely, you should run <code>rndc-confgen</code> without the <code>-a</code> option and set up <code>rndc.conf</code> and <code>named.conf</code> files, as directed. <code>-b <i>keysize</i></code> Specify the size of the authentication key in bits. The <i>keysize</i> argument must be between 1 and 512 bits; the default is 128. <code>-c <i>keyfile</i></code> Used with the <code>-a</code> option to specify an alternate location for <code>rndc.key</code>. <code>-h</code> Print a short summary of the options and arguments to <code>rndc-confgen</code>. <code>-k <i>keyname</i></code> Specify the key name of the <code>rndc</code> authentication key. The <i>keyname</i> argument must be a valid domain name. The default is <code>rndc-key</code>. <code>-p <i>port</i></code> Specify the command channel port where <code>named</code> listens for connections from <code>rndc</code>. The default is 953. <code>-r <i>randomfile</i></code> Specify a source of random data for generating the authorization. By default, <code>/dev/random</code> is used. The <i>randomdev</i> argument specifies the name of a character device or file containing random data to be used instead of the default. The special value <code>keyboard</code> indicates that keyboard input should be used.

- s *address*
Specify the IP address where named listens for command channel connections from rndc. The default is the loopback address 127.0.0.1.
- t *chrootdir*
Used with the -a option to specify a directory where named will run after the root directory is changed with [chroot\(2\)](#). An additional copy of the rndc.key will be written relative to this directory so that it will be found by the named in the new directory.
- u *user*
Used with the -a option to set the owner of the rndc.key file generated. If -t is also specified only the file in the chroot area has its owner changed.

示例

示例 1 Create Automatic rndc Configuration

The following command creates an automatic rndc configuration, so that rndc can be used immediately.

```
# rndc-confgen -a
```

示例 2 Print a Sample rndc.conf File

The following command prints a sample rndc.conf file with corresponding controls and key statements. These statements can subsequently be manually inserted in the file named.conf.

```
# rndc-confgen
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	network/dns/bind
Interface Stability	Volatile

另请参见

[chroot\(2\)](#), [named\(1M\)](#), [rndc\(1M\)](#), [rndc.conf\(4\)](#), [attributes\(5\)](#)

See the BIND 9 Administrator's Reference Manual. As of the date of publication of this man page, this document is available at <https://www.isc.org/software/bind/documentation>.

引用名

roleadd – administer a new role account on the system

用法概要

```
roleadd [-c comment] [-d dir] [-e expire] [-f inactive]
        [-g group] [-G group [, group...]] [-m [-k skel_dir]]
        [-u uid [-o]] [-s shell] [-S repository]
        [-A authorization [, authorization...]] [-K key=value] role

roleadd -D [-b base_dir] [-e expire] [-f inactive]
        [-g group] [-A authorization [, authorization...]]
        [-P profile [, profile...]] [-K key=value]]
```

描述

roleadd adds a role entry to the `passwd` and `shadow` and `user_attr` databases specified by the `-S` option. The default repository is `files`. The `-A` and `-P` options respectively assign authorizations and profiles to the role. Roles cannot be assigned to other roles. The `-K` option adds a `key=value` pair to `user_attr` for a role. Multiple `key=value` pairs can be added with multiple `-K` options.

roleadd also creates supplementary group memberships for the role (`-G` option) and creates the home directory (`-m` option) for the role if requested. The new role account remains locked until the [passwd\(1\)](#) command is executed.

Specifying roleadd `-D` with the `-g`, `-b`, `-f`, `-e`, or `-K` option (or any combination of these option) sets the default values for the respective fields. See the `-D` option. Subsequent roleadd commands without the `-D` option use these arguments.

The system file entries created with this command have a limit of 512 characters per line. Specifying long arguments to several options can exceed this limit.

The role (`role`) field accepts a string of no more than eight bytes consisting of characters from the set of alphabetic characters, numeric characters, period (`.`), underscore (`_`), and hyphen (`-`). The first character should be alphabetic and the field should contain at least one lower case alphabetic character. A warning message is written if these restrictions are not met. A future Solaris release might refuse to accept role fields that do not meet these requirements.

The `role` field must contain at least one character and must not contain a colon (`:`) or a newline (`\n`).

An administrator must be granted the User Management Profile to be able to create a new role. The authorizations required to set the various fields in `passwd`, `shadow` and `user_attr` can be found in [passwd\(4\)](#), [shadow\(4\)](#), and [user_attr\(4\)](#). The authorizations required to assign groups can be found in [group\(4\)](#).

选项

The following options are supported:

`-A authorization`

One or more comma separated authorizations defined in [auth_attr\(4\)](#). Only a user or role who has grant rights to the authorization can assign it to an account

-b *base_dir*

The default base directory for the system if *-d dir* is not specified. *base_dir* is concatenated with the account name to define the home directory. If the *-m* option is not used, *base_dir* must exist.

-c *comment*

Any text string. It is generally a short description of the role. This information is stored in the role's passwd entry.

-d *dir*

The home directory of the new role. It defaults to *base_dir/account_name*, where *base_dir* is the base directory for new login home directories and *account_name* is the new role name. This creates an *auto_home* entry for the user.

The argument to the option can be specified as *server:dir* where *server* is the hostname of the machine on which the home directory resides and *dir* is the path to the user's home directory. If the server is a remote host then the home directory needs to be created on the remote host for the system to mount it, when the user logs in. If no server name is specified then the home directory will be created on the host where the command is executed, when the *-m* option is used.

-D

Display the default values for *group*, *base_dir*, *skel_dir*, *shell*, *inactive*, *expire* and *key=value* pairs. When used with the *-g*, *-b*, *-f*, or *-K*, options, the *-D* option sets the default values for the specified fields. The default values are:

```
group
    other (GID of 1)
```

```
base_dir
    /export/home
```

```
skel_dir
    /etc/skel
```

```
shell
    /bin/pfsh
```

```
inactive
    0
```

```
expire
    Null
```

```
auths
    Null
```

```
profiles
    Null
```

key=value (pairs defined in [user_attr\(4\)](#))
not present

-e *expire*

Specify the expiration date for a role. After this date, no user is able to access this role. The expire option argument is a date entered using one of the date formats included in the template file `/etc/datemsk`. See [getdate\(3C\)](#).

If the date format that you choose includes spaces, it must be quoted. For example, you can enter `10/6/90` or `October 6, 1990`. A null value (`" "`) defeats the status of the expired date. This option is useful for creating temporary roles.

-f *inactive*

The maximum number of days allowed between uses of a role ID before that ID is declared invalid. Normal values are positive integers. A value of `0` defeats the status.

-g *group*

An existing group's integer ID or character-string name. Without the `-D` option, it defines the new role's primary group membership and defaults to the default group. You can reset this default value by invoking `roleadd -D -g group`.

-G *group*

An existing group's integer ID or character-string name. It defines the new role's supplementary group membership. Duplicates between *group* with the `-g` and `-G` options are ignored. No more than `NGROUPS_MAX` groups can be specified.

-k *skel_dir*

A directory that contains skeleton information (such as `.profile`) that can be copied into a new role's home directory. This directory must already exist. The system provides the `/etc/skel` directory that can be used for this purpose.

-K *key=value*

A *key=value* pair to add to the role's attributes. Multiple `-K` options can be used to add multiple *key=value* pairs. The generic `-K` option with the appropriate key can be used instead of the specific implied key options (`-A` and `-P`). See [user_attr\(4\)](#) for a list of valid *key=value* pairs. The “type” key is not a valid key for this option. Keys can not be repeated.

-m

Create the new role's home directory if it does not already exist. If the directory already exists, it must have read, write, and execute permissions by *group*, where *group* is the role's primary group. If the server name specified to the `-d` option is a remote host then the system will not attempt to create the home directory.

If the directory does not already exist, a new ZFS dataset will be created. In the global zone, the dataset is created as `rpool/export/home/rolename`. For non-global zones, the dataset will be created as `ROOT-dataset/export/home/rolename`. The mountpoint for the ZFS dataset is `/export/home/rolename` by default. If `-d path` is specified and it is a path on the local machine, the dataset will be mounted at the specified location. The role is delegated

permissions to create ZFS snapshots and promote them. The newly created dataset will inherit the encryption setting from its parent. If it is encrypted, the role is granted permission to change its wrapping key.

- o
This option allows a UID to be duplicated (non-unique).
- P *profile*
One or more comma-separated execution profiles defined in [prof_attr\(4\)](#).
- s *shell*
Full pathname of the program used as the user's shell on login. It defaults to an empty field causing the system to use `/bin/pfsh` as the default. The value of *shell* must be a valid executable file.
- S *repository*
The valid repositories are `files`, `ldap`. The repository specifies which name service will be updated. The default repository is `files`. When the repository is `files`, the authorizations, profiles, and roles can be present in other name service repositories and can be assigned to a user in the `files` repository. When the repository is `ldap`, all the assignable attributes must be present in the `ldap` repository.
- u *uid*
The UID of the new role. This UID must be a non-negative decimal integer below `MAXUID` as defined in `<sys/param.h>`. The UID defaults to the next available (unique) number above the highest number currently assigned. For example, if UIDs 100, 105, and 200 are assigned, the next default UID number is 201. (UIDs from 0-99 are reserved for possible use in future applications.)

文件

`/etc/datemsk`

`/etc/passwd`

`/etc/shadow`

`/etc/group`

`/etc/skel`

`/usr/include/limits.h`

`/etc/user_attr`

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Committed

另请参见

[passwd\(1\)](#), [pfexec\(1\)](#), [profiles\(1\)](#), [roles\(1\)](#), [users\(1B\)](#), [groupadd\(1M\)](#), [groupdel\(1M\)](#), [groupmod\(1M\)](#), [grpck\(1M\)](#), [logins\(1M\)](#), [pwck\(1M\)](#), [userdel\(1M\)](#), [usermod\(1M\)](#), [getdate\(3C\)](#), [auth_attr\(4\)](#), [group\(4\)](#), [passwd\(4\)](#), [prof_attr\(4\)](#), [shadow\(4\)](#), [user_attr\(4\)](#), [attributes\(5\)](#)

诊断

In case of an error, `roleadd` prints an error message and exits with a non-zero status.

The following indicates that `login` specified is already in use:

UX: `roleadd: ERROR: login is already in use. Choose another.`

The following indicates that the `uid` specified with the `-u` option is not unique:

UX: `roleadd: ERROR: uid uid is already in use. Choose another.`

The following indicates that the `group` specified with the `-g` option is already in use:

UX: `roleadd: ERROR: group group does not exist. Choose another.`

The following indicates that the `uid` specified with the `-u` option is in the range of reserved UIDs (from 0-99):

UX: `roleadd: WARNING: uid uid is reserved.`

The following indicates that the `uid` specified with the `-u` option exceeds `MAXUID` as defined in `<sys/param.h>`:

UX: `roleadd: ERROR: uid uid is too big. Choose another.`

The following indicates that the `/etc/passwd` or `/etc/shadow` files do not exist:

UX: `roleadd: ERROR: Cannot update system files - login cannot be created.`

The following indicates that the user executing the command does not have sufficient authorization to perform the operation:

UX: `roleadd: ERROR: Permission denied.`

引用名	roledel – delete a role's login from the system
用法概要	roledel [-r] [-S <i>repository</i>] <i>role</i>
描述	The <code>roledel</code> utility deletes a role account from the system and makes the appropriate account-related changes to the system file and file system. <code>roledel</code> also removes the role from each user's list of assumable roles. An administrator must be granted the User Management Profile to be able to delete an existing role.
选项	The following options are supported: -r Remove the role's home directory from the system. This directory must exist. The files and directories under the home directory will no longer be accessible following successful execution of the command. The ZFS dataset that was created for the role's home directory will be removed. The <code>auto_home</code> entry for the role will be deleted. -S <i>repository</i> The valid repositories are <code>files</code>, <code>ldap</code>. The repository specifies which name service will be updated. The default repository is <code>files</code>.
操作数	The following operands are supported: <i>role</i> An existing role name to be deleted.
退出状态	The following exit values are returned: 0 Successful completion. 2 Invalid command syntax. A usage message for the <code>roledel</code> command is displayed. 6 The account to be removed does not exist. 8 The account to be removed is in use. 10 Cannot update the <code>/etc/group</code> or <code>/etc/user_attr</code> file but the login is removed from the <code>/etc/passwd</code> file. 12 Cannot remove or otherwise modify the home directory.
文件	<code>/etc/passwd</code> system password file <code>/etc/shadow</code> system file containing roles' encrypted passwords and related information <code>/etc/group</code> system file containing group definitions <code>/etc/user_attr</code> system file containing additional role attributes

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [auths\(1\)](#), [passwd\(1\)](#), [profiles\(1\)](#), [roles\(1\)](#), [users\(1B\)](#), [groupadd\(1M\)](#), [groupdel\(1M\)](#), [groupmod\(1M\)](#), [logins\(1M\)](#), [roleadd\(1M\)](#), [rolemod\(1M\)](#), [useradd\(1M\)](#), [userdel\(1M\)](#), [usermod\(1M\)](#), [passwd\(4\)](#), [prof_attr\(4\)](#), [user_attr\(4\)](#), [attributes\(5\)](#)

附注 The `roledel` utility only deletes an account definition that is in the `group`, `passwd`, `shadow`, and `user_attr` databases in the repository.

引用名	rolemod – modify a role's login information on the system
用法概要	<pre> rolemod [-u uid [-o]] [-g group] [-G [+ -]group [, group...]] [-d dir [-m]] [-s shell] [-c comment] [-l new_name] [-f inactive] [-e expire] [-A [+ -]authorization [, authorization]] [-S repository] [-P [+ -]profile [, profile]] [-K key[+ -]=value] role </pre>
描述	The <code>rolemod</code> utility modifies a role's login information on the system. It changes the definition of the specified login and makes the appropriate login-related system file and file system changes. The system file entries created with this command have a limit of 512 characters per line. Specifying long arguments to several options may exceed this limit. An administrator must be granted the User Security Profile to be able to modify the security attributes for an existing role. To be able to modify non-security attributes of an existing user requires the User Management Profile. The authorizations required to set the various fields in <code>passwd</code>, <code>shadow</code> and <code>user_attr</code> can be found in passwd(4), shadow(4), user_attr(4). The authorizations required to assign groups can be found in group(4).
选项	The following options are supported: -A [+]authorization One or more comma separated authorizations as defined in auth_attr(4). Only role with grant rights to the authorization can assign it to an account. This replaces any existing authorization setting. If no authorization list is specified, the existing setting is removed. A prefix <code>+</code> adds the authorization to the existing authorization; a prefix <code>-</code> removes the authorization from the existing authorization. With no prefix, the value for <i>authorization</i> replaces the existing authorization. -c comment Specify a comment string. <i>comment</i> can be any text string. It is generally a short description of the login, and is currently used as the field for the user's full name. This information is stored in the user's <code>/etc/passwd</code> entry. -d dir Specify the new home directory of the role. It defaults to <i>base_dir/login</i>, where <i>base_dir</i> is the base directory for new login home directories, and <i>login</i> is the new login. This creates or modifies an <code>auto_home</code> entry for the user. The argument to the option can be specified as <i>server:dir</i> where <i>server</i> is the hostname of the machine on which the home directory resides and <i>dir</i> is the path to the user's home directory. If the server is a remote host then the home directory needs to be created on the remote host for the system to mount it, when the user logs in. If no server name is specified then the home directory will be created on the host where the command is executed, when the <code>-m</code> option is used.

-e *expire*

Specify the expiration date for a role. After this date, no role will be able to access this login. The expire option argument is a date entered using one of the date formats included in the template file `/etc/datemsk`. See [getdate\(3C\)](#).

For example, you may enter `10/6/90` or `October 6, 1990`. A value of `''` defeats the status of the expired date.

-f *inactive*

Specify the maximum number of days allowed between uses of a login ID before that login ID is declared invalid. Normal values are positive integers. A value of `0` defeats the status.

-g *group*

Specify an existing group's integer ID or character-string name. It redefines the role's primary group membership.

-G *[+|-]group*

An existing group's integer ID or character-string name. It defines the new user's supplementary group membership. Duplicates between group with the `-g` and `-G` options are ignored. No more than `NGROUPS_MAX` groups can be specified. GIDs 0-99 are reserved for allocation by the Solaris Operating System.

A prefix `+` adds the group to the existing group; a prefix `-` removes the group from the existing group. With no prefix, *group* replaces the existing group.

-K *key[+|-]=value*

Replace existing or add to a role's *key=value* pair attributes. Multiple `-K` options can be used to replace or add multiple *key=value* pairs. However, keys must not be repeated. The generic `-K` option with the appropriate key may be used instead of the specific implied key options (`-A` and `-P`). See [user_attr\(4\)](#) for a list of valid *key=value* pairs. If no value is specified, the existing key is removed.

The keyword type can be specified with the value `role` or the value `normal`. When using the value `normal`, the account changes from a role user to a normal user; using the value `role` keeps the account a role user.

A prefix `+` adds the value to the existing value; a prefix `-` removes the value from the existing value. With no prefix, *value* replaces the existing value.

The prefix `+/-` operation is applicable only to the following keys: `auths`, `profiles`, `roles`, `project`, `limitpriv`, and `defaultpriv`.

-l *new_logname*

Specify the new login name for the role. The *new_logname* argument is a string no more than eight bytes consisting of characters from the set of alphabetic characters, numeric characters, period (`.`), underline (`_`), and hyphen (`-`). The first character should be alphabetic and the field should contain at least one lower case alphabetic character. A warning message will be written if these restrictions are not met. A future Solaris release

may refuse to accept login fields that do not meet these requirements. The *new_logname* argument must contain at least one character and must not contain a colon (:) or NEWLINE (\n).

-m

Move the role's home directory to the new directory specified with the -d option. If the directory already exists, it must have permissions read/write/execute by *group*, where *group* is the role's primary group. If the server name specified to the -d option is a remote host then the system will not attempt to create the home directory.

If the directory does not already exist, a new ZFS dataset will be created. In the global zone, the dataset is created as `rpool/export/home/rolename`. For non-global zones, the dataset will be created as `ROOT-dataset/export/home/rolename`. The mountpoint for the ZFS dataset is `/export/home/rolename` by default. If -d *path* is specified and it is a path on the local machine, the dataset will be mounted at the specified location. The role is delegated permissions to create ZFS snapshots and promote them. The newly created dataset will inherit the encryption setting from its parent. If it is encrypted, the role is granted permission to change its wrapping key.

-o

This option allows the specified UID to be duplicated (non-unique).

-P [+|-]*profile*

One or more comma-separated execution profiles defined in [auth_attr\(4\)](#). This replaces any existing profile setting. If no profile list is specified, the existing setting is removed.

A prefix + adds the profile to the existing profile; a prefix - removes the profile from the existing profile. With no prefix, *profile* replaces the existing profile.

-s *shell*

Specify the full pathname of the program that is used as the role's shell on login. The value of *shell* must be a valid executable file.

-S *repository*

The valid repositories are `files`, `ldap`. The repository specifies which name service will be updated. The default repository is `files`. When the repository is `files`, the authorizations, profiles, and roles can be present in other name service repositories and can be assigned to a user in the `files` repository. When the repository is `ldap`, all the assignable attributes must be present in the `ldap` repository.

-u *uid*

Specify a new UID for the role. It must be a non-negative decimal integer less than `MAXUID` as defined in `<param.h>`. The UID associated with the role's home directory is not modified with this option; a role will not have access to their home directory until the UID is manually reassigned using [chown\(1\)](#).

操作数

The following operands are supported:

`login`

An existing login name to be modified.

退出状态

In case of an error, `rolemod` prints an error message and exits with one of the following values:

2

The command syntax was invalid. A usage message for the `rolemod` command is displayed.

3

An invalid argument was provided to an option.

4

The *uid* given with the `-u` option is already in use.

5

The password files contain an error. [pwconv\(1M\)](#) can be used to correct possible errors. See [passwd\(4\)](#).

6

The login to be modified does not exist, the *group* does not exist, or the login shell does not exist.

8

The login to be modified is in use.

9

The *new_logname* is already in use.

10

Cannot update the `/etc/group` or `/etc/user_attr` file. Other update requests will be implemented.

11

Insufficient space to move the home directory (`-m` option). Other update requests will be implemented.

12

Unable to complete the move of the home directory to the new home directory.

文件

`/etc/group`

system file containing group definitions

`/etc/datemsk`

system file of date formats

`/etc/passwd`

system password file

`/etc/shadow`

system file containing users' and roles' encrypted passwords and related information

/etc/user_attr
system file containing additional user and role attributes

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [chown\(1\)](#), [passwd\(1\)](#), [users\(1B\)](#), [groupadd\(1M\)](#), [groupdel\(1M\)](#), [groupmod\(1M\)](#), [logins\(1M\)](#), [pwconv\(1M\)](#), [roleadd\(1M\)](#), [roledel\(1M\)](#), [useradd\(1M\)](#), [userdel\(1M\)](#), [usermod\(1M\)](#), [getdate\(3C\)](#), [auth_attr\(4\)](#), [group\(4\)](#), [passwd\(4\)](#), [shadow\(4\)](#), [user_attr\(4\)](#), [attributes\(5\)](#)

引用名 root_archive – manage bootable miniroot archives

用法概要 /boot/solaris/bin/root_archive pack *archive* *root*
/boot/solaris/bin/root_archive unpack *archive* *root*

描述 The root_archive utility is used to manipulate boot archives. The utility can pack and unpack boot archives and image miniroots. Both ufs and hsf s (iso9660) format archives can be unpacked, although only ufs format is generated when packing.

For normal, boot-related system administration, [bootadm\(1M\)](#) is recommended.

Subcommands The root_archive command has the following subcommands:

pack *archive* *root* Pack the contents of the root directory into the boot archive *archive*.

unpack *archive* *root* Unpack the contents of the boot archive named *archive* to the directory named *root*.

退出状态 The following exit values are returned:

0 The command completed successfully.

1 The command exited due to an error.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [cpio\(1\)](#), [bootadm\(1M\)](#), [mount\(1M\)](#), [attributes\(5\)](#), [lofi\(7D\)](#)

引用名	route – manually manipulate the routing tables										
用法概要	<pre>route [-fnvq] sub-command [[modifiers] args] route [-fnvq] [-p [-R root-dir]] add delete [modifiers] destination gateway [args] route [-fnvq] change get [modifiers] destination [gateway [args]] route [-fn] monitor [modifiers] route [-fnvq] flush [modifiers] route -p [-R root-dir] show</pre>										
描述	route manually manipulates the network routing tables. These tables are normally maintained by the system routing daemon, such as in.routed(1M) and in.ripngd(1M). route supports a limited number of general options, but a rich command language. Users can specify an arbitrary request that can be delivered by means of the programmatic interface discussed in route(7P). route uses a routing socket and the new message types RTM_ADD, RTM_DELETE, RTM_GET, and RTM_CHANGE. While only superusers can modify routing tables, the RTM_GET operation is allowed for non-privileged users.										
选项	<table><tr><td>-f</td><td>Flush the routing tables of all gateway entries. If you use the -f option in conjunction with any of the route sub-commands, route flushes the gateways before performing the sub-command. Specify the table to flush by placing the inet or inet6 modifier immediately after the -f option. If unspecified, flushing IPv4 (inet) routes is the default.</td></tr><tr><td>-n</td><td>Prevent attempts to print host and network names symbolically when reporting actions. This option is useful when name servers are unavailable.</td></tr><tr><td>-p</td><td>Make changes to the network route tables persistent across system restarts. The operation is applied to the network routing tables first and, if successful, is then applied to the list of saved routes used at system startup. In determining whether an operation was successful, a failure to add a route that already exists or to delete a route that is not in the routing table is ignored. Particular care should be taken when using host or network names in persistent routes, as network-based name resolution services are not available at the time routes are added at startup.</td></tr><tr><td>-q</td><td>Suppress all output.</td></tr><tr><td>-R root-dir</td><td>Specify an alternate root directory where route applies changes. This option is ignored unless used in conjunction with the -p option. When -R is specified, route changes are applied only to the list of saved routes to be used at startup, <i>not</i> to the network routing tables. In addition, certain checks, such as the</td></tr></table>	-f	Flush the routing tables of all gateway entries. If you use the -f option in conjunction with any of the route sub-commands, route flushes the gateways before performing the sub-command. Specify the table to flush by placing the inet or inet6 modifier immediately after the -f option. If unspecified, flushing IPv4 (inet) routes is the default.	-n	Prevent attempts to print host and network names symbolically when reporting actions. This option is useful when name servers are unavailable.	-p	Make changes to the network route tables persistent across system restarts. The operation is applied to the network routing tables first and, if successful, is then applied to the list of saved routes used at system startup. In determining whether an operation was successful, a failure to add a route that already exists or to delete a route that is not in the routing table is ignored. Particular care should be taken when using host or network names in persistent routes, as network-based name resolution services are not available at the time routes are added at startup.	-q	Suppress all output.	-R root-dir	Specify an alternate root directory where route applies changes. This option is ignored unless used in conjunction with the -p option. When -R is specified, route changes are applied only to the list of saved routes to be used at startup, <i>not</i> to the network routing tables. In addition, certain checks, such as the
-f	Flush the routing tables of all gateway entries. If you use the -f option in conjunction with any of the route sub-commands, route flushes the gateways before performing the sub-command. Specify the table to flush by placing the inet or inet6 modifier immediately after the -f option. If unspecified, flushing IPv4 (inet) routes is the default.										
-n	Prevent attempts to print host and network names symbolically when reporting actions. This option is useful when name servers are unavailable.										
-p	Make changes to the network route tables persistent across system restarts. The operation is applied to the network routing tables first and, if successful, is then applied to the list of saved routes used at system startup. In determining whether an operation was successful, a failure to add a route that already exists or to delete a route that is not in the routing table is ignored. Particular care should be taken when using host or network names in persistent routes, as network-based name resolution services are not available at the time routes are added at startup.										
-q	Suppress all output.										
-R root-dir	Specify an alternate root directory where route applies changes. This option is ignored unless used in conjunction with the -p option. When -R is specified, route changes are applied only to the list of saved routes to be used at startup, <i>not</i> to the network routing tables. In addition, certain checks, such as the										

existence of network interfaces used with `-ifp`, are skipped. This can be useful from within JumpStart scripts, where the root directory of the system being modified is in a location other than `/`.

-v Print additional details in verbose mode.

Subcommands

The following subcommands are supported:

add	Add a route.
change	Change aspects of a route (such as its gateway).
delete	Delete a specific route.
flush	Remove all gateway entries from the routing table.
get	Look up and display the route for a destination.
monitor	Continuously report any changes to the routing information base, routing lookup misses, or suspected network partitionings.
show	Display the list of routes to be applied at system startup. Can be used only in conjunction with the -p option.

The add and delete sub-commands have the following syntax:

```
route [ -fnvq ] cmd destination gateway [metric/netmask]
```

where *cmd* is add or delete, *destination* is the destination host or network, and *gateway* is the next-hop intermediary through which packets should be routed. Modifiers described in OPERANDS can be placed anywhere on the command line.

The get and change sub-commands have the following syntax:

```
route [ -fnvq ] cmd destination [gateway [metric/netmask]]
```

where *cmd* is get or change, *destination* is the destination host or network, and *gateway* is the next-hop intermediary through which packets should be routed. Modifiers described in OPERANDS can be placed anywhere on the command line.

The `monitor` sub-command has the following syntax:

```
route monitor [ -inet | -inet6 ]
```

操作数

route executes its sub-commands on routes to destinations by way of gateways.

Destinations and Gateways

By default, destination and gateway addresses are interpreted as IPv4 addresses. All symbolic names are tried first as a host name, using `getipnodebyname(3SOCKET)`. If this lookup fails in the `AF_INET` case, `getnetbyname(3SOCKET)` interprets the name as that of a network.

Including an optional modifier on the command line before the address changes how the route sub-command interprets it.

The following modifiers are supported:

- inet Force the address to be interpreted as an IPv4 address, that is, under the AF_INET address family.
- inet6 Force the address to be interpreted as an IPv6 address, that is, under the AF_INET6 address family.

For IPv4 addresses, routes to a particular host are by default distinguished from those to a network by interpreting the Internet address specified as the destination. If the destination has a *local address part* (that is, the portion not covered by the netmask) of 0, or if the destination is resolved as the symbolic name of a network, then the route is assumed to be to a network; otherwise, it is presumed to be a route to a host.

You can force this selection by using one of the following modifiers:

- host Force the destination to be interpreted as a host.
- net Force the destination to be interpreted as a network.

For example:

Destination	Destination Equivalent
128.32	-host 128.0.0.32
128.32.130	-host 128.32.0.130
-net 128.32	128.32.0.0
-net 128.32.130	128.32.130.0

Two modifiers avoid confusion between addresses and keywords (for example., host used as a symbolic host name). You can distinguish a *destination* by preceding it with the -dst modifier. You can distinguish a gateway address by using the -gateway modifier. If the destination is directly reachable by way of an interface requiring no intermediary IP router to act as a gateway, this can be indicated by using the -interface or -iface modifier.

In the following example, the route does not refer to an external gateway (router), but rather to one of the machine's interfaces. Packets with IP destination addresses matching the destination and mask on such a route are sent out on the interface identified by the gateway address. For interfaces using the ARP protocol, this type of route is used to specify that all matching destinations are local to the physical link. That is, a host could be configured to ARP for all addresses, without regard to the configured interface netmask, by adding a default route using this command. For example:

```
example# route add default hostname -interface
```

where gateway address *hostname* is the name or IP address associated with the network interface over which all matching packets should be sent. On a host with a single network interface, *hostname* is usually the same as the *nodename* returned by the `uname -n` command. See [uname\(1\)](#).

For backward compatibility with older systems, directly reachable routes can also be specified by placing a `0` after the gateway address:

```
example# route add default hostname 0
```

This value was once a route metric, but this metric is no longer used. If the value is specified as `0`, then the destination is directly reachable (equivalent to specifying `-interface`). If it is non-zero but cannot be interpreted as a subnet mask, then a gateway is used (default).

With the `AF_INET` address family or an IPv4 address, a separate subnet mask can be specified. This can be specified in one of the following ways:

- IP address following the gateway address. This is typically specified in *decimal dot* notation as for [inet_addr\(3SOCKET\)](#) rather than in symbolic form.
- IP address following the `-netmask` qualifier.
- Slash character and a decimal length appended to the destination address.

If a subnet mask is not specified, the mask used is the subnet mask of the output interface selected by the gateway address, if the classful network of the destination is the same as the classful network of the interface. Otherwise, the classful network mask for the destination address is used.

Each of the following examples creates an IPv4 route to the destination `192.0.2.32` subnet with a subnet mask of `255.255.255.224`:

```
example# route add 192.0.2.32/27 somegateway
example# route add 192.0.2.32 -netmask 255.255.255.224 somegateway
example# route add 192.0.2.32 somegateway 255.255.255.224
```

For IPv6, only the slash format is accepted. The following example creates an IPv6 route to the destination `3ffe::` with a netmask of 16 one-bits followed by 112 zero-bits.

```
example# route add -inet6 3ffe::/16 somegateway
```

In cases where the gateway does not uniquely identify the output interface (for example, when several interfaces have the same address), you can use the `-ifp ifname` modifier to specify the interface by name. For example, `-ifp lo0` associates the route with the `lo0` interface. If the named interface is an underlying interface in an IPMP (IP multipathing) group, then requests to add a route will automatically be translated to the corresponding IPMP IP interface, and requests to delete or change a route on an underlying interface will fail.

When the routing table contains several equal routes, that is, routes for the same destination and mask, then IP attempts to spread the traffic over those routes. The spreading is such that an individual transport connection uses the same route to avoid packet reordering as seen by, for example, TCP. The details of the spreading algorithm is not documented and is likely to evolve over time.

Routing Flags

Routes have associated flags that influence operation of the protocols when sending to destinations matched by the routes. These flags can be set (and in some cases cleared, indicated by ~) by including the following modifiers on the command line:

Modifier	Flag	Description
-interface	~RTF_GATEWAY	Destination is directly reachable
-iface	~RTF_GATEWAY	Alias for interface modifier
-static	RTF_STATIC	Manually added route
-nostatic	~RTF_STATIC	Pretend route was added by kernel or routing daemon
-reject	RTF_REJECT	Emit an ICMP unreachable when matched
-blackhole	RTF_BLACKHOLE	Silently discard packets
-proto1	RTF_PROTO1	Set protocol specific routing flag #1
-proto2	RTF_PROTO2	Set protocol specific routing flag #2
-private	RTF_PRIVATE	Do not advertise this route
-multirt	RTF_MULTIRT	Creates the specified redundant route
-setsrc	RTF_SETSRC	Assigns the default source address
-indirect	RTF_INDIRECT	Allows adding routes where gateway is not on-link

The optional `-indirect` modifier allows adding routes where the gateway is not directly reachable. When an indirect route is the best match for a packet to be sent or forwarded, then IP proceeds to look up that gateway to find a route that is directly reachable. The `-indirect` modifier can be used even if the gateway is directly reachable.

The optional modifiers `-rtt`, `-rttvar`, `-sendpipe`, `-recvpipe`, `-mtu`, `-hopcount`, `-expire`, and `-sssthresh` provide initial values to quantities maintained in the routing entry by transport level protocols, such as TCP. These can be individually locked either by preceding each modifier to be locked by the `-lock` meta-modifier, or by specifying that all ensuing metrics can be locked by the `-lockrest` meta-modifier.

Some transport layer protocols can support only some of these metrics. The following optional modifiers are supported:

`-expire` Lifetime for the entry. This optional modifier is not currently supported.

- hopcount Maximum hop count. This optional modifier is not currently supported.
- mtu Maximum MTU in bytes.
- recvpipe Receive pipe size in bytes.
- rtt Round trip time in microseconds.
- rttvar Round trip time variance in microseconds.
- sendpipe Send pipe size in bytes.
- ssthresh Send pipe size threshold in bytes.
- secattr Security attributes of the route. This modifier is available only if the system is configured with the Solaris Trusted Extensions feature.

The -secattr modifier has the following format:

`min_sl=val,max_sl=val,doi=val,cipso`

or:

`sl=VAL,doi=VAL,cipso`

In the first form, above, the *val* for `min_sl` and `max_sl` is a sensitivity label in either hex or string form. The *val* for `doi` is a non-negative integer. The route will apply only for packets with the same domain of interpretation as defined by the `doi` value and within the accreditation range defined by the `min_sl` and `max_sl` values. The `cipso` keyword is optional and set by default. Valid `min_sl`, `max_sl` and `doi` keyword/value pairs are mandatory. Note that if *val* contains a space, it must be protected by double quotes.

The second form, above, is equivalent to specifying the first form with the same `VAL` for `min_sl` and `max_sl`. The second form should be used for the `get` command, because `get` uses only a single sensitivity label.

Compatibility

The modifiers `host` and `net` are taken to be equivalent to `-host` and `-net`. To specify a symbolic address that matches one of these names, use the `dst` or `gateway` keyword to distinguish it. For example: `-dst host`

The following two flags are also accepted for compatibility with older systems, but have no effect.

Modifier	Flag
-cloning	RTF_CLONING
-xresolve	RTF_XRESOLVE

The -ifa hostname modifier is also accepted, but has no effect.

- 文件
- /etc/defaultrouter

List of default routers
- /etc/hosts

List of host names and net addresses
- /etc/networks

List of network names and addresses

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[uname\(1\)](#), [in.ripngd\(1M\)](#), [in.routed\(1M\)](#), [netstat\(1M\)](#), [routed\(1M\)](#), [ioctl\(2\)](#), [getipnodebyname\(3SOCKET\)](#), [getnetbyname\(3SOCKET\)](#), [inet_addr\(3SOCKET\)](#), [defaultrouter\(4\)](#), [hosts\(4\)](#), [networks\(4\)](#), [attributes\(5\)](#), [arp\(7P\)](#), [ip\(7P\)](#), [route\(7P\)](#), [routing\(7P\)](#)

- 诊断
- add [host | network] *destination:gateway flags*

The specified route is being added to the tables. The values printed are from the routing table entry supplied in the [ioctl\(2\)](#) call. If the gateway address used was not the primary address of the gateway (the first one returned by [getipnodebyname\(3SOCKET\)](#)) the gateway address is printed numerically as well as symbolically.

delete [host | network] *destination:gateway flags*

As add, but when deleting or changing an entry.

change [host | network] *destination:gateway flags*

As add, but when deleting or changing an entry.

destination done

When the -f flag is specified, or the flush sub-command is used, each routing table entry deleted is indicated with a message of this form.

Network is unreachable

An attempt to add a route failed because the gateway listed was not on a directly-connected network. Give the next-hop gateway instead.

not in table

A delete operation was attempted for an entry that is not in the table.

entry exists

An add operation was attempted for a route that already exists in the kernel.

系统管理命令－第2单元

1877

routing table overflow

An operation was attempted, but the system was unable to allocate memory to create the new entry.

insufficient privileges

An attempt to add, delete, change, or flush a route failed because the calling process does not have appropriate privileges.

附注

Specifying that destinations are local (with the `-interfacemodifier`) assumes that the routers implement proxy ARP, meaning that they respond to ARP queries for all reachable destinations. Normally, using either router discovery or RIP is more reliable and scalable than using proxy ARP. See [in.routed\(1M\)](#) for information related to RIP.

Combining the all destinations are local route with subnet or network routes can lead to unpredictable results. The search order as it relates to the all destinations are local route are undefined and can vary from release to release.

引用名	routeadm – IP forwarding and routing configuration
用法概要	<pre>routeadm [-p <i>option</i>]] routeadm [-R <i>root-dir</i>] [-e <i>option</i> ...] [-d <i>option</i>...] [-r <i>option</i>...] [-s <i>var=value</i>] routeadm [-l <i>fmri</i>] routeadm [-m <i>fmri key=value</i> [<i>key=value</i>]...] routeadm [-u]</pre>
描述	The routeadm command is used to administer system-wide configuration for IP forwarding and routing. IP forwarding is the passing of IP packets from one network to another; IP routing is the use of a routing protocol to determine routes. IP routing functions are also represented as services within the service management facility (SMF), and can be administered by means of svcadm(1M) also, using the following fault management resource identifiers (FMRIs): <pre>svc:/network/routing/route:default svc:/network/routing/ripng:default</pre> See EXAMPLES for relevant examples. In addition to enabling and disabling routing and forwarding, routeadm is used to interact with SMF-based routing daemon services. Routing daemon services are identified by the presence of a routeadm application property group, which routeadm uses in administering the given service. Routing daemon services can also specify properties relating to their operation in the routing application property group; these can be modified by means of routeadm -m. If an FMRI for a service without such a property group is specified, an error is issued and the operation is not carried out. If a routing daemon has not been converted to SMF, the <code>ipv4[or 6]-routing-daemon</code>, <code>ipv4[or 6]-routing-daemon-args</code>, and <code>ipv4[or 6]-routing-stop-cmd</code> variables can be used to specify the appropriate daemon for IPv4 or IPv6 routing. routeadm will then run that daemon using the <code>svc:/network/routing/legacy-routing:ipv4[or 6]</code> service as appropriate. This conversion process occurs when you issue an enable (-e), disable (-d) or an update (-u) command. The first usage, in the SYNOPSIS above, reports the current configuration.
选项	The following command-line options are supported: <pre>-p <i>option</i></pre> Print the configuration in parseable format. If <i>option</i> is specified, only the configuration for the specified option or variable is displayed.

-R *root-dir*

Specify an alternate root directory where routeadm applies changes. This can be useful from within JumpStart scripts, where the root directory of the system being modified is mounted elsewhere.

注 – The root file system of any non-global zones must not be referenced with the -R option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See [zones\(5\)](#).

-e *option...*

Enable the specified option. The effect is to prepare the associated services for enabling. By means of the `routing-svcs` variable, the routing daemons are specified to be enabled on subsequent boot or when routeadm -u is run.

-d *option...*

Disable the specified option. The effect is to prepare the associated services for disabling. By means of the `routing-svcs` variable, the routing daemons are specified to be disabled on subsequent boot or when routeadm -u is run.

-l *fmri*

List all properties in the routing application property group for the SMF routing daemon service.

-m *fmri key=value*

Change property value of property *key* to *value* in routing application property group for the SMF routing daemon service. For multi-valued properties, the property name can be used multiple times in the modify operation, and each associated value will be added.

-r *option...*

Revert the specified option to the system default. The system defaults are specified in the description of each *option*.

-u

Apply the currently configured options to the running system. These options might include enabling or disabling IP forwarding and launching or killing routing daemons, if any are specified. It does not alter the state of the system for those settings that have been set to default. This option is meant to be used by administrators who do not want to reboot to apply their changes. In addition, this option upgrades non-SMF configurations from the invocations of daemon stop commands, which might include a set of arguments, to a simple enabling of the appropriate service.

-s *key=value*

Specify string values for specific variables in a comma-separated list with no intervening spaces. If invalid options are specified, a warning message is displayed and the program exits. The following variables can be specified:

`routing-svcs=fmrilist`

Specifies the routing daemon services to be enabled. Routing daemon services are determined to be IPv4 or IPv6 (and so enabled or disabled when routeadm -e/-d

ipv4(6)-routing is run) on the basis of property values in the routeadm application property group. Default: route:default ripng:default

ipv4-routing-daemon=<full_path_to_routing_daemon>

Specifies the routing daemon to be started when ipv4-routing is enabled. The routing daemon specified must be an executable binary or shell-script. If the specified program maps to an SMF service, the service will be used, and daemon arguments to the program will be transferred to the properties of the service at enable time. Default: ""

ipv4-routing-daemon-args=<args>

Specifies the startup arguments to be passed to the ipv4-routing-daemon when ipv4-routing is enabled. Default: no arguments

ipv4-routing-stop-cmd=<command>

Specifies the command to be executed to stop the routing daemon when ipv4-routing is disabled. <command> can be an executable binary or shell-script, or a string that can be parsed by [system\(3C\)](#). Default: ""

ipv6-routing-daemon=<full_path_to_routing_daemon>

Specifies the routing daemon to be started when ipv6-routing is enabled. The routing daemon specified must be an executable binary or shell-script. If the specified program maps to an SMF service, the service will be used, and daemon arguments to the program will be transferred to the properties of the service at enable time. Default: ""

ipv6-routing-daemon-args=<args>

Specifies the startup arguments to be passed to the ipv6-routing-daemon when ipv6-routing is enabled. Default: ""

ipv6-routing-stop-cmd=<command>

Specifies the command to be executed to stop the routing daemon when ipv6-routing is disabled. <command> can be an executable binary or shell-script, or a string that can be parsed by [system\(3C\)](#). Default: ""

Multiple -e, -d, and -r options can be specified on the command line. Changes made by -e, -d, and -r are persistent, but are not applied to the running system unless routeadm is called later with the -u option.

Use the following options as arguments to the -e, -d, and -r options (shown above as *option...*).

ipv4-forwarding

Controls the global forwarding configuration for all IPv4 interfaces. The system default is disabled. If enabled, IP will forward IPv4 packets to and from interfaces when appropriate. If disabled, IP will not forward IPv4 packets to and from interfaces when appropriate.

ipv4-routing

Determines whether an IPv4 routing daemon is run. The system default is enabled unless the /etc/defaultrouter file exists (see [defaultrouter\(4\)](#)), in which case the default is

disabled. The value of this option reflects the state of all IPv4 routing services, such that if any IPv4 routing service is enabled, `ipv4-routing` is enabled. This allows users to interact with routing services using `svcadm(1M)`, as well as through `routeadm`. IPv4 routing services, specified by means of the `routing-svcs` variable, will be prepared for enable on next boot when the user explicitly enables `ipv4-routing`. The SMF routing daemon service for `in.routed` (`svc:/network/routing/route:default`) is specified by default.

`ipv6-forwarding`

Controls the global forwarding configuration for all IPv6 interfaces. The system default is disabled. If enabled, IP will forward IPv6 packets to and from interfaces when appropriate. If disabled, IP will not forward IPv6 packets to and from interfaces when appropriate.

`ipv6-routing`

Determines whether an IPv6 routing daemon is run. The system default is disabled. The value of this option reflects the state of all IPv6 routing services, such that, if any IPv6 routing service is enabled, `ipv6-routing` is enabled. This allows users to interact with routing services via `svcadm(1M)` as well as through `routeadm`. IPv6 routing services, specified by means of the `routing-svcs` variable, will be prepared for enable on next boot when the user explicitly enables `ipv6-routing`. The SMF routing daemon service for `in.ripngd` (`svc:/network/routing/ripng:default`) is specified by default.

The forwarding and routing settings are related but not mutually dependent. For example, a router typically forwards IP packets and uses a routing protocol, but nothing would prevent an administrator from configuring a router that forwards packets and does not use a routing protocol. In that case, the administrator would enable forwarding, disable routing, and populate the router's routing table with static routes.

The forwarding settings are global settings. Each interface also has an `IFF_ROUTER` forwarding flag that determines whether packets can be forwarded to or from a particular interface. That flag can be independently controlled by means of `ifconfig(1M)`'s `router` option. When the global forwarding setting is changed (that is, `-u` is issued to change the value from enabled to disabled or vice-versa), all interface flags in the system are changed simultaneously to reflect the new global policy. Interfaces configured by means of DHCP automatically have their interface-specific `IFF_ROUTER` flag cleared.

When a new interface is plumbed by means of `ifconfig`, the value of the interface-specific forwarding flag is set according to the current global forwarding value. Thus, the forwarding value forms the “default” for all new interfaces.

示例

示例 1 Enabling IPv4 Forwarding

IPv4 forwarding is disabled by default. The following command enables IPv4 forwarding:

```
example# routeadm -e ipv4-forwarding
```

示例 2 Apply Configured Settings to the Running System

In the previous example, a system setting was changed, but will not take effect until the next reboot unless a command such as the following is used:

```
example# routeadm -u
```

示例 3 Making a Setting Revert to its Default

To make the setting changed in the first example revert to its default, enter the following:

```
example# routeadm -r ipv4-forwarding
example# routeadm -u
```

示例 4 Starting in.routed with the -q Flag

Setting the -q flag is represented in the SMF service by setting the quiet_mode property to true. The following sequence of commands starts in.routed with the -q flag:

```
example# routeadm -m route:default quiet_mode=true
example# routeadm -e ipv4-routing -u
```

See [in.routed\(1M\)](#) for details of property names and how they relate to daemon behavior.

退出状态

The following exit values are returned:

- 0 Successful completion.
- !=0 An error occurred while obtaining or modifying the system configuration.

文件

/etc/inet/routing.conf Parameters for IP forwarding and routing. (Not to be edited.)

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[ifconfig\(1M\)](#), [in.routed\(1M\)](#), [ipadm\(1M\)](#), [svcadm\(1M\)](#), [gateways\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

引用名

rpcbind – universal addresses to RPC program number mapper

用法概要

```
rpcbind [-d] [-w]
```

描述

rpcbind is a server that converts RPC program numbers into universal addresses. It must be running on the host to be able to make RPC calls on a server on that machine.

When an RPC service is started, it tells rpcbind the address at which it is listening, and the RPC program numbers it is prepared to serve. When a client wishes to make an RPC call to a given program number, it first contacts rpcbind on the server machine to determine the address where RPC requests should be sent.

rpcbind should be started before any other RPC service. Normally, standard RPC servers are started by port monitors, so rpcbind must be started before port monitors are invoked.

When rpcbind is started, it checks that certain name-to-address translation-calls function correctly. If they fail, the network configuration databases can be corrupt. Since RPC services cannot function correctly in this situation, rpcbind reports the condition and terminates.

rpcbind maintains an open transport end for each transport that it uses for indirect calls. This is the UDP port on most systems.

The rpcbind service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/rpc/bind
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). rpcbind can only be started by the superuser or someone in the Primary Administrator role.

The configuration properties of this service can be modified with [svccfg\(1M\)](#).

The following SMF property is used to allow or disallow access to rpcbind by remote clients:

```
config/local_only = true
```

The default value, true, shown above, disallows remote access; a value of false allows remote access. See EXAMPLES.

The FMRI svc:/network/rpc/bind property group config contains the following property settings:

enable_tcpwrappers	Specifies that the TCP wrappers facility is used to control access to TCP services. The value true enables checking. The default value for enable_tcpwrappers is false. If the enable_tcpwrappers parameter is enabled, then all calls to rpcbind originating from non-local addresses are automatically wrapped by the TCP wrappers facility. The syslog facility code daemon is used to log

	allowed connections (using the info severity level) and denied traffic (using the warning severity level). See syslog.conf(4) for a description of syslog codes and severity levels. The Interface Stability of the TCP wrappers facility and its configuration files is Volatile. As the TCP wrappers facility is not controlled by Sun, intrarelease incompatibilities are not uncommon. See attributes(5) .
verbose_logging	Specifies whether the TCP wrappers facility logs all calls or just the denied calls. The default is false. This option has no effect if TCP wrappers are not enabled.
allow_indirect	Specifies whether rpcbind allows indirect calls at all. By default, rpcbind allows most indirect calls, except to a number of standard services (keyserv, automount, mount, nfs, rquota, and selected NIS and rpcbind procedures). Setting allow_indirect to false causes all indirect calls to be dropped. The default is true. NIS broadcast clients rely on this functionality on NIS servers.

选项

The following options are supported:

- d Run in debug mode. In this mode, rpcbind does not fork when it starts. It prints additional information during operation, and aborts on certain errors. With this option, the name-to-address translation consistency checks are shown in detail.
- w Do a warm start. If rpcbind aborts or terminates on SIGINT or SIGTERM, it writes the current list of registered services to /var/run/portmap.file and /var/run/rpcbind.file. Starting rpcbind with the -w option instructs it to look for these files and start operation with the registrations found in them. This allows rpcbind to resume operation without requiring all RPC services to be restarted.

示例

示例 1 Allowing Remote Access

The following sequence of commands allows remote access to rpcbind.

```
# svccfg -s svc:/network/rpc/bind setprop config/local_only = false
# svcadm refresh svc:/network/rpc/bind
```

文件

/var/run/portmap.file	Stores the information for RPC services registered over IP based transports for warm start purposes.
/var/run/rpcbind.file	Stores the information for all registered RPC services for warm start purposes.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	See below.

TCP wrappers is Volatile.

另请参见

[smf\(5\)](#), [rpcinfo\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [rpcbind\(3NSL\)](#), [syslog.conf\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

For information on the TCP wrappers facility, see the `hosts_access(4)` man page, delivered as part of the Solaris operating environment in `/usr/sfw/man` and available in the `security/tcp-wrapper` package.

附注

Terminating `rpcbind` with `SIGKILL` prevents the warm-start files from being written.

All RPC servers are restarted if the following occurs: `rpcbind` crashes (or is killed with `SIGKILL`) and is unable to write the warm-start files; `rpcbind` is started without the `-w` option after a graceful termination. Otherwise, the warm start files are not found by `rpcbind`.

- 引用名

rpc.bootparamd, bootparamd – boot parameter server
- 用法概要

```
/usr/sbin/rpc.bootparamd [-d]
```
- 描述

rpc.bootparamd is a server process that provides information from a bootparams database to diskless clients at boot time. See [bootparams\(4\)](#)

The source for the bootparams database is determined by the [nsswitch.conf\(4\)](#) file (on the machine running the rpc.bootparamd process).
- 选项

The following options are supported:

 - d Display debugging information.
- 文件

/etc/bootparams	boot parameter data base
/etc/nsswitch.conf	configuration file for the name-service switch
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/boot/network

另请参见

[svcs\(1\)](#), [svcadm\(1M\)](#), [bootparams\(4\)](#), [nsswitch.conf\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注

A diskless client requires service from at least one rpc.bootparamd process running on a server that is on the same IP subnetwork as the diskless client.

Some routines that compare hostnames use case-sensitive string comparisons; some do not. If an incoming request fails, verify that the case of the hostname in the file to be parsed matches the case of the hostname called for, and attempt the request again.

The rpc.bootparamd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/rpc/bootparams
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名

rpcinfo – report RPC information

用法概要

```
rpcinfo [-m | -s] [host]
rpcinfo -p [host]
rpcinfo -T transport host prognum [versnum]
rpcinfo -l [-T transport] host prognum versnum
rpcinfo [-n portnum] -u host prognum [versnum]
rpcinfo [-n portnum] -t host prognum [versnum]
rpcinfo -a serv_address -T transport prognum [versnum]
rpcinfo -b [-T transport] prognum versnum
rpcinfo -d [-T transport] prognum versnum
```

描述

rpcinfo makes an RPC call to an RPC server and reports what it finds.

In the first synopsis, `rpcinfo` lists all the registered RPC services with `rpcbind` on *host*. If *host* is not specified, the local host is the default. If `-s` is used, the information is displayed in a concise format.

In the second synopsis, `rpcinfo` lists all the RPC services registered with `rpcbind`, version 2. Note that the format of the information is different in the first and the second synopsis. This is because the second synopsis is an older protocol used to collect the information displayed (version 2 of the `rpcbind` protocol).

The third synopsis makes an RPC call to procedure 0 of *prognum* and *versnum* on the specified *host* and reports whether a response was received. *transport* is the transport which has to be used for contacting the given service. The remote address of the service is obtained by making a call to the remote `rpcbind`.

The *prognum* argument is a number that represents an RPC program number (see [rpc\(4\)](#)).

If a *versnum* is specified, `rpcinfo` attempts to call that version of the specified *prognum*. Otherwise, `rpcinfo` attempts to find all the registered version numbers for the specified *prognum* by calling version 0, which is presumed not to exist; if it does exist, `rpcinfo` attempts to obtain this information by calling an extremely high version number instead, and attempts to call each registered version. Note that the version number is required for `-b` and `-d` options.

The EXAMPLES section describe other ways of using `rpcinfo`.

选项

<code>-T transport</code>	Specify the transport on which the service is required. If this option is not specified, <code>rpcinfo</code> uses the transport specified in the <code>NETPATH</code> environment variable, or if that is unset or <code>NULL</code> , the transport in the netconfig(4) database is used. This is a generic option, and can be used in conjunction with other options as shown in the SYNOPSIS.

-
- a *serv_address* Use *serv_address* as the (universal) address for the service on *transport* to ping procedure 0 of the specified *prognum* and report whether a response was received. The -T option is required with the -a option. If *versnum* is not specified, *rpcinfo* tries to ping all available version numbers for that program number. This option avoids calls to remote *rpcbind* to find the address of the service. The *serv_address* is specified in universal address format of the given transport.
 - b Make an RPC broadcast to procedure 0 of the specified *prognum* and *versnum* and report all hosts that respond. If *transport* is specified, it broadcasts its request only on the specified transport. If broadcasting is not supported by any transport, an error message is printed. Use of broadcasting should be limited because of the potential for adverse effect on other systems.
 - d Delete registration for the RPC service of the specified *prognum* and *versnum*. If *transport* is specified, unregister the service on only that transport, otherwise unregister the service on all the transports on which it was registered. Only the owner of a service can delete a registration, except the superuser, who can delete any service.
 - l Display a list of entries with a given *prognum* and *versnum* on the specified *host*. Entries are returned for all transports in the same protocol family as that used to contact the remote *rpcbind*.
 - m Display a table of statistics of *rpcbind* operations on the given *host*. The table shows statistics for each version of *rpcbind* (versions 2, 3 and 4), giving the number of times each procedure was requested and successfully serviced, the number and type of remote call requests that were made, and information about RPC address lookups that were handled. This is useful for monitoring RPC activities on *host*.
 - n *portnum* Use *portnum* as the port number for the -t and -u options instead of the port number given by *rpcbind*. Use of this option avoids a call to the remote *rpcbind* to find out the address of the service. This option is made obsolete by the -a option.
 - p Probe *rpcbind* on *host* using version 2 of the *rpcbind* protocol, and display a list of all registered RPC programs. If *host* is not specified, it defaults to the local host. This option is not useful for IPv6; use -s (see below) instead. Note that version 2 of the *rpcbind* protocol was previously known as the portmapper protocol.
 - s Display a concise list of all registered RPC programs on *host*. If *host* is not specified, it defaults to the local host.

- t Make an RPC call to procedure 0 of *prognum* on the specified *host* using TCP, and report whether a response was received. This option is made obsolete by the -T option as shown in the third synopsis.
- u Make an RPC call to procedure 0 of *prognum* on the specified *host* using UDP, and report whether a response was received. This option is made obsolete by the -T option as shown in the third synopsis.

示例

示例1 RPC services.

To show all of the RPC services registered on the local machine use:

```
example% rpcinfo
```

To show all of the RPC services registered with rpcbind on the machine named klaxon use:

```
example% rpcinfo klaxon
```

The information displayed by the above commands can be quite lengthy. Use the -s option to display a more concise list:

```
example% rpcinfo -s klaxon
```

program	vrsn	netid(s)	service	owner
100000	2,3,4	tcp,udp,ticlts,ticots,ticotsord	rpcbind	superuser
100008	1	ticotsord,ticots,ticlts,udp,tcp	walld	superuser
100002	2,1	ticotsord,ticots,ticlts,udp,tcp	rusersd	superuser
100001	2,3,4	ticotsord,ticots,tcp,ticlts,udp	rstatd	superuser
100012	1	ticotsord,ticots,ticlts,udp,tcp	sprayd	superuser
100007	3	ticotsord,ticots,ticlts,udp,tcp	ypbind	superuser
100029	1	ticotsord,ticots,ticlts	keyserv	superuser
100078	4	ticotsord,ticots,ticlts	-	superuser
100024	1	ticotsord,ticots,ticlts,udp,tcp	status	superuser
100021	2,1	ticotsord,ticots,ticlts,udp,tcp	nlockmgr	superuser
100020	1	ticotsord,ticots,ticlts,udp,tcp	llockmgr	superuser

To show whether the RPC service with program number *prognum* and version *versnum* is registered on the machine named klaxon for the transport TCP use:

```
example% rpcinfo -T tcp klaxon prognum versnum
```

To show all RPC services registered with version 2 of the rpcbind protocol on the local machine use:

```
example% rpcinfo -p
```

To delete the registration for version 1 of the walld (program number 100008) service for all transports use:

```
example# rpcinfo -d 100008 1
```

or

```
example# rpcinfo -d walld 1
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[rpcbind\(1M\)](#), [rpc\(3NSL\)](#), [netconfig\(4\)](#), [rpc\(4\)](#), [attributes\(5\)](#)

引用名	rpc.mdcommd – multi-node disk set services
用法概要	/usr/sbin/rpc.mdcommd
描述	rpc.mdcommd is an rpc(4) daemon that functions as a server process. rpc.mdcommd manages communication among hosts participating in a multi-node disk set configuration. rpc.mdcommd is invoked by inetd(1M).
退出状态	The following exit values are returned: 0 Successful completion. >0 An error occurred.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见	svcs(1), inetadm(1M), inetd(1M), metaset(1M), svcadm(1M), rpc(3NSL), rpc(4), services(4), attributes(5), smf(5) 《Solaris Volume Manager Administration Guide》

附注	The rpc.mdcommd service is managed by the service management facility, smf(5), under the service identifier: svc:/network/rpc/mdcomm Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M). Responsibility for initiating and restarting this service is delegated to inetd(1M). Use inetadm(1M) to make configuration changes and to view configuration information for this service. The service's status can be queried using the svcs(1) command.

引用名	rpc.metad – remote metaset services
用法概要	/usr/sbin/rpc.metad
描述	rpc.metad is an rpc(4) daemon (functioning as a server process) that is used to manage local copies of metadevice diskset information. The rpc.metad daemon is controlled by inetadm(1M) .
退出状态	The following exit values are returned: 0 Successful completion. >0 An error occurred.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm

另请参见 [svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [metaset\(1M\)](#), [rpc.metamhd\(1M\)](#), [svcadm\(1M\)](#), [rpc\(3NSL\)](#), [services\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

《Solaris Volume Manager Administration Guide》

附注 The rpc.metad service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/network/rpc/meta:default

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名 rpc.metamedd – remote mediator services

用法概要 /usr/sbin/rpc.metamedd

描述 rpc.metamedd is an [rpc\(4\)](#) server which is used to manage mediator information for use in 2-string HA configurations. The rpc.metamedd daemon is controlled by [inetadm\(1M\)](#).

退出状态 The following exit values are returned:

0 Successful completion.

>0 An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	storage/svm
Interface Stability	Committed

另请参见 [svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [svcadm\(1M\)](#), [rpc\(4\)](#), [services\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

Sun Cluster documentation, 《[Solaris Volume Manager Administration Guide](#)》

附注 The rpc.metamedd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/network/rpc/metamed:default

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名	rpc.metamhd – remote multihost disk services				
用法概要	/usr/sbin/rpc.metamhd				
描述	rpc.metamhd is an rpc(4) daemon (functioning as a server process) that is used to manage multi-hosted disks. The <code>rpc.metamhd</code> daemon is controlled by inetadm(1M) .				
退出状态	The following exit values are returned: 0 Successful completion. >0 An error occurred.				
属性	See attributes(5) for descriptions of the following attributes:				
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>storage/svm</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	storage/svm
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	storage/svm				
另请参见	svcs(1) , inetadm(1M) , inetd(1M) , metaset(1M) , rpc.metad(1M) , svcadm(1M) , rpc(3NSL) , services(4) , attributes(5) , smf(5) 《Solaris Volume Manager Administration Guide》				
附注	The <code>rpc.metamhd</code> service is managed by the service management facility, smf(5), under the service identifier: <code>svc:/network/rpc/metamh:default</code> Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M). Responsibility for initiating and restarting this service is delegated to inetd(1M). Use inetadm(1M) to make configuration changes and to view configuration information for this service. The service's status can be queried using the svcs(1) command.				

引用名

rpc.rexd, rexd – RPC-based remote execution server

用法概要

/usr/sbin/rpc.rexd [-s]

描述

rpc . rexd is the Sun RPC server for remote program execution. This daemon is started by [inetd\(1M\)](#) whenever a remote execution request is made.

For non-interactive programs, the standard file descriptors are connected directly to TCP connections. Interactive programs involve pseudo-terminals, in a fashion that is similar to the login sessions provided by [rlogin\(1\)](#). This daemon may use NFS to mount file systems specified in the remote execution request.

There is a 10240 byte limit for arguments to be encoded and passed from the sending to the receiving system.

选项

The following option is supported:

-s Secure. When specified, requests must have valid DES credentials. If the request does not have a DES credential it is rejected. The default publickey credential is rejected. Only newer [on\(1\)](#) commands send DES credentials.

If access is denied with an authentication error, you may have to set your publickey with the [chkey\(1\)](#) command.

Specifying the -s option without presenting secure credentials will result in an error message: Unix too weak auth (DesONLY)!

安全

rpc . rexd uses [pam\(3PAM\)](#) for account and session management. The PAM configuration policy, listed through /etc/pam . conf, specifies the modules to be used for rpc . rexd. Here is a partial pam . conf file with rpc . rexd entries for account and session management using the UNIX module.

rpc.rexd	account requisite	pam_roles.so.1
rpc.rexd	account required	pam_projects.so.1
rpc.rexd	account required	pam_unix_account.so.1
rpc.rexd	session required	pam_unix_session.so.1

If there are no entries for the rpc . rexd service, the entries for the “other” service will be used. rpc . rexd uses the getpuid () call to determine whether the given user is a legal user.

文件

/dev/pts/n	Pseudo-terminals used for interactive mode
/etc/passwd	Authorized users
/tmp_rex/rexd??????	Temporary mount points for remote file systems

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/nis

另请参见 [chkey\(1\)](#), [on\(1\)](#), [rlogin\(1\)](#), [svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [svcadm\(1M\)](#), [pam\(3PAM\)](#), [pam.conf\(4\)](#), [publickey\(4\)](#), [attributes\(5\)](#), [pam_authtok_check\(5\)](#), [pam_authtok_get\(5\)](#), [pam_authtok_store\(5\)](#), [pam_dhkeys\(5\)](#), [pam_passwd_auth\(5\)](#), [pam_unix_account\(5\)](#), [pam_unix_auth\(5\)](#), [pam_unix_session\(5\)](#), [smf\(5\)](#)

诊断 Diagnostic messages are normally printed on the console, and returned to the requestor.

附注 Root cannot execute commands using rexd client programs such as [on\(1\)](#).

The `rpc.rexd` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/network/rpc/rex:default`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名 `rpc.rstatd`, `rstatd` – kernel statistics server

用法概要 `/usr/lib/netsvd/rstat/rpc.rstatd`

描述 `rpc.rstatd` is a server which returns performance statistics obtained from the kernel. [rup\(1\)](#) uses `rpc.rstatd` to collect the uptime information that it displays.

`rpc.rstatd` is an RPC service.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-servers

另请参见 [rup\(1\)](#), [svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [svcadm\(1M\)](#), [services\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The `rpc.rstatd` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/network/rpc/rstat:default`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

- 引用名

rpc.rusersd, rusersd – network username server
- 用法概要

/usr/lib/netsvc/rusers/rpc.rusersd
- 描述

rpc . rusersd is a server that returns a list of users on the host. The rpc . rusersd daemon may be started by [inetd\(1M\)](#).
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-servers

另请参见

[svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [svcadm\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注

The rpc . rusersd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/rpc/rusers:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名 `rpc.rwalld, rwalld` – network rwall server

用法概要 `/usr/lib/netsvc/rwall/rpc.rwalld`

描述 `rpc.rwalld` is a server that handles [rwall\(1M\)](#) requests. It is implemented by calling [wall\(1M\)](#) on all the appropriate network machines. The `rpc.rwalld` daemon may be started by [inetd\(1M\)](#).

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-servers

另请参见 [svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [rwall\(1M\)](#), [svcadm\(1M\)](#), [wall\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The `rpc.rwalld` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/network/rpc/wall:default`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名 `rpc.smserverd` – removable media device server

用法概要 `/usr/lib/smedia/rpc.smserverd`

描述 `rpc.smserverd` is a server that handles requests from client applications, such as volume management software, for access to removable media devices. In addition to volume management software, [rmformat\(1\)](#) and the CDE Filemanager (when performing removable media operations) are `rpc.smserverd` clients. The `rpc.smserverd` daemon is started by [inetd\(1M\)](#) when a client makes a call to a Solaris-internal library to access a SCSI, IDE, or USB device. The daemon is not started if a client attempts to access a PCMCIA device. Once started, the daemon remains active until such time as it is idle for three minutes or more.

The `rpc.smserverd` daemon is provided for the exclusive use of the client applications mentioned above. It has no external, customer-accessible interfaces, including no configuration file.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/storage/removable-media

另请参见 [svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [svcadm\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The `rpc.smserverd` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/network/rpc/smserver`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名 `rpc.sprayd`, `sprayd` – spray server

用法概要 `/usr/lib/netsvc/spray/rpc.sprayd`

描述 `rpc.sprayd` is a server that records the packets sent by [spray\(1M\)](#). The `rpc.sprayd` daemon may be started by [inetd\(1M\)](#).

The service provided by `rpc.sprayd` is not useful as a networking benchmark as it uses unreliable connectionless transports, (udp for example). It can report a large number of packets dropped when the drops were caused by the program sending packets faster than they can be buffered locally (before the packets get to the network medium).

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-servers

另请参见 [svcs\(1\)](#), [inetadm\(1M\)](#), [inetd\(1M\)](#), [spray\(1M\)](#), [svcadm\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The `rpc.sprayd` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/network/rpc/spray:default`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). Responsibility for initiating and restarting this service is delegated to [inetd\(1M\)](#). Use [inetadm\(1M\)](#) to make configuration changes and to view configuration information for this service. The service's status can be queried using the [svcs\(1\)](#) command.

引用名 `rpc.yppasswdd`, `yppasswdd` – server for modifying NIS password file

用法概要

```

/usr/lib/netsvc/yp/rpc.yppasswdd [-D directory]
    [-nogecos] [-noshell] [-nopw]
    [-m argument1 argument2...]

/usr/lib/netsvc/yp/rpc.yppasswdd
    [passwordfile adjunctfile]
    [-nogecos] [-noshell] [-nopw]
    [-m argument1 argument2...]

```

描述

`rpc.yppasswdd` is a server that handles password change requests from [yppasswd\(1\)](#). It changes a password entry in the `passwd`, `shadow`, and `security/passwd.adjunct` files. The `passwd` and `shadow` files provide the basis for the `passwd.byname` and `passwd.byuid` maps. The `passwd.adjunct` file provides the basis for the `passwd.adjunct.byname` and `passwd.adjunct.byuid` maps. Entries in the `passwd`, `shadow` or `passwd.adjunct` files are changed only if the password presented by [yppasswd\(1\)](#) matches the encrypted password of the entry. All password files are located in the `PWDIR` directory.

If the `-D` option is given, the `passwd`, `shadow`, or `passwd.adjunct` files are placed under the directory path that is the argument to `-D`.

If the `-noshell`, `-nogecos` or `-nopw` options are given, these fields cannot be changed remotely using `chfn`, `chsh`, or [passwd\(1\)](#).

If the `-m` option is given, a [make\(1S\)](#) is performed in `/var/yp` after any of the `passwd`, `shadow`, or `passwd.adjunct` files are modified. All arguments following the flag are passed to `make`.

The second of the listed syntaxes is provided only for backward compatibility. If the second syntax is used, the `passwordfile` is the full pathname of the password file and `adjunctfile` is the full pathname of the optional `passwd.adjunct` file. If a shadow file is found in the same directory as `passwordfile`, the `shadowfile` is used as described above. Use of this syntax and the discovery of a `shadowfile` file generates diagnostic output. The daemon, however, starts normally.

The first and second syntaxes are mutually exclusive. You cannot specify the full pathname of the `passwd`, `passwd.adjunct` files and use the `-D` option at the same time.

The daemon is started automatically on the master server of the `passwd` map by [ypstart\(1M\)](#), which is invoked at boot time by the `svcs:/network/nis/server:default` service.

The server does not insist on the presence of a shadow file unless there is no `-D` option present or the directory named with the `-D` option is `/etc`. In addition, a `passwd.adjunct` file is not necessary. If the `-D` option is given, the server attempts to find a `passwd.adjunct` file in the `security` subdirectory of the named directory. For example, in the presence of `-D /var/yp` the server checks for a `/var/yp/security/passwd.adjunct` file.

If only a `passwd` file exists, then the encrypted password is expected in the second field. If both a `passwd` and a `passwd.adjunct` file exist, the encrypted password is expected in the second field of the adjunct file with `##username` in the second field of the `passwd` file. If all three files are in use, the encrypted password is expected in the shadow file. Any deviation causes a password update to fail.

If you remove or add a shadow or `passwd.adjunct` file after `rpc.yppasswdd` has started, you must stop and restart the daemon to enable it to recognize the change. See [ypstart\(1M\)](#) for information on restarting the daemon.

The `rpc.yppasswdd` daemon considers a shell that has a name that begins with 'r' to be a restricted shell. By default, the daemon does not check whether a shell begins with an 'r'. However, you can tell it to do so by uncommenting the `check_restricted_shell_name=1` line in `/etc/default/yppasswdd`. The result will be to restrict a user's ability to change from his default shell. See [yppasswdd\(4\)](#).

On start up, `yppasswdd` checks for the existence of a NIS to LDAP (N2L) configuration file, `/var/yp/NISLDAPmapping`. If the configuration file is present, the daemon runs in N2L mode. If the file is not present, `yppasswdd` runs in traditional, non-N2L mode.

In N2L mode, changes are written directly to the Directory Information Tree (DIT). If the changes are written successfully, the NIS map is updated. The NIS source files, `passwd`, `shadow`, and `passwd.adjunct`, for example, are not updated. Thus, in N2L mode, the `-D` option is meaningless. In N2L mode, `yppasswdd` propagates changes by calling [yppush\(1M\)](#) instead of [ypmake\(1M\)](#). The `-m` option is thus unused.

During an NIS-to-LDAP transition, the `yppasswdd` daemon uses the N2L-specific map, `ageing.byname`, to read and write password aging information to the DIT. If you are not using password aging, then the `ageing.byname` mapping is ignored.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/nis

另请参见 [svcs\(1\)](#), [make\(1S\)](#), [passwd\(1\)](#), [yppasswdd\(1\)](#), [inetd\(1M\)](#), [svcadm\(1M\)](#), [ypmake\(1M\)](#), [yppush\(1M\)](#), [ypstart\(1M\)](#), [NISLDAPmapping\(4\)](#), [passwd\(4\)](#), [shadow\(4\)](#), [ypfiles\(4\)](#), [yppasswdd\(4\)](#), [ypserv\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 If `make` has not been installed and the `-m` option is given, the daemon outputs a warning and proceeds, effectively ignoring the `-m` flag.

When using the `-D` option, you should make sure that the `PWDIR` of the `/var/yp/Makefile` is set accordingly.

The second listed syntax is supplied only for backward compatibility and might be removed in a future release of this daemon.

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications PLC, and cannot be used without permission.

The NIS server service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svcs:/network/nis/server:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	rpc.yppupdated, yppupdated – server for changing NIS information				
用法概要	<code>/usr/lib/netsvc/yp/rpc.yppupdated [-is]</code>				
描述	yppupdated is a daemon that updates information in the Network Information Service (NIS). yppupdated consults the updaters(4) file in the <code>/var/yp</code> directory to determine which NIS maps should be updated and how to change them. By default, the daemon requires the most secure method of authentication available to it, either DES (secure) or UNIX (insecure). On start up, yppupdated checks for the existence of a NIS to LDAP (N2L) configuration file, <code>/var/yp/NISLDAPmapping</code>. If the file is present, yppupdated generates an informational message and exits. yppupdated is not supported in N2L mode.				
选项	-i Accept RPC calls with the insecure AUTH_UNIX credentials. This allows programmatic updating of the NIS maps in all networks. -s Accept only calls authenticated using the secure RPC mechanism (AUTH_DES authentication). This disables programmatic updating of the NIS maps unless the network supports these calls.				
文件	<table><tr><td><code>/var/yp/updaters</code></td><td>Configuration file for <code>rpc.updated</code> command.</td></tr><tr><td><code>/var/yp/NISLDAPmapping</code></td><td>Configuration file for N2L</td></tr></table>	<code>/var/yp/updaters</code>	Configuration file for <code>rpc.updated</code> command.	<code>/var/yp/NISLDAPmapping</code>	Configuration file for N2L
<code>/var/yp/updaters</code>	Configuration file for <code>rpc.updated</code> command.				
<code>/var/yp/NISLDAPmapping</code>	Configuration file for N2L				
属性	See attributes(5) for descriptions of the following attributes:				

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/nis
Interface Stability	Committed

另请参见 [keyserv\(1M\)](#), [updaters\(4\)](#), [NISLDAPmapping\(4\)](#), [attributes\(5\)](#)

《Oracle Solaris Administration: Naming and Directory Services》

附注

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two services remains the same. Only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications PLC, and it must not be used without permission.

引用名	rquotad – remote quota server
用法概要	/usr/lib/nfs/rquotad
描述	rquotad is an rpc(4) server which returns quotas for a user of a local file system which is mounted by a remote machine over the NFS. The results are used by quota(1M) to display user quotas for remote file systems. The rquotad daemon is normally invoked by inetd(1M) .
用法	See largefile(5) for the description of the behavior of rquotad when encountering files greater than or equal to 2 Gbyte (2 ³¹ bytes).
文件	quotas quota file at a UFS file system root
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/file-system/nfs

另请参见	svcs(1) , automountd(1M) , inetadm(1M) , inetd(1M) , mount_nfs(1M) , quota(1M) , share_nfs(1M) , svcadm(1M) , rpc(4) , services(4) , attributes(5) , largefile(5) , smf(5) 《安装 Oracle Solaris 11 系统》
附注	The rquotad service is managed by the service management facility, smf(5), under the service identifier: <code>svc:/network/nfs/rquota</code> Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M). Responsibility for initiating and restarting this service is delegated to inetd(1M). Use inetadm(1M) to make configuration changes and to view configuration information for this service. The service's status can be queried using the svcs(1) command. If it is disabled, it is enabled by mount_nfs(1M), share_nfs(1M), and automountd(1M) unless its application/auto_enable property is set to false.

引用名

rsh, restricted_shell – restricted shell command interpreter

用法概要

```
/usr/lib/rsh [-acefhiknpirstuvx] [argument]...
```

描述

rsh is a limiting version of the standard command interpreter sh, used to restrict logins to execution environments whose capabilities are more controlled than those of sh (see [sh\(1\)](#) for complete description and usage).

When the shell is invoked, it scans the environment for the value of the environmental variable, SHELL. If it is found and rsh is the file name part of its value, the shell becomes a restricted shell.

The actions of rsh are identical to those of sh, except that the following are disallowed:

- changing directory (see [cd\(1\)](#)),
- setting the value of \$PATH,
- specifying path or command names containing /,
- redirecting output (> and >>).

The restrictions above are enforced after *.profile* is interpreted.

A restricted shell can be invoked in one of the following ways:

1. rsh is the file name part of the last entry in the */etc/passwd* file (see [passwd\(4\)](#));
2. the environment variable SHELL exists and rsh is the file name part of its value; the environment variable SHELL needs to be set in the *.login* file;
3. the shell is invoked and rsh is the file name part of argument 0;
4. the shell is invoked with the *-r* option.

When a command to be executed is found to be a shell procedure, rsh invokes sh to execute it. Thus, it is possible to provide to the end-user shell procedures that have access to the full power of the standard shell, while imposing a limited menu of commands; this scheme assumes that the end-user does not have write and execute permissions in the same directory.

The net effect of these rules is that the writer of the *.profile* (see [profile\(4\)](#)) has complete control over user actions by performing guaranteed setup actions and leaving the user in an appropriate directory (probably *not* the login directory).

The system administrator often sets up a directory of commands (that is, */usr/rbin*) that can be safely invoked by a restricted shell. Some systems also provide a restricted editor, *red*.

退出状态

Errors detected by the shell, such as syntax errors, cause the shell to return a non-zero exit status. If the shell is being used non-interactively execution of the shell file is abandoned. Otherwise, the shell returns the exit status of the last command executed.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [Intro\(1\)](#), [cd\(1\)](#), [login\(1\)](#), [rsh\(1\)](#), [sh\(1\)](#), [exec\(2\)](#), [passwd\(4\)](#), [profile\(4\)](#), [attributes\(5\)](#)

附注 The restricted shell, `/usr/lib/rsh`, should not be confused with the remote shell, `/usr/bin/rsh`, which is documented in [rsh\(1\)](#).

引用名 `rtc` – provide all real-time clock and UTC-lag management

用法概要 `/usr/sbin/rtc [-c] [-z zone-name]`

描述 On x86 systems, the `rtc` command reconciles the difference in the way that time is established between UNIX and MS-DOS systems. UNIX systems utilize Universal Coordinated Time (UTC), while MS-DOS systems utilize local time.

Without arguments, `rtc` displays the currently configured time zone string. The currently configured time zone string is based on what was last recorded by `rtc -z zone-name`.

The `rtc` command is not normally run from a shell prompt; it is generally invoked by the system. Commands such as [date\(1\)](#) and [rdate\(1M\)](#), which are used to set the time on a system, invoke `/usr/sbin/rtc -c` to ensure that daylight savings time (DST) is corrected for properly.

- 选项**
- `-c` This option checks for DST and makes corrections if necessary. It is normally run once a day by a cron job.

If there is no RTC time zone or `/etc/rtc_config` file, this option will do nothing.
 - `-z zone-name` This option, which is normally run by the system at software installation time, is used to specify the time zone in which the RTC is to be maintained. It updates the configuration file `/etc/rtc_config` with the name of the specified zone and the current UTC lag for that zone. If there is an existing `rtc_config` file, this command will update it. If not, this command will create it.

文件 `/etc/rtc_config` The data file used to record the time zone and UTC lag. This file is completely managed by `/usr/sbin/rtc`, and it is read by the kernel.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Architecture	x86
Availability	system/core-os

另请参见 [date\(1\)](#), [rdate\(1M\)](#), [attributes\(5\)](#)

引用名

用法概要

描述

选项

rtquery – query routing daemons for their routing tables

rtquery [-np1] [-w timeout] [-r addr] [-a secret] host...

rtquery [-t operation] host...

The rtquery command is used to query a RIP network routing daemon, [in.routed\(1M\)](#) or GateD, for its routing table by sending a request or poll command. The routing information in any routing response packets returned is displayed numerically and symbolically.

By default, rtquery uses the request command. When the -p option is specified, rtquery uses the poll command, an undocumented extension to the RIP protocol supported by GateD. When querying GateD, the poll command is preferred over the request command because the response is not subject to Split Horizon and/or Poisoned Reverse, and because some versions of GateD do not answer the request command. in.routed does not answer the poll command, but recognizes requests coming from rtquery and so answers completely.

The rtquery command is also used to turn tracing on or off in in.routed.

The following options are supported:

-a passwd=XXX

-a md5_passwd=XXX|KeyID Causes the query to be sent with the indicated cleartext or MD5 password.

-n Displays only the numeric network and host addresses instead of both numeric and symbolic names.

-p Uses the poll command to request full routing information from GateD. This is an undocumented extension RIP protocol supported only by GateD.

-r addr Asks about the route to destination addr.

-t operation Changes tracing, where operation is one of the actions listed below. Requests from processes not running with UID 0 or on distant networks are generally ignored by the daemon except for a message in the system log. GateD is likely to ignore these debugging requests.

on=tracefile Turns tracing on, directing tracing into the specified file. That file must have been specified when the daemon was started or have the name, /var/log/in.routed.trace.

more Increases the debugging level.

off Turns off tracing.

- dump

Dumps the daemon's routing table to the current trace file.
- w *timeout*

Changes the delay for an answer from each host. By default, each host is given 15 seconds to respond.
- 1

Queries using RIP version 1 instead of RIP version 2.

退出状态

The following exit values are returned:

0

Successful completion.

>0

An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/routing

另请参见

[in.routed\(1M\)](#), [route\(1M\)](#), [gateways\(4\)](#), [attributes\(5\)](#), [icmp\(7P\)](#), [inet\(7P\)](#), [udp\(7P\)](#)

Routing Information Protocol, RIPv1, RFC 1058

Routing Information Protocol, RIPv2, RFC 2453, STD 0056

引用名	runacct – run daily accounting																						
用法概要	<code>/usr/lib/acct/runacct [mmdd [state]]</code>																						
描述	runacct is the main daily accounting shell procedure. It is normally initiated using <code>cron</code>. runacct processes connect, fee, disk, and process accounting files. It also prepares summary files for <code>prdaily</code> or billing purposes. runacct is distributed only to source code licensees. runacct takes care not to damage active accounting files or summary files in the event of errors. It records its progress by writing descriptive diagnostic messages into <code>active</code>. When an error is detected, a message is written to <code>/dev/console</code>, mail (see mail(1)) is sent to <code>root</code> and <code>adm</code>, and runacct terminates. runacct uses a series of lock files to protect against re-invocation. The files <code>lock</code> and <code>lock1</code> are used to prevent simultaneous invocation, and <code>lastdate</code> is used to prevent more than one invocation per day. runacct breaks its processing into separate, restartable <i>states</i> using <code>statefile</code> to remember the last <i>state</i> completed. It accomplishes this by writing the <i>state</i> name into <code>statefile</code>. runacct then looks in <code>statefile</code> to see what it has done and to determine what to process next. <i>states</i> are executed in the following order: <table><tr><td>SETUP</td><td>Move active accounting files into working files.</td></tr><tr><td>WTMPFIX</td><td>Verify integrity of <code>wtmpx</code> file, correcting date changes if necessary.</td></tr><tr><td>CONNECT</td><td>Produce connect session records in <code>tacct.h</code> format.</td></tr><tr><td>PROCESS</td><td>Convert process accounting records into <code>tacct.h</code> format.</td></tr><tr><td>MERGE</td><td>Merge the connect and process accounting records.</td></tr><tr><td>FEES</td><td>Convert output of chargefee into <code>tacct.h</code> format, merge with connect, and process accounting records.</td></tr><tr><td>DISK</td><td>Merge disk accounting records with connect, process, and fee accounting records.</td></tr><tr><td>MERGETACCT</td><td>Merge the daily total accounting records in <code>daytacct</code> with the summary total accounting records in <code>/var/adm/acct/sum/tacct</code>.</td></tr><tr><td>CMS</td><td>Produce command summaries.</td></tr><tr><td>USEREXIT</td><td>Any installation dependent accounting programs can be included here.</td></tr><tr><td>CLEANUP</td><td>Clean up temporary files and exit. To restart runacct after a failure, first check the <code>active</code> file for diagnostics, then fix any corrupted data files, such as <code>pacct</code> or <code>wtmpx</code>. The <code>lock</code>, <code>lock1</code>, and <code>lastdate</code> files must be removed before runacct can be restarted. The argument <i>mmdd</i> is necessary if runacct is being restarted. <i>mmdd</i> specifies the month and day for which runacct will rerun the accounting. The entry point for processing is based on the contents of <code>statefile</code>; to override this, include the desired <i>state</i> on the command line to designate where processing should begin.</td></tr></table>	SETUP	Move active accounting files into working files.	WTMPFIX	Verify integrity of <code>wtmpx</code> file, correcting date changes if necessary.	CONNECT	Produce connect session records in <code>tacct.h</code> format.	PROCESS	Convert process accounting records into <code>tacct.h</code> format.	MERGE	Merge the connect and process accounting records.	FEES	Convert output of chargefee into <code>tacct.h</code> format, merge with connect, and process accounting records.	DISK	Merge disk accounting records with connect, process, and fee accounting records.	MERGETACCT	Merge the daily total accounting records in <code>daytacct</code> with the summary total accounting records in <code>/var/adm/acct/sum/tacct</code> .	CMS	Produce command summaries.	USEREXIT	Any installation dependent accounting programs can be included here.	CLEANUP	Clean up temporary files and exit. To restart runacct after a failure, first check the <code>active</code> file for diagnostics, then fix any corrupted data files, such as <code>pacct</code> or <code>wtmpx</code> . The <code>lock</code> , <code>lock1</code> , and <code>lastdate</code> files must be removed before runacct can be restarted. The argument <i>mmdd</i> is necessary if runacct is being restarted. <i>mmdd</i> specifies the month and day for which runacct will rerun the accounting. The entry point for processing is based on the contents of <code>statefile</code> ; to override this, include the desired <i>state</i> on the command line to designate where processing should begin.
SETUP	Move active accounting files into working files.																						
WTMPFIX	Verify integrity of <code>wtmpx</code> file, correcting date changes if necessary.																						
CONNECT	Produce connect session records in <code>tacct.h</code> format.																						
PROCESS	Convert process accounting records into <code>tacct.h</code> format.																						
MERGE	Merge the connect and process accounting records.																						
FEES	Convert output of chargefee into <code>tacct.h</code> format, merge with connect, and process accounting records.																						
DISK	Merge disk accounting records with connect, process, and fee accounting records.																						
MERGETACCT	Merge the daily total accounting records in <code>daytacct</code> with the summary total accounting records in <code>/var/adm/acct/sum/tacct</code> .																						
CMS	Produce command summaries.																						
USEREXIT	Any installation dependent accounting programs can be included here.																						
CLEANUP	Clean up temporary files and exit. To restart runacct after a failure, first check the <code>active</code> file for diagnostics, then fix any corrupted data files, such as <code>pacct</code> or <code>wtmpx</code> . The <code>lock</code> , <code>lock1</code> , and <code>lastdate</code> files must be removed before runacct can be restarted. The argument <i>mmdd</i> is necessary if runacct is being restarted. <i>mmdd</i> specifies the month and day for which runacct will rerun the accounting. The entry point for processing is based on the contents of <code>statefile</code> ; to override this, include the desired <i>state</i> on the command line to designate where processing should begin.																						

示例

示例1 Starting runacct

The following example starts runacct:

```
example% nohup runacct 2> /var/adm/acct/nite/fd2log &
```

示例2 Restarting runacct

The following example restarts runacct:

```
example% nohup runacct 0601 2>> /var/adm/acct/nite/fd2log &
```

示例3 Restarting runacct at a Specific State

The following example restarts runacct at a specific state:

```
example% nohup runacct 0601 MERGE 2>> /var/adm/acct/nite/fd2log &
```

文件

/var/adm/wtmpx History of user access and administration information

/var/adm/pacctincr

/var/adm/acct/nite/active

/var/adm/acct/nite/daytacct

/var/adm/acct/nite/lock

/var/adm/acct/nite/lock1

/var/adm/acct/nite/lastdate

/var/adm/acct/nite/statefile

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/accounting/legacy-accounting

另请参见

[acctcom\(1\)](#), [mail\(1\)](#), [acct\(1M\)](#), [acctcms\(1M\)](#), [acctcon\(1M\)](#), [acctmerg\(1M\)](#), [acctprc\(1M\)](#), [acctsh\(1M\)](#), [cron\(1M\)](#), [fwtmp\(1M\)](#), [acct\(2\)](#), [acct.h\(3HEAD\)](#), [utmpx\(4\)](#), [attributes\(5\)](#)

附注

It is not recommended to restart runacct in the *SETUP state*. Run SETUP manually and restart using:

```
runacct mmd WTMPFIX
```

If runacct failed in the *PROCESS state*, remove the last ptacct file because it will not be complete.

The runacct command can process a maximum of

- 6000 distinct sessions

- 1000 distinct terminal lines
- 2000 distinct login names

during a single invocation of the command. If at some point the actual number of any one of these items exceeds the maximum, the command will not succeed.

Do not invoke `runacct` at the same time as `ckpacct`, as there may be a conflict if both scripts attempt to execute `turnacct switch` simultaneously.

引用名

rwall – write to all users over a network

用法概要

```
/usr/sbin/rwall hostname...  
  
/usr/sbin/rwall -n netgroup...  
  
/usr/sbin/rwall -h hostname -n netgroup
```

描述

rwall reads a message from standard input until EOF. It then sends this message, preceded by the line:

Broadcast Message . . .

to all users logged in on the specified host machines. With the -n option, it sends to the specified network groups.

选项

-n netgroup

Send the broadcast message to the specified network groups.

-h hostname

Specify the hostname, the name of the host machine.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-clients

另请参见

[inetd\(1M\)](#), [wall\(1M\)](#), [attributes\(5\)](#)

附注

The timeout is fairly short to allow transmission to a large group of machines (some of which may be down) in a reasonable amount of time. Thus the message may not get through to a heavily loaded machine.

引用名	sar, sa1, sa2, sadc – system activity report package	
用法概要	<pre>/usr/lib/sa/sadc [t n] [ofile] /usr/lib/sa/sa1 [t n] /usr/lib/sa/sa2 [-aAbcdgkmpqruvwy] [-e time] [-f filename] [-i sec] [-s time]</pre>	
描述	System activity data can be accessed at the special request of a user (see sar(1)) and automatically, on a routine basis, as described here. The operating system contains several counters that are incremented as various system actions occur. These include counters for CPU utilization, buffer usage, disk and tape I/O activity, TTY device activity, switching and system-call activity, file-access, queue activity, inter-process communications, and paging. For more general system statistics, use iostat(1M), sar(1), or vmstat(1M). sadc and two shell procedures, sa1 and sa2, are used to sample, save, and process this data. sadc, the data collector, samples system data <i>n</i> times, with an interval of <i>t</i> seconds between samples, and writes in binary format to <i>ofile</i> or to standard output. The sampling interval <i>t</i> should be greater than 5 seconds; otherwise, the activity of sadc itself may affect the sample. If <i>t</i> and <i>n</i> are omitted, a special record is written. This facility can be used at system boot time, when booting to a multi-user state, to mark the time at which the counters restart from zero. For example, when accounting is enabled, the svc: /system/sar:default service writes the restart mark to the daily data file using the command entry: <pre>su sys -c "/usr/lib/sa/sadc /var/adm/sa/sa'date +%d"</pre> The shell script sa1, a variant of sadc, is used to collect and store data in the binary file /var/adm/sa/sadd, where dd is the current day. The arguments <i>t</i> and <i>n</i> cause records to be written <i>n</i> times at an interval of <i>t</i> seconds, or once if omitted. The following entries in /var/spool/cron/crontabs/sys will produce records every 20 minutes during working hours and hourly otherwise: <pre>0 * * * 0-6 /usr/lib/sa/sa1 20,40 8-17 * * 1-5 /usr/lib/sa/sa1</pre> See crontab(1) for details. The shell script sa2, a variant of sar, writes a daily report in the file /var/adm/sa/sardd. See the OPTIONS section in sar(1) for an explanation of the various options. The following entry in /var/spool/cron/crontabs/sys will report important activities hourly during the working day: <pre>5 18 * * 1-5 /usr/lib/sa/sa2 -s 8:00 -e 18:01 -i 1200 -A</pre>	
文件	<pre>/tmp/sa.adrfl</pre>	address file
	<pre>/var/adm/sa/sadd</pre>	daily data file
	<pre>/var/adm/sa/sardd</pre>	daily report file

/var/spool/cron/crontabs/sys used for performance collection

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/accounting/legacy-accounting

另请参见 [crontab\(1\)](#), [sar\(1\)](#), [svcs\(1\)](#), [timex\(1\)](#), [iostat\(1M\)](#), [svcadm\(1M\)](#), [vmstat\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

 《Oracle Solaris 管理：常见任务》

附注 The sar service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/system/sar

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	sasinfo – Serial Attached SCSI HBA 端口命令行接口
用法概要	<pre> sasinfo hba [-v] [HBA_Name]... sasinfo hba-port [-lvy] [-a HBA_Name] [HBA_port_name]... sasinfo expander [-vt] [-p HBA_port_Name] [Expander_SAS_Addr]... sasinfo target-port [-s -v] [Target_port_SAS_Addr]... sasinfo logical-unit lu [-v] [device_path]... sasinfo [-V] sasinfo [-?]</pre>
描述	sasinfo 实用程序是一个命令行接口，用于收集有关支持存储管理 HBA API (Storage Management HBA API, SM-HBA) 的 Serial Attached SCSI-2 (SAS-2) 主机总线适配器 (host bus adapter, HBA) 的管理信息。此实用程序报告 HBA 端口和扩展器设备的属性，以及可以连接到那些 HBA 端口的 SCSI 目标设备。 sasinfo 是作为一组子命令实现的，如下所述。
子命令	sasinfo 支持下列子命令。 hba 列出指定的 <i>HBA_name</i> 所指的 HBA 的信息。如果未指定 <i>HBA_name</i>，则会列出主机上的所有 HBA。 hba-port 列出指定的 <i>HBA_port_Name</i> 所指的 HBA 端口的信息。如果未指定 <i>HBA_port_Name</i>，则会列出主机上的所有 HBA 端口。请注意，HBA 端口可以动态配置/取消配置，因此，HBA 上的 HBA 端口数量可能会随着与目标设备的连接的建立而改变。如果未建立与存储设备的连接，则不会配置 HBA 端口。 expander 列出指定的 <i>Expander_SAS_Addr</i> 所指的扩展器的分层视图。如果未指定任何参数，则会列出在主机上的所有 HBA 端口上可见的所有扩展器。 target-port 列出指定的 <i>Target_port_SAS_Addr</i> 地址所指的那些目标端口的信息。如果未指定任何参数，则会列出在主机上可见的所有目标端口。 logical-unit lu 列出指定的 <i>device_path</i> 所指的逻辑单元。如果未指定 <i>device_path</i>，则会列出所有 SAS 逻辑单元，包括 SAS 管理协议 (SAS Management Protocol, SMP) 目标端口。
选项	支持以下选项。 -?, --help 显示用法信息

-a、**--hba *HBA_Name***

通过主机上的 HBA 的 *HBA_Name* 检索 HBA 端口信息。-a 选项仅可与 **hba-port** 子命令一起使用。

-l、**--phy-linkstat**

为指定的 *HBA_port_Name* 所指的 HBA 端口或所有 HBA 端口（如果未指定 HBA 端口）上的 phy 列出链路错误统计信息。此选项仅可与 **hba-port** 子命令一起使用。

-p *HBA_port_SAS_Addr*、**--port *HBA_port_SAS_Addr***

通过主机上的本地 HBA 端口的 *HBA_port_SAS_Addr* 检索远程端口信息。-p 选项仅可与 **expander** 子命令一起使用。

-s、**--scsi**

列出请求显示的目标端口的 SCSI 属性。此选项仅可用于 **target-port** 子命令。

-v、**--verbose**

根据所调用的子命令显示硬件的详细信息，例如 SAS 地址、拓扑设备，等等。

-V、**--version**

显示版本信息。

-y、**--phy**

列出 *HBA_port_Name* 指定的 HBA 端口或所有 HBA 端口（如果未指定 HBA 端口）上的 phy 信息。此选项仅可与 **hba-port** 子命令一起使用。

示例

示例1 列出所有 HBA

以下命令列出主机上的所有 HBA。

```
# sasinfo hba
HBA Name: SUNW-pmcs-0
HBA Name: SUNW-pmcs-1
```

示例2 详细列出所有 HBA

以下命令列出主机上的所有 HBA，并包括相关详细信息。请注意，每个 HBA 配置有两个 HBA 端口，每个 HBA 端口连接到一个存储设备。

```
# sasinfo hba -v
HBA Name: SUNW-pmcs-0
  Manufacturer: sun
  Model: SAS Gen-2
  Firmware Version: 1.1
  FCode/BIOS Version: 1.1
  Serial Number: 111-11111
  Driver Name: smvsl
  Driver Version: 1.1
  Number of HBA Ports: 2
HBA Name: SUNW-pmcs-1
  Manufacturer: sun
  Model: SAS Gen-2
```

示例2 详细列出所有 HBA (续)

```
Firmware Version: 1.1
FCode/BIOS Version: 1.1
Serial Number: 111-11111
Driver Name: smvsl
Driver Version: 1.1
Number of HBA Ports: 2
```

示例3 列出所有 HBA 端口

以下命令列出主机上的所有 HBA 端口。

```
# sasinfo hba-port
HBA Name: SUNW-pmcs-0
  HBA Port Name: /dev/cfg/c1
  HBA Port Name: /dev/cfg/c2
HBA Name: SUNW-pmcs-1
  HBA Port Name: /dev/cfg/c3
  HBA Port Name: /dev/cfg/c4
```

示例4 详细列出所有 HBA 端口

以下命令列出主机上的所有 HBA 端口，并包括相关详细信息。

```
# sasinfo hba-port -v
HBA Name: SUNW-pmcs-0
  HBA Port Name: /dev/cfg/c1
    Type: sas-device
    State: online
    Local SAS Address: 5000c50000d756aa
    Attached Port SAS Address: 50800201a5a502bf
    Number of Phys: 4
  HBA Port Name: /dev/cfg/c25000c50000d756cc
    Type: sas-device
    State: online
    Local SAS Address: 5000c50000d756aa
    Attached Port SAS Address: 50800201a5a503bf
    Number of Phys: 4
HBA Name: SUNW-pmcs-1
  HBA Port Name: /dev/cfg/c3
    Type: sas-device
    State: online
    Local SAS Address: 5000c50000d756cc
    Attached Port SAS Address: 50800201a5a504bf
    Number of Phys: 4
  HBA Port Name: /dev/cfg/c4
    Type: sas-device
    State: online
    Local SAS Address: 5000c50000d756cc
```

示例 4 详细列出所有 HBA 端口 (续)

```
Attached Port SAS Address: 50800201a5a505bf
Number of Phys: 4
```

示例 5 列出所有 HBA 端口的 phy 信息

以下命令列出所有 HBA 端口的 phy 信息。

```
# sasinfo hba-port -y
HBA Name: SUNW-pmcs-0
  HBA Port Name: /dev/cfg/c1
    Phy Information:
      Identifier: 0
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit
      Identifier: 1
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit
      Identifier: 2
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit
      Identifier: 3
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit
  HBA Port Name: /dev/cfg/c2
    Phy Information:
      Identifier: 4
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit
      Identifier: 5
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit
      Identifier: 6
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
```

示例 5 列出所有 HBA 端口的 phy 信息 (续)

```

        NegotiatedLinkRate: 3Gbit
    Identifier: 7
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit
HBA Name: SUNW-pmcs-1
HBA Port Name: /dev/cfg/c3
Phy Information:
    Identifier: 0
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit
    Identifier: 1
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit
    Identifier: 2
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit
    Identifier: 3
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit
HBA Port Name: /dev/cfg/c4
Phy Information:
    Identifier: 4
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit
    Identifier: 5
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit
    Identifier: 6
        State: enabled
        HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
        ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
        NegotiatedLinkRate: 3Gbit

```

示例 5 列出所有 HBA 端口的 phy 信息 (续)

```
Identifier: 7
State: enabled
HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
NegotiatedLinkRate: 3Gbit
```

示例 6 列出特定 HBA 端口的 phy 链路错误统计信息
以下命令列出一个特定端口的 phy 链路错误统计信息。

```
# sasinfo hba-port -ly /dev/cfg/c1
HBA Name: SUNW-pmcs-0
HBA Port Name: /dev/cfg/c1
Phy Information:
Identifier: 0
State: enabled
HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
NegotiatedLinkRate: 3Gbit
Link Error Statistics:
Invalid Dword: 0
Running Disparity Error: 0
Loss of Dword Sync: 0
Reset Problem: 0
Identifier: 1
State: enabled
HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
NegotiatedLinkRate: 3Gbit
Link Error Statistics:
Invalid Dword: 0
Running Disparity Error: 0
Loss of Dword Sync: 0
Reset Problem: 0
Identifier: 2
State: enabled
HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
NegotiatedLinkRate: 3Gbit
Link Error Statistics:
Invalid Dword: 0
Running Disparity Error: 0
Loss of Dword Sync: 0
Reset Problem: 0
Identifier: 3
State: enabled
HardwareLinkRate(Min/Max): 1.5Gbit/3Gbit
```


示例 6 列出特定 HBA 端口的 phy 链路错误统计信息 (续)

```

ProgrammedLinkRate(Min/Max): 1.5Gbit/3Gbit
NegotiatedLinkRate: 3Gbit
Link Error Statistics:
  Invalid Dword: 0
  Running Disparity Error: 0
  Loss of Dword Sync: 0
  Reset Problem: 0

```

示例 7 列出连接到 HBA 端口的扩展器

以下命令列出在指定 HBA 端口上可见的所有扩展器。

```

# sasinfo expander -p /dev/cfg/c1
HBA Name: SUNW-pmcs-0
HBA Port Name: /dev/cfg/c1
  Expander SAS Address(Tier 1): 50800201a5a502bf
    Expander SAS Address(Tier 2): 50800201a5a50233
      Expander SAS Address(Tier 2); 5000c5000d2da812
    Expander SAS Address(Tier 1): 50800201a5a503bf
      Expander SAS Address(Tier 2): 50800201a5a502d2
        Expander SAS Address(Tier 2); 5000c5000d2da823

```

示例 8 列出扩展器的详细信息

以下命令列出在指定 HBA 端口上可见的所有扩展器的详细信息。

```

# sasinfo expander -v -p /dev/cfg/c1
HBA Name: SUNW-pmcs-0
HBA Port Name: /dev/cfg/c1
  Expander SAS Address(Tier 1): 50800201a5a502bf
    OS Device Name: /dev/smp/expd0
      State: online
        Expander SAS Address(Tier 2): 50800201a5a50233
          OS Device Name: /dev/smp/expd2
            State: online
              Expander SAS Address(Tier 2); 5000c5000d2da812
                OS Device Name: /dev/smp/expd3
                  State: online
  Expander SAS Address(Tier 1): 50800201a5a503bf
    OS Device Name: /dev/smp/expd1
      State: online
        Expander SAS Address(Tier 2): 50800201a5a502d2
          OS Device Name: /dev/smp/expd3
            State: online
              Expander SAS Address(Tier 2); 5000c5000d2da823
                OS Device Name: /dev/smp/expd4
                  State: online

```

示例9 列出连接到扩展器的目标端口

以下命令列出连接到指定 HBA 端口的扩展器上连接的所有目标端口。

```
# sasinfo expander -t -p /dev/cfg/c1
HBA Name: SUNW-pmcs-0
  HBA Port Name: /dev/cfg/c1
    Expander SAS Address(Tier 1): 50800201a5a502bf
      Target Port SAS Address: 50800201a5a504f1
        Expander SAS Address(Tier 2): 50800201a5a50233
          Target Port SAS Address: 50800201a5a502be
            Target Port SAS Address: 5000c5000d2da8b2
          Expander SAS Address(Tier 2); 5000c5000d2da812
            Target Port SAS Address: 50800201a5a502be
              Target Port SAS Address: 50800201a5a508b2
        Expander SAS Address(Tier 1): 50800201a5a503bf
          Target Port SAS Address: 50800201a5a50421
            Expander SAS Address(Tier 2): 50800201a5a502d2
              Target Port SAS Address: 50800201a5a503be
                Target Port SAS Address: 5000c5000d2da7be
              Expander SAS Address(Tier 2); 5000c5000d2da823
                Target Port SAS Address: 50800201a5a503be
                  Target Port SAS Address: 5000c5000d2da7be
```

示例10 列出目标端口信息

以下命令列出在主机上搜索到所有目标端口。

```
# sasinfo target-port
Target Port SAS Address: 5000c5000bae4009
Target Port SAS Address: 5000c5000baef4b1
Target Port SAS Address: 5000c5000bae3fe1
Target Port SAS Address: 5000c5000bae49d9
Target Port SAS Address: 5000c5000bae36c5
```

示例11 列出包括拓扑详细信息在内的目标端口信息

以下命令列出所有目标端口以及它们所连接到的 HBA 端口和扩展器。

```
# sasinfo target-port -v
Target Port SAS Address: 5000c5000bae4009
  Type: SAS Device
    HBA Port Name: /dev/cfg/c7
      Expander Device SAS Address: 500e004aaaaaaa3f
        HBA Port Name: /dev/cfg/c5
          Expander Device SAS Address: 500e004aaaaaaa3f
Target Port SAS Address: 5000c5000baef4b1
  Type: SAS Device
    HBA Port Name: /dev/cfg/c7
      Expander Device SAS Address: 500e004aaaaaaa3f
        HBA Port Name: /dev/cfg/c5
```

示例 11 列出包括拓扑详细信息在内的目标端口信息 (续)

```

    Expander Device SAS Address: 500e004aaaaaaa3f
Target Port SAS Address: 5000c5000bae3fe1
Type: SAS Device
HBA Port Name: /dev/cfg/c7
    Expander Device SAS Address: 500e004aaaaaaa3f
HBA Port Name: /dev/cfg/c5
    Expander Device SAS Address: 500e004aaaaaaa3f
Target Port SAS Address: 5000c5000bae49d9
Type: SAS Device
HBA Port Name: /dev/cfg/c7
    Expander Device SAS Address: 500e004aaaaaaa3f
HBA Port Name: /dev/cfg/c5
    Expander Device SAS Address: 500e004aaaaaaa3f
Target Port SAS Address: 5000c5000bae36c5
Type: SAS Device
HBA Port Name: /dev/cfg/c7
    Expander Device SAS Address: 500e004aaaaaaa3f
HBA Port Name: /dev/cfg/c5
    Expander Device SAS Address: 500e004aaaaaaa3f

```

示例 12 列出目标端口及其 SCSI 信息

以下命令列出所有目标端口详细信息，包括每个目标端口的 SCSI 信息。

```

# sasinfo target-port -s 5000c5000bae4009
Target Port SAS Address: 5000c5000bae4009
Type: SAS Device
    HBA Port Name: /dev/cfg/c7
        Expander Device SAS Address: 500e004aaaaaaa3f
            LUN : 0
                OS Device Name : /dev/rdisk/c6t5000C5000BAE400Bd0s2
                Vendor: Sun
                Product: J4400
                Device Type: Disk
    HBA Port Name: /dev/cfg/c5
        Expander Device SAS Address: 500e004aaaaaaa3f
            LUN : 0
                OS Device Name : /dev/rdisk/c6t5000C5000BAE400Bd0s2
                Vendor: Sun
                Product: J4400
                Device Type: Disk

```

示例 13 列出逻辑单元

以下命令列出主机上的逻辑单元。

```

# sasinfo logical-unit
OS Device Name: /dev/rdisk/c4t50020F2300B4904Ed0s2

```

示例 13 列出逻辑单元 (续)

OS Device Name: /dev/rdisk/c4t50020F230000B4AFd0s2

示例 14 列出逻辑单元的其他信息

以下命令使用 -v 选项列出设备 /dev/rmt/On 的其他逻辑单元信息。

```
# sasinfo lu -v
OS Device Name: /dev/rdisk/c4t50020F2300B4904Ed0s2
  HBA Port Name: /dev/cfg/c1
    Target Port SAS Address: 50020f2300b4904e
      LUN: 0
    Vendor: Sun
    Product: J4400
    Device Type: Disk
OS Device Name: /dev/rdisk/c4t50020F230000B4AFd0s2
  HBA Port Name: /dev/cfg/c1
    Target Port SAS Address: 50020f230063100b
      LUN: 0
    Vendor: Sun
    Product: J4400
    Device Type: Disk
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/storage/sas-utilities
接口稳定性	Committed（已确定）

另请参见 [attributes\(5\)](#)

引用名	savecore – save a crash dump of the operating system
用法概要	<code>/usr/bin/savecore [-Lvd] [-f <i>dumpfile</i>] [<i>directory</i>]</code>
描述	The savecore utility saves a crash dump of the kernel (assuming that one was made) and writes a reboot message in the shutdown log. It is invoked by the dumpadm service each time the system boots. savecore can be configured by dumpadm(1M) to save crash dump data in either a compressed or uncompressed format. For the compressed format, savecore saves the crash dump data in the file <i>directory/vmdump.N</i>, where <i>N</i> in the pathname is replaced by a number which increments by one each time savecore is run in <i>directory</i>. The compressed file can be uncompressed in a separate step using the <code>-f <i>dumpfile</i></code> option. For the uncompressed format, savecore saves the crash dump data in the file <i>directory/vmcore.N</i> and the kernel's namelist in <i>directory/unix.N</i>. Before writing out a crash dump, savecore reads a number from the file <i>directory/minfree</i>. This is the minimum number of kilobytes that must remain free on the file system containing <i>directory</i>. If after saving the crash dump the file system containing <i>directory</i> would have less free space the number of kilobytes specified in <i>minfree</i>, the crash dump is not saved. If the <i>minfree</i> file does not exist, savecore assumes a <i>minfree</i> value of 1 megabyte. The savecore utility also logs a reboot message using facility LOG_AUTH (see syslog(3C)). If the system crashed as a result of a panic, savecore logs the panic string too.
选项	The following options are supported: -d Disregard dump header valid flag. Force savecore to attempt to save a crash dump even if the header information stored on the dump device indicates the dump has already been saved. -f <i>dumpfile</i> Save a crash dump from the specified file instead of from the system's current dump device. When given <i>directory/vmdump.N</i>, uncompress the file to <i>vmcore.N</i> and <i>unix.N</i>, where <i>N</i> is the same number as in the compressed name. This option may also be useful if the information stored on the dump device has been copied to an on-disk file by means of the dd(1M) command. -L Save a crash dump of the live running Solaris system, without actually rebooting or altering the system in any way. This option forces savecore to save a live snapshot of the system to the dump device, and then immediately to retrieve the data and to write it out to a new set of crash dump files in the specified directory. Live system crash dumps can only be performed if you have configured your system to have a dedicated dump device using dumpadm(1M).

-v Verbose. Enables verbose error messages from savecore.

The following operands are supported:

文件

- 属性

See `attributes(5)` for descriptions of the following attributes:

另请参见

附注

If the dump device is also being used as a swap device, you must run `savecore` very soon after booting, before the swap space containing the crash dump is overwritten by programs currently running.

引用名	sbdadm – SCSI Block Disk 命令行接口
用法概要	<pre>sbdadm create-lu [-s, --size size] filename</pre> <pre>sbdadm delete-lu filename</pre> <pre>sbdadm import-lu filename</pre> <pre>sbdadm list-lu</pre> <pre>sbdadm modify-lu [-s, --size size] lu_name filename</pre>
描述	sbdadm 命令创建和管理向 SCSI 目标模式框架 (SCSI Target Mode Framework, STMF) 进行了注册的基于 SCSI 块设备的逻辑单元。
子命令	sbdadm 命令支持下面列出的子命令。注意，如果您输入一个问号作为命令的参数 (sbdadm ?)，sbdadm 将显示帮助。 create-lu [-s, --size size] filename 创建一个可以向 STMF 进行注册的逻辑单元。对于 -s 选项，size 是一个整数，后跟下列字母之一，表示大小的单位： - k 千字节 - m 兆字节 - g 千兆字节 - t 兆兆字节 - p 千兆兆字节 - e 艾字节 如果未指定 size，则大小缺省为 filename 的大小。 指定的大小可以超过文件或设备的大小。 delete-lu filename 删除使用 sbdadm create-lu 创建的现有逻辑单元。这可以有效地从 STMF 框架中卸载逻辑单元。逻辑单元上的任何现有数据将保持不变。 import-lu filename 将之前使用 sbdadm create-lu 创建的且后来又使用 sbdadm delete-lu 从 STMF 中删除的逻辑单元导入并装入到 STMF 中。成功后，该逻辑单元将重新可供 STMF 使用。filename 是在 sbdadm create-lu 命令中为该逻辑单元使用的文件名称。 list-lu 列出使用 sbdadm create-lu 命令创建的所有逻辑单元。 modify-lu [-s, --size size] lu_name filename 修改使用 sbdadm create-lu 命令创建的逻辑单元的属性。对于 -s 选项，size 是一个整数值，后跟单位说明符。单位说明符同上文的 create-lu 部分所述。指定了此选项时，逻辑单元的当前大小将更改为新的大小。

指定的大小可以超过 *lu_name* 表示的文件或设备的大小。

操作数

sbdadm 使用以下操作数：

- filename*
现有文件的名称，或原始块设备的全限定路径名。
- lu_name*
逻辑单元的 32 字节十六进制表示。

示例

示例1 创建逻辑单元
以下命令系列将创建一个 10 千兆字节的逻辑单元。

```
# touch /export/lun/0
# sbdadm create-lu -s 10g /export/lun/0
# sbdadm create-lu /dev/rdisk/clt1d0s0
```

示例2 列出逻辑单元
以下命令将列出所有逻辑单元。

```
# sbdadm list-lu
Found 2 LU(s)

          GUID          DATA SIZE    SOURCE
-----
6000ae4000144f21d92c47b0dd650002  10737352704  /export/lun/0
6000ae4000144f21d92c47b0de300032  134283264    /dev/rdisk/clt1d0s0
```

退出状态

- 0
成功完成。
- 非零
出现错误。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/storage/scsi-target-mode-framework
接口稳定性	Obsolete Committed（已过时且已确定）

另请参见

[stmfadm\(1M\)](#)、[attributes\(5\)](#)

引用名	scadm – administer System Controller (SC)
用法概要	<code>/usr/platform/platform-name/sbin/scadm subcommand [option] [argument] . . .</code>
描述	The <code>scadm</code> utility administers the System Controller (SC). This utility allows the host server to interact with the SC. The <code>scadm</code> utility <i>must</i> be run as root. The interface, output, and location in the directory hierarchy for <code>scadm</code> are uncommitted and might change. <i>platform-name</i> is the name of the platform implementation. Use the <code>uname -i</code> command to identify the platform implementation. See uname(1). The <code>scadm</code> utility has a number of subcommands. Some subcommands have specific options and arguments associated with them. See SUBCOMMANDS, OPTIONS, OPERANDS, and USAGE.
SUBCOMMANDS	Subcommands immediately follow the <code>scadm</code> command on the command line, and are separated from the command by a SPACE. The following subcommands are supported consolehistory Display the SC's console log. The SC maintains a running log which captures all console output. This log is maintained as a first-in, first-out buffer: New console output may displace old console output if the buffer is full. By default, only the last eight kilobytes of the console log file are displayed. The optional <code>-a</code> argument specifies that the entire console log file be displayed. It is possible for the SC to fill this log more quickly than the <code>consolehistory</code> subcommand can read it. This means that it is possible for some log data to be lost before it can be displayed. If this happens, the <code>consolehistory</code> subcommand displays “scadm: lost <number> bytes of console log data” in the log output, to indicate that some data was lost. The format for the <code>consolehistory</code> subcommand is: <pre>scadm consolehistory [-a]</pre> The <code>consolehistory</code> subcommand is not available on all platforms. If this command is used on a platform that does not support it, <code>scadm</code> prints this message: <pre>scadm: command/option not supported</pre> and exit with non-zero status. date Display the SC's time and date The format for the <code>date</code> subcommand is:

scadm date

download

Program the SC's firmware.

There are two parts to the firmware, the boot monitor and the main image.

By default, The `scadm` command's download programs the main firmware image. The boot argument selects programming of the boot monitor.

The format for the download subcommand is:

```
scadm download [boot] file
```

fruhistory

Display the contents of the “field replacable unit” log maintained by the SC. By default, only the last eight kilobytes of the fru history log file are displayed. The data in contained this log contains snapshots of the SC's “showfru” command, taken whenever the system is reset, or a hot-plug event is detected by the SC.

The optional `-a` argument specifies that the entire fru log file be displayed.

It is possible for the SC to fill this log more quickly than the `fruhistory` subcommand can read it. This means that it is possible for some log data to be lost before it can be displayed. If this happens, the `fruhistory` subcommand displays “scadm: lost <number> bytes of fru log data” in the log output, to indicate that some data was lost.

The format for the `fruhistory` subcommand is:

```
scadm fruhistory [-a]
```

The `fruhistory` subcommand is not available on all platforms. If this command is used on a platform which does not support it, `scadm` prints this message:

```
scadm: command/option not supported
```

and exit with non-zero status.

help

Display a list of commands.

The format for the `help` subcommand is:

```
scadm help
```

loghistory

Display the most recent entries in the SC event log. The optional `-a` argument causes the entire event log history to be displayed. The `-a` argument is available only on platforms which support large log files. On platforms which do not support large log files, this flag has no additional effect.

It is possible for the SC to fill this log more quickly than the `loghistory` subcommand can read it. This means that it is possible for some log data to be lost before it can be displayed. If this happens, the `loghistory` subcommand displays “scadm: lost <number> events” in the log output, to indicate that some data was lost.

The format for the `loghistory` subcommand is:

```
scadm loghistory [-a]
```

resetrsc

Reset the SC. There are two types of resets allowed, a hard reset and a soft reset. The hard reset is done by default. The soft reset can be selected by using the `-s` option.

The format for the `resetrsc` subcommand is:

```
scadm resetrsc [-s]
```

send_event

Manually send a text based event. The SC can forward the event to the SC event log. You can configure the `-c` option to send a critical warning to email, alert to logged in SC users, and `syslog`. Critical events are logged to `syslog(3C)`. There is an 80 character limit to the length of the associated text message.

The format for the `send_event` subcommand is:

```
scadm send_event [-c] "message"
```

set

Set SC configuration variables to a value.

Examples of SC configuration variables include: SC IP address `netsc_ipaddr` and SC Customer Information `sc_customerinfo`. See the output from the `scadm help` command for a complete list of SC configuration variables.

The format for the `set` subcommand is:

```
scadm set variable value
```

show

Display the current SC configuration variable settings. If no variable is specified, `scadm` shows all variable settings.

The format for the `show` subcommand is:

```
scadm show [variable]
```

shownetwork

Display the current network configuration parameters for SC.

The format for the `shownetwork` subcommand is:

```
scadm shownetwork
```

useradd

Add user accounts to the SC. The SC supports up to sixteen separate users.

The format for the `useradd` subcommand is:

```
scadm useradd username
```

userdel

Delete a user account from SC.

The format for the `userdel` subcommand is:

```
scadm userdel username
```

userpassword

Set a password for the user account specified. This password overrides any existing password currently set. There is no verification of the old password before setting the new password.

The format for the `userpassword` subcommand is:

```
scadm userpassword username
```

userperm

Set the permission level for the user.

The format for the `userperm` subcommand is:

```
scadm userperm username [aucr]
```

usershow

Display details on the specified user account. If a username is not specified, all user accounts are displayed.

The format for the `usershow` subcommand is:

```
scadm usershow username
```

version

Display the version numbers of the SC and its components.

The format for the `version` subcommand is:

```
scadm version [-v]
```

选项

The `resetrsc`, `send_event`, and `version` subcommands have associated options. Options follow subcommands on the command line and are separated from the subcommand by a SPACE.

The `resetrsc` subcommand supports the following options:

-s

Perform a soft reset instead of a hard reset. A hard reset physically resets the SC hardware. The SC software jumps to the boot firmware, simulating a reset, for a soft reset.

The `send_event` subcommand supports the following options:

-c

Send a critical event. Without the -c, `send_event` sends a warning.

The `version` subcommand supports the following options:

-v

Display a verbose output of version numbers and associated information.

The `consolehistory`, `fruhistory`, and `loghistory` subcommands support the following option:

-a

Display the entire log. These subcommands normally display only the most recent log data. This flag causes them to display the entire log.

操作数

The `download`, `send_event`, `set`, `show`, `useradd`, `userdel`, `userperm`, `usershow`, `userpassword`, and `userperm` subcommands have associated arguments (operands).

If the subcommand has an option, the arguments follow the option on the command line and is separated from the option by a SPACE. If the subcommand does not have an option, the arguments follow the subcommand on the command line and are separated from the subcommand by a SPACE. If there are more than one arguments, they are separated from each other by a SPACE.

The `download` subcommand supports the following arguments:

boot

Program the boot monitor portion of the flash. The main portion of the flash is programmed without any arguments

file

Specify *file* as the path to where the boot or main firmware image resides for download.

Examples of *file* are:

`/usr/platform/platform_type/lib/image/alommainfw`

or

`/usr/platform/platform_type/lib/image/alombootfw`

The `send_event` subcommand supports the following arguments:

“message”

Describe event using the text contained in *message*. Enclose *message* in quotation marks.

The `set` subcommand supports the following arguments:

variable

Set SC configuration *variable*.

value

Set SC configuration variable to *value*.

The show subcommand supports the following arguments:

variable

Display the value of that particular variable.

The useradd subcommand supports the following arguments:

username

Add new SC account *username*.

The userdel subcommand supports the following arguments:

username

Remove SC account *username*.

The userperm subcommand supports the following arguments:

-aucr

Set permissions for SC user accounts. If no permissions are specified, all four permissions are disabled and read only access is assigned.

The following are the definitions for permissions:

a

Allow user to administer or change the SC configuration variables

u

Allow user to use the user commands to modify SC accounts

c

Allow user to connect to console.

r

Allow user to reset SC and to power on and off the host.

username

Change permissions on SC account *username*.

The -usershow subcommand supports the following arguments:

username

Display information on SC account *username*. If *username* is not specified, all accounts are displayed.

The userpassword subcommand supports the following arguments:

username

Set SC password for *username*.

The userperm subcommand supports the following arguments:

username

Change SC permissions for *username*.

示例

示例 1 Displaying the SC's Date and Time

The following command displays the SC's date and time.

```
scadm date
```

示例 2 Setting the SC's Configuration Variables

The following command sets the SC's configuration variable `netsc_ipaddr` to `192.168.1.2`:

```
scadm set netsc_ipaddr 192.168.1.2
```

示例 3 Displaying the Current SC's Configuration Settings:

The following command displays the current SC configuration settings:

```
scadm show
```

示例 4 Displaying the Current Settings for a Variable

The following command displays the current settings for the `sys_hostname` variable:

```
scadm show sys_hostname
```

示例 5 Sending a Text-Based Critical Event

The following command sends a critical event to the SC logs, alerts the current SC users, and sends an event to [syslog\(3C\)](#):

```
scadm send_event -c "The UPS signaled a loss in power"
```

示例 6 Sending an Informational Text-Based Event

The following command sends a non-critical informational text based event to the SC event log:

```
scadm send_event "The disk is close to full capacity"
```

示例 7 Adding a User To the SC

The following command adds user `rsccroot` to the SC:

```
scadm useradd rsccroot
```

示例 8 Deleting a User From the SC

The following command deletes user `olduser` from the SC:

```
scadm userdel olduser
```

示例 9 Displaying User Details

The following command displays details of all user accounts:

```
scadm usershow
```

示例 10 Displaying Details for a Specific User

The following command displays details of user account `rscroot`:

```
scadm usershow rscroot
```

示例 11 Setting the User Permission Level

The following command sets the full permission level for user `rscroot` to `aucr`:

```
scadm userperm rscroot aucr
```

示例 12 Setting the User Permission Level

The following command sets only console access for user `newuser` to `c`:

```
scadm userperm newuser c
```

示例 13 Setting the User Permission Level

The following command sets the permission level for user `newuser` to read only access:

```
scadm userperm newuser
```

示例 14 Displaying the Current Network Parameters

The following command displays the current network configuration parameters for the SC:

```
scadm shownetwork
```

示例 15 Viewing the Consolehistory

The following command displays the content console in the SC event log:

```
scadm consolehistory [-a]
```

示例 16 Viewing the Fruhistory

The following command displays the content of the “field replacable unit” in the SC event log:

```
scadm fruhistory [-a]
```

示例 17 Viewing the Loghistory

The following command displays the most recent entries in the SC event log:

```
scadm loghistory [-a]
```


示例 18 Displaying Verbose Information

The following command displays verbose version information on the SC and its components:

```
scadm version -v
```

退出状态

The following exit values are returned:

- 0
Successful completion.
- non-zero
An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/library/platform
Interface Stability	Committed

另请参见

[uname\(1\)](#), [syslog\(3C\)](#), [attributes\(5\)](#)

引用名 sckmd – Sun cryptographic key management daemon

用法概要 /usr/platform/sun4u/lib/sckmd

描述 sckmd is a server process that resides on a high-end system domain to maintain the Internet Protocol Security (IPsec) Security Associations (SAs) needed to secure communications between a Service Processor or System Controller (SC) and platform management software running within a domain. The [cvcd\(1M\)](#) and [dcs\(1M\)](#) daemons use these Security Associations. See [ipsec\(7P\)](#) for a description of Security Associations.

The sckmd daemon receives SAs from the Service Processor or SC and installs these SAs in a domain's Security Association Database (SADB) using [pf_key\(7P\)](#).

sckmd starts up at system boot time as an SMF service. The FMRI for the sckmd service is:

```
svc:/platform/sun4u/sckmd:default
```

A domain supports only one running sckmd process at a time.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWscmx.u, service/key-management/sun-fire-15000
Interface Stability	Committed

另请参见 [cvcd\(1M\)](#), [dcs\(1M\)](#), [ipsecconf\(1M\)](#), [ipsecalgs\(1M\)](#), [attributes\(5\)](#), [ipsec\(7P\)](#), [ipsecah\(7P\)](#), [ipsecesp\(7P\)](#), [pf_key\(7P\)](#)

附注 The sckmd service is used only on Sun Fire high-end systems and the SPARC Enterprise Server family. It provides a mechanism for exchanging IPsec keys between a domain and its System Controller (SC) or Service Processor. These platforms use IPsec to secure the communications between the SC or Service Processor and certain platform-specific daemons in the domain. Such daemons currently include [cvcd\(1M\)](#) and [dcs\(1M\)](#).

The documentation for each platform that supports sckmd describes how to configure its use of IPsec for such communications. Also, the documentation for each specific application describes how to configure its security policies and IPsec options in a manner appropriate for the target platform. Refer to the platform- and application-specific documentation for detailed information.

引用名	scmadm – 存储高速缓存管理器管理实用程序
用法概要	<pre>scmadm scmadm -h scmadm -e scmadm -d scmadm -v scmadm -C [parameter [= [value]]...] scmadm -o {system cd device} [option] scmadm -m {cd diskname all}</pre>
描述	scmadm 命令提供了用于控制和收集存储设备高速缓存相关信息的各种选项。
选项	如果未指定任何选项，scmadm 将显示所配置的包含磁盘名称、选项和全局选项的高速缓存描述符的列表。scmadm 命令支持以下选项： -h 显示 scmadm 命令的用法信息。 -e 读取配置并使用那些参数启用存储设备高速缓存。请参见 dscfg(1M)。 -d 关闭存储设备高速缓存。 -v 显示高速缓存的版本号。 -C [parameter[=[value]] ...] 设置或显示配置参数。如果没有为 -c 选项指定任何参数，则会显示当前的高速缓存配置参数。如果指定了 <i>parameter</i>，则会显示 <i>parameter</i> 的当前值。如果指定了 <i>parameter=value</i>，则会显示 <i>parameter</i> 的当前值并将该参数的值更改为 <i>value</i>。如果省略了 <i>value</i> 或者将 <i>value</i> 指定为空字符串、" " 或 "-"，则会将参数从配置中删除，且系统将使用缺省值。可以在 scmadm 命令的单个调用中指定多个参数。对配置参数所做的更改在高速缓存下次重新启动时才会生效。 -o {system cd diskname} [option] 为系统或者为由 cd 或 <i>diskname</i> 指定的高速缓存设备设置或显示选项。如果指定了 选项 <i>rdcache</i> 或 <i>nordcache</i>，则将为系统或指定的高速缓存设备设置该选项。选项保存为配置的一部分，因此选项持久有效。请参见 dscfg(1M)。要通知系统“忘记”所保存的某个选项，请使用 forget（忘记）选项。这不会更改选项，只会把选项从所保存的配置中删除。如果未指定任何选项，则会显示当前选项。<i>rdcache</i> 选项是缺省设置。<i>option</i> 的定义如下： <i>rdcache</i> 数据块可能会被再次引用，应该保留在高速缓存中。

nordcache
数据块不可能被再次引用，应该将其作为最近最少使用的项对待，以便其他块可以更长久地保留在高速缓存中。

-m { cd | *diskname* | all }
为由 *cd* 或 *diskname* 指定的设备显示高速缓存描述符和磁盘名称映射；如果指定了 *all*（全部），则会为系统上的所有存储设备显示高速缓存描述符和磁盘名称映射。

退出状态
scmadm 命令成功时返回 0，发生错误时返回非零值。

文件
/dev/sdbc

属性
有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	storage/avs/avs-cache-management
接口稳定性	Committed（已确定）

另请参见
[dscfg\(1M\)](#)、[attributes\(5\)](#)

诊断
如果没有足够的连续内存，则 *scmadm* 将失败。

引用名	sconadm – register system information																				
用法概要	<pre> /usr/sbin/sconadm register -a [-e softwareUpdate -E softwareUpdate] [-h <i>hostname</i>] [-l <i>logfile</i>] [-N] [-p <i>proxy_host[:proxy_port]</i>] [-r <i>registration_profile</i>] [-u <i>username</i>] [-x <i>proxy_username</i>] /usr/sbin/sconadm proxy [-l <i>logfile</i>] [-p <i>proxy_host[:proxy_port]</i>] [-r <i>registration_profile</i>] [-x <i>proxy_username</i>] </pre>																				
描述	The <code>sconadm</code> utility is a command-line version of the Basic Registration GUI. In the first form of the command in the SYNOPSIS, <code>sconadm</code> uses the <code>register</code> subcommand to register a host with a registration server. In the second form, <code>sconadm</code> uses the <code>proxy</code> subcommand to configure all of the components for software update to use an HTTP web proxy. The parameters specified with <code>-u</code>, <code>-e</code> (or <code>-E</code>), <code>-h</code>, <code>-p</code>, and <code>-x</code> override values specified in your registration profile. A template for this profile, owned by root, with read-only permissions, is stored in <code>/usr/lib/breg/data/RegistrationProfile.properties</code>. See registration_profile(4). For the <code>proxy</code> subcommand, the proxy password is stored in the <code>RegistrationProfile.properties</code> file, available if proxy authentication is needed. Storage in the profile prevents proxy passwords from being exposed as part of a listing of processes on a system.																				
选项	The following options are supported: <table> <tr> <td><code>-a</code></td><td>Accept “Terms of Use and Binary Code License”. Absence of this option means that you do not accept the license.</td></tr> <tr> <td><code>-e softwareUpdate</code></td><td>Enable client to be managed at the Sun-hosted Update Connection Service.</td></tr> <tr> <td><code>-E softwareUpdate</code></td><td>Disable client's ability to be managed at the Sun-hosted Update Connection Service.</td></tr> <tr> <td><code>-h <i>hostname</i></code></td><td>Hostname of the machine you want to register.</td></tr> <tr> <td><code>-l <i>logfile</i></code></td><td>Pathname of log file.</td></tr> <tr> <td><code>-N</code></td><td>Never register.</td></tr> <tr> <td><code>-p <i>proxy_host[:proxy_port]</i></code></td><td>Proxy host name and optional proxy port number.</td></tr> <tr> <td><code>-r <i>registration_profile</i></code></td><td>Pathname to a registration profile.</td></tr> <tr> <td><code>-u <i>username</i></code></td><td>User name (a Sun Online Account).</td></tr> <tr> <td><code>-x <i>proxy_username</i></code></td><td>User name on the proxy host.</td></tr> </table>	<code>-a</code>	Accept “Terms of Use and Binary Code License”. Absence of this option means that you do not accept the license.	<code>-e softwareUpdate</code>	Enable client to be managed at the Sun-hosted Update Connection Service.	<code>-E softwareUpdate</code>	Disable client's ability to be managed at the Sun-hosted Update Connection Service.	<code>-h <i>hostname</i></code>	Hostname of the machine you want to register.	<code>-l <i>logfile</i></code>	Pathname of log file.	<code>-N</code>	Never register.	<code>-p <i>proxy_host[:proxy_port]</i></code>	Proxy host name and optional proxy port number.	<code>-r <i>registration_profile</i></code>	Pathname to a registration profile.	<code>-u <i>username</i></code>	User name (a Sun Online Account).	<code>-x <i>proxy_username</i></code>	User name on the proxy host.
<code>-a</code>	Accept “Terms of Use and Binary Code License”. Absence of this option means that you do not accept the license.																				
<code>-e softwareUpdate</code>	Enable client to be managed at the Sun-hosted Update Connection Service.																				
<code>-E softwareUpdate</code>	Disable client's ability to be managed at the Sun-hosted Update Connection Service.																				
<code>-h <i>hostname</i></code>	Hostname of the machine you want to register.																				
<code>-l <i>logfile</i></code>	Pathname of log file.																				
<code>-N</code>	Never register.																				
<code>-p <i>proxy_host[:proxy_port]</i></code>	Proxy host name and optional proxy port number.																				
<code>-r <i>registration_profile</i></code>	Pathname to a registration profile.																				
<code>-u <i>username</i></code>	User name (a Sun Online Account).																				
<code>-x <i>proxy_username</i></code>	User name on the proxy host.																				

示例 Unless specified otherwise, the commands below require root privileges or privileges equivalent to root. See [privileges\(5\)](#).

示例 1 Registering a New System

Assume a file `registrationprofile.properties` in `/tmp` that contains the following:

```
userName=user123
password=abc123
hostName=
subscriptionKey=
portalEnabled=false
proxyHostName=
proxyPort=
proxyUserName=
proxyPassword=
```

To register a new system using the profile above, you enter:

```
/usr/sbin/sconadm register -a -r /tmp/registrationprofile.properties
```

示例 2 Reregistering a System with a Different User

Assume a file `registrationprofile.properties` in `/tmp` with the contents shown below. Note the changed specification for `userName` and `password`.

```
userName=newuser
password=newpassword
hostName=
subscriptionKey=
portalEnabled=false
proxyHostName=
proxyPort=
proxyUserName=
proxyPassword=
```

To reregister a new system using the profile above, you enter the same command you entered to register the system:

```
/usr/sbin/sconadm register -a -r /tmp/registrationprofile.properties
```

示例 3 Reregistering a System, Adding a Sun Subscription Key

Modify `registrationprofile.properties` as follows:

```
userName=newuser
password=newpassword
hostName=
subscriptionKey=abc12345678
portalEnabled=false
proxyHostName=
```

示例 3 Reregistering a System, Adding a Sun Subscription Key (续)

```
proxyPort=
proxyUserName=
proxyPassword=
```

Run the command:

```
/usr/sbin/sconadm register -a -r /tmp/registrationprofile.properties
```

示例 4 Reregistering and Enabling Access to all Update Connection Services

Modify `registrationprofile.properties` as follows:

```
userName=newuser
password=newpassword
hostName=
subscriptionKey=abc12345678
portalEnabled=false
proxyHostName=
proxyPort=
proxyUserName=
proxyPassword=
```

Note that `portalEnabled` is set to `false`. Run the command:

```
/usr/sbin/sconadm register -a -r /tmp/registrationprofile.properties \
-e softwareUpdate
```

示例 5 Never Registering

To never register a system, enter:

```
/usr/sbin/sconadm register -N
```

示例 6 Using a Proxy Server With Proxy Authentication

Edit `registrationprofile.properties` as follows:

```
userName=
password=
hostName=
subscriptionKey=
portalEnabled=
proxyHostName=webcache.mycompany.com
proxyPort=8080
proxyUserName=myCompanyProxyUserName
proxyPassword=myCompanyProxyPassword
```

Run the command:

示例 6 Using a Proxy Server With Proxy Authentication (续)

```
/usr/sbin/sconadm proxy -r /tmp/registrationprofile.properties
```

示例 7 Changing Proxy Host Settings

Edit `registrationprofile.properties` as follows:

```
userName=  
password=  
hostName=  
subscriptionKey=  
portalEnabled=  
proxyHostName=webcache.mycompany.com  
proxyPort=8080  
proxyUserName=myCompanyProxyUserName  
proxyPassword=myCompanyProxyPassword
```

Run the command:

```
/usr/sbin/sconadm proxy -r /tmp/registrationprofile.properties
```

Then, change the `proxyHostName` value by running the following command:

```
/usr/sbin/sconadm proxy -r /tmp/registrationprofile.properties \  
-p newproxy.mycompany.com
```

After the preceding command all proxies use `newproxy.mycompany.com`.

示例 8 Resetting a System Not to Use a Proxy

Edit `registrationprofile.properties` as follows:

```
userName=  
password=  
hostName=  
subscriptionKey=  
portalEnabled=  
proxyHostName=  
proxyPort=  
proxyUserName=  
proxyPassword=
```

Note that values for all proxy fields are null.

Run the command:

```
/usr/sbin/sconadm proxy -r /tmp/registrationprofile.properties
```


- 退出状态 0 Success.
- >0 An error occurred.
- 文件 /usr/lib/breg/data/RegistrationProfile.properties
- Registration profile template.
- 属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWbrg
Interface Stability	Committed

另请参见 [registration_profile\(4\)](#), [attributes\(5\)](#), [privileges\(5\)](#)

引用名

sdpadm – SDP system configuration administration

用法概要

```
/usr/sbin/sdpadm status | enable | disable
```

描述

The `sdpadm` command is no longer shipped with the Oracle Solaris operating system. SDP sockets are enabled by default. This man page describes a utility that was shipped in prior builds of the operating system.

The `sdpadm` command is used to display the system state of the Sockets Direct Protocol (SDP) protocol. The `sdpadm` command can optionally be used to set the state of the SDP protocol. See [sdp\(7D\)](#).

By default the SDP protocol is disabled on the system. It can be enabled by using `sdpadm enable`.

选项

The following subcommands are supported:

<code>status</code>	Displays the system status of the SDP protocol
<code>enable</code>	Enables the SDP protocol
<code>disable</code>	Disables the SDP protocol

用法

The required privileges to change the state of the SDP protocol are controlled by the network configuration policy. If a user does not have the correct privileges to set the SDP policy, `sdpadm` returns the current state of SDP without having changed the state.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/spdadm
Interface Stability	Obsolete

另请参见

[attributes\(5\)](#), [sdp\(7D\)](#)

Infiniband Specification Volume 1 (<http://www.infinibandta.org>)

引用名

sendmail – send mail over the internet

用法概要

```
/usr/lib/sendmail [-Ac] [-Am] [-ba] [-bD] [-bd] [-bi] [-bl]
                  [-bm] [-bp] [-bP] [-bs] [-bt] [-bv] [-B type] [-C file]
                  [-D logfile] [-d X] [-F fullname] [-f name] [-G]
                  [-h N] [-L tag] [-M xvalue] [-N notifications] [-n]
                  [-Ooption =value] [-o xvalue] [-p protocol]
                  [-Q [reason]] [-q [time]] [-q Xstring] [-R ret]
                  [-r name] [-t] [-V envid] [-v] [-X logfile]
                  [address]...
```

描述

The sendmail utility sends a message to one or more people, routing the message over whatever networks are necessary. sendmail does internetwork forwarding as necessary to deliver the message to the correct place.

sendmail is not intended as a user interface routine. Other programs provide user-friendly front ends. sendmail is used only to deliver pre-formatted messages.

With no flags, sendmail reads its standard input up to an EOF, or a line with a single dot, and sends a copy of the letter found there to all of the addresses listed. It determines the network to use based on the syntax and contents of the addresses.

Local addresses are looked up in the local [aliases\(4\)](#) file, or in a name service as defined by the [nsswitch.conf\(4\)](#) file, and aliased appropriately. In addition, if there is a .forward file in a recipient's home directory, sendmail forwards a copy of each message to the list of recipients that file contains. Refer to the NOTES section for more information about .forward files.

Aliasing can be prevented by preceding the address with a backslash.

There are several conditions under which the expected behavior is for the alias database to be either built or rebuilt. This cannot occur under any circumstances unless root owns *and* has exclusive write permission to the /etc/mail/aliases* files.

If a message is found to be undeliverable, it is returned to the sender with diagnostics that indicate the location and nature of the failure; or, the message is placed in a dead.letter file in the sender's home directory.

Service Management

The sendmail service is managed by the service management facility, [smf\(5\)](#), under the service identifiers:

```
svc:/network/smtp:sendmail
svc:/network/sendmail-client:default
```

Administrative actions on these services, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The services' status can be queried using the [svcs\(1\)](#) command.

These are separate services rather than instances of the same service so that other services can properly express any dependencies. In particular, here are some guidelines about which service/instance should be depended on for which purposes:

- For a service that uses `sendmail` to send mail, an optional dependency on the service `svc:/network/sendmail-client` might be in order.
- For a service that needs to receive mail in general, but does not depend on `sendmail` being the particular SMTP receiver, a dependency on the service `svc:/network/smtp` might be in order.
- For a service that needs to interact with `sendmail` in particular, such as a `Milter`, a dependency on the instance `svc:/network/smtp:sendmail` might be in order.

For the last two, note the difference, as the latter has the `":sendmail"` instance specification, whereas the former does not, thus representing the more general service.

Enabling Access to Remote Clients

On an unmodified system, access to `sendmail` by remote clients is enabled and disabled through the service management facility (see [smf\(5\)](#)). In particular, remote access is determined by the value of the `local_only` SMF property:

```
svc:/network/smtp:sendmail/config/local_only = true
```

A setting of `true`, as above, disallows remote access; `false` allows remote access. The default value is `true`.

The following example shows the sequence of SMF commands used to enable `sendmail` to allow access to remote systems:

```
# svccfg -s svc:/network/smtp:sendmail setprop config/local_only = false
# svcadm refresh svc:/network/smtp:sendmail
# svcadm restart svc:/network/smtp:sendmail
```

See [svcadm\(1M\)](#) and [svccfg\(1M\)](#).

Note, however, on a system where any of the [sendmail\(4\)](#) files have been customized, setting this property might not have the intended effect. See [sendmail\(4\)](#) for details.

Automated Rebuilding of Configuration Files

See [sendmail\(4\)](#) for details on which service properties can be set to automate (re)building of configuration files when the service is started.

Restricting Host Access

`sendmail` uses TCP Wrappers to restrict access to hosts. It uses the service name of `sendmail` for `hosts_access()`. For more information on TCP Wrappers, see [tcpd\(1M\)](#) and [hosts_access\(4\)](#) in the `security/tcp-wrapper` package. [tcpd\(1M\)](#) and [hosts_access\(4\)](#) are not part of the Solaris man pages.

Startup Options

The `/etc/default/sendmail` file stores startup options for `sendmail` so that the options are not removed when a host is upgraded. See also [sendmail\(4\)](#) for details on which service properties can be set to automate (re)building of configuration files when the service is started.

You can use the following variables in the `/etc/default/sendmail` startup file:

CLIENTOPTIONS=string

Selects additional options to be used with the client daemon, which looks in the client-only queue (`/var/spool/clientmqueue`) and acts as a client queue runner. No syntax checking is done, so be careful when making changes to this variable.

CLIENTQUEUEINTERVAL=#

Similar to the `QUEUEINTERVAL` option, `CLIENTQUEUEINTERVAL` sets the time interval for mail queue runs. However, the `CLIENTQUEUEINTERVAL` option controls the functions of the client daemon, instead of the functions of the master daemon. Typically, the master daemon is able to deliver all messages to the SMTP port. However, if the message load is too high or the master daemon is not running, then messages go into the client-only queue, `/var/spool/clientmqueue`. The client daemon, which checks in the client-only queue, then acts as a client queue processor.

ETRNL_HOSTS=string

Enables an SMTP client and server to interact immediately without waiting for the queue run intervals, which are periodic. The server can immediately deliver the portion of its queue that goes to the specified hosts. For more information, refer to the [etrn\(1M\)](#) man page.

MODE=-bd

Selects the mode to start `sendmail` with. Use the `-bd` option or leave it undefined.

OPTIONS=string

Selects additional options to be used with the master daemon. No syntax checking is done, so be careful when making changes to this variable.

QUEUEINTERVAL=#

Sets the interval for mail queue runs on the master daemon. `#` can be a positive integer that is followed by either `s` for seconds, `m` for minutes, `h` for hours, `d` for days, or `w` for weeks. The syntax is checked before `sendmail` is started. If the interval is negative or if the entry does not end with an appropriate letter, the interval is ignored and `sendmail` starts with a queue interval of 15 minutes.

QUEUEOPTIONS=p

Enables one persistent queue runner that sleeps between queue run intervals, instead of a new queue runner for each queue run interval. You can set this option to `p`, which is the only setting available. Otherwise, this option is not set.

Mail Filter API

`sendmail` supports a mail filter API called “milter”. For more information, see `/usr/include/libmilter/README` and <http://www.milter.org>

选项

The following options are supported:

-Ac

Uses `submit.cf` even if the operation mode does not indicate an initial mail submission.

-Am

Uses `sendmail.cf` even if the operation mode indicates an initial mail submission.

- ba
Goes into ARPANET mode. All input lines must end with a RETURN-LINEFEED, and all messages are generated with a RETURN-LINEFEED at the end. Also, the From: and Sender: fields are examined for the name of the sender.
- bd
Runs as a daemon in the background, waiting for incoming SMTP connections.
- bD
Runs as a daemon in the foreground, waiting for incoming SMTP connections.
- bi
Initializes the [aliases\(4\)](#) database. Root must own *and* have exclusive write permission to the /etc/mail/aliases* files for successful use of this option.
- bl
Runs as a daemon (like -bd) but accepts only loopback SMTP connections.
- bm
Delivers mail in the usual way (default).
- bp
Prints a summary of the mail queues.
- bP
Prints the number of entries in the queues. This option is only available with shared memory support.
- bs
Uses the SMTP protocol as described in RFC 2821. This flag implies all the operations of the -ba flag that are compatible with SMTP.
- bt
Runs in address test mode. This mode reads addresses and shows the steps in parsing; it is used for debugging configuration tables.
- bv
Verifies names only. Does not try to collect or deliver a message. Verify mode is normally used for validating users or mailing lists.
- B *type*
Indicates body *type* (7BIT or 8BITMIME).
- C *file*
Uses alternate configuration file.
- D *logfile*
Send debugging output to the indicated log file instead of stdout.
- d *X*
Sets debugging value to *X*.

- f *name*
Sets the name of the “from” person (that is, the sender of the mail).
- F *fullname*
Sets the full name of the sender.
- G
When accepting messages by way of the command line, indicates that they are for relay (gateway) submission. When this flag is set, `sendmail` might complain about syntactically invalid messages, for example, unqualified host names, rather than fixing them. `sendmail` does not do any canonicalization in this mode.
- h *N*
Sets the hop count to *N*. The hop count is incremented every time the mail is processed. When it reaches a limit, the mail is returned with an error message, the victim of an aliasing loop.
- L *tag*
Sets the identifier used in `syslog` messages to the supplied *tag*.
- M*xvalue*
Sets macro *x* to the specified *value*.
- n
Does not do aliasing.
- N *notifications*
Tags all addresses being sent as wanting the indicated *notifications*, which consists of the word “NEVER” or a comma-separated list of “SUCCESS”, “FAILURE”, and “DELAY” for successful delivery, failure and a message that is stuck in a queue somewhere. The default is “FAILURE,DELAY”.
- o*xvalue*
Sets option *x* to the specified *value*. Processing Options are described below.
- O*option=value*
Sets *option* to the specified *value* (for long from names). Processing Options are described below.
- p *protocol*
Sets the sending protocol. The *protocol* field can be in form *protocol: host* to set both the sending protocol and the sending host. For example: `-pUUCP: uunet` sets the sending *protocol* to UUCP and the sending host to uunet. Some existing programs use `-oM` to set the *r* and *s* macros; this is equivalent to using `-p`.
- q[*time*]
Processes saved messages in the queue at given intervals. If *time* is omitted, processes the queue once. *time* is given as a tagged number, where *s* is seconds, *m* is minutes, *h* is hours, *d* is days, and *w* is weeks. For example, `-q1h30m` or `-q90m` would both set the timeout to one hour thirty minutes.

By default, sendmail runs in the background. This option can be used safely with `-bd`.

`-qp [time-]`

Similar to `-q[time]`, except that instead of periodically forking a child to process the queue, sendmail forks a single persistent child for each queue that alternates between processing the queue and sleeping. The sleep time (*time*) is specified as the argument; it defaults to 1 second. The process always sleeps at least 5 seconds if the queue was empty in the previous queue run.

`-qf`

Processes saved messages in the queue once and does not `fork(2)`, but runs in the foreground.

`-qG name`

Processes jobs in queue group called *name* only.

`-q[!]I substr`

Limits processed jobs to those containing *substr* as a substring of the queue ID or not when `!` is specified.

`-q[!]Q substr`

Limits processed jobs to those quarantined jobs containing *substr* as a substring of the quarantine *reason* or not when `!` is specified.

`-q[!]R substr`

Limits processed jobs to those containing *substr* as a substring of one of the recipients or not when `!` is specified.

`-q[!]S substr`

Limits processed jobs to those containing *substr* as a substring of the sender or not when `!` is specified.

`-Q[reason]`

Quarantines a normal queue item with the given reason or unquarantines a quarantined queue item if no reason is given. This should only be used with some sort of item matching as described above.

`-r name`

An alternate and obsolete form of the `-f` flag.

`-R ret`

Identifies the information you want returned if the message bounces. *ret* can be HDRS for headers only or FULL for headers plus body.

`-t`

Reads message for recipients. `To:`, `Cc:`, and `Bcc:` lines are scanned for people to send to. The `Bcc:` line is deleted before transmission. Any addresses in the argument list is suppressed. The NoRecipientAction Processing Option can be used to change the behavior when no legal recipients are included in the message.

-v

Goes into verbose mode. Alias expansions are announced, and so forth.

-V *envid*

The indicated *envid* is passed with the envelope of the message and returned if the message bounces.

-X *logfile*

Logs all traffic in and out of sendmail in the indicated *logfile* for debugging mailer problems. This produces a lot of data very quickly and should be used sparingly.

Processing Options

There are a number of “random” options that can be set from a configuration file. Options are represented by a single character or by multiple character names. The syntax for the single character names of is:

Oxvalue

This sets option *x* to be *value*. Depending on the option, *value* may be a string, an integer, a boolean (with legal values t, T, f, or F; the default is TRUE), or a time interval.

The multiple character or long names use this syntax:

O Longname=argument

This sets the option *Longname* to be *argument*. The long names are beneficial because they are easier to interpret than the single character names.

Not all processing options have single character names associated with them. In the list below, the multiple character name is presented first followed by the single character syntax enclosed in parentheses.

AliasFile (*Afile*)

Specifies possible alias files.

AliasWait (*a N*)

If set, waits up to *N* minutes for an “@:” entry to exist in the [aliases\(4\)](#) database before starting up. If it does not appear in *N* minutes, issues a warning. Defaults to 10 minutes.

AllowBogusHELO

Allows a HELO SMTP command that does not include a host name. By default this option is disabled.

BadRcptThrottle=*N*

If set and more than the specified number of recipients in a single SMTP envelope are rejected, sleeps for one second after each rejected RCPT command.

BlankSub (*Bc*)

Sets the blank substitution character to *c*. Unquoted spaces in addresses are replaced by this character. Defaults to SPACE (that is, no change is made).

CACertFile

File containing one CA cert.

CACertPath

Path to directory with certs of CAs.

CheckAliases (n)

Validates the RHS of aliases when rebuilding the [aliases\(4\)](#) database.

CheckpointInterval (CN)

Checkpoints the queue every *N* (default 10) addresses sent. If your system crashes during delivery to a large list, this prevents retransmission to any but the last *N* recipients.

ClassFactor (zfact)

The indicated factor *fact* is multiplied by the message class (determined by the Precedence: field in the user header and the P lines in the configuration file) and subtracted from the priority. Thus, messages with a higher Priority: are favored. Defaults to 1800.

ClientCertFile

File containing the cert of the client, that is, this cert is used when sendmail acts as client.

ClientKeyFile

File containing the private key belonging to the client cert.

ClientPortOptions

Sets client SMTP options. The options are key=value pairs. Known keys are:

Addr Address Mask

Address Mask defaults to INADDR_ANY. The address mask can be a numeric address in dot notation or a network name.

Family

Address family (defaults to INET).

Listen

Size of listen queue (defaults to 10).

Port

Name/number of listening port (defaults to smtp).

RcvBufSize

The size of the TCP/IP receive buffer.

SndBufSize

The size of the TCP/IP send buffer.

Modifier

Options (flags) for the daemon. Can be:

h

Uses name of interface for HELO command.

If `h` is set, the name corresponding to the outgoing interface address (whether chosen by means of the `Connection` parameter or the default) is used for the HELO/EHLO command.

`ColonOkInAddr`

If set, colons are treated as a regular character in addresses. If not set, they are treated as the introducer to the RFC 822 “group” syntax. This option is on for version 5 and lower configuration files.

`ConnectionCacheSize (kN)`

The maximum number of open connections that are to be cached at a time. The default is 1. This delays closing the current connection until either this invocation of `sendmail` needs to connect to another host or it terminates. Setting it to 0 defaults to the old behavior, that is, connections are closed immediately.

`ConnectionCacheTimeout (Ktimeout)`

The maximum amount of time a cached connection is permitted to idle without activity. If this time is exceeded, the connection is immediately closed. This value should be small (on the order of ten minutes). Before `sendmail` uses a cached connection, it always sends a NOOP (no operation) command to check the connection. If the NOOP command fails, it reopens the connection. This keeps your end from failing if the other end times out. The point of this option is to be a good network neighbor and avoid using up excessive resources on the other end. The default is five minutes.

`ConnectionRateThrottle`

The maximum number of connections permitted per second. After this many connections are accepted, further connections are delayed. If not set or ≤ 0 , there is no limit.

`ConnectionRateWindowSize`

Define the length of the interval for which the number of incoming connections is maintained. The default is 60 seconds.

`ControlSocketName`

Name of the control socket for daemon management. A running `sendmail` daemon can be controlled through this Unix domain socket. Available commands are: `help`, `restart`, `shutdown`, and `status`. The `status` command returns the current number of daemon children, the free disk space (in blocks) of the queue directory, and the load average of the machine expressed as an integer. If not set, no control socket is available. For the sake of security, this Unix domain socket must be in a directory which is accessible only by root; `/var/spool/mqueue/.smcontrol` is recommended for the socket name.

`CRLFile`

File containing certificate revocation status, useful for X.509v3 authentication.

`DaemonPortOptions (Options)`

Sets server SMTP options. The options are *key=value* pairs. Known keys are:

`Name`

User-definable name for the daemon (defaults to “Daemon#”). Used for error messages and logging.

Addr

Address mask (defaults INADDR_ANY).

The address mask may be a numeric address in dot notation or a network name.

Family

Address family (defaults to INET).

InputMailFilters

List of input mail filters for the daemon.

Listen

Size of listen queue (defaults to 10).

Modifier

Options (flags) for the daemon; can be a sequence (without any delimiters) of:

a

Requires authentication.

b

Binds to interface through which mail has been received.

c

Performs hostname canonification (.cf).

f

Requires fully qualified hostname (.cf).

h

Uses name of interface for HELO command.

u

Allows unqualified addresses (.cf).

C

Does not perform hostname canonification.

E

Disallows ETRN (see RFC 2476).

Name

User-definable name for the daemon (defaults to Daemon#). Used for error messages and logging.

Port

Name/number of listening port (defaults to smtp).

ReceiveSize

The size of the TCP/IP receive buffer.

SendSize

The size of the TCP/IP send buffer.

children

Maximum number of children per daemon. See `MaxDaemonChildren`.

DeliveryMode

Delivery mode per daemon. See `DeliveryMode`.

refuseLA

RefuseLA per daemon.

delayLA

DelayLA per daemon.

queueLA

QueueLA per daemon.

`sendmail` listens on a new socket for each occurrence of the `DaemonPortOptions` option in a configuration file.

DataFileBufferSize

Sets the threshold, in bytes, before a memory-based queue data file becomes disk-based. The default is 4096 bytes.

DeadLetterDrop

Defines the location of the system-wide dead.letter file, formerly hard-coded to `/var/tmp/dead.letter`. If this option is not set (the default), `sendmail` does not attempt to save to a system-wide dead.letter file in the event it cannot bounce the mail to the user or postmaster. Instead, it renames the `qf` file as it has in the past when the dead.letter file could not be opened.

DefaultCharSet

Sets the default character set to use when converting unlabeled 8 bit input to MIME.

DefaultUser (*gid*) or (*uid*)

Sets the default group ID for mailers to run in to *gid* or set the default userid for mailers to *uid*. Defaults to 1. The value can also be given as a symbolic group or user name.

DelayLA=*LA*

When the system load average exceeds *LA*, `sendmail` sleeps for one second on most SMTP commands and before accepting connections.

DeliverByMin=*time*

Sets minimum time for Deliver By SMTP Service Extension (RFC 2852). If 0, no time is listed, if less than 0, the extension is not offered, if greater than 0, it is listed as minimum time for the EHLO keyword DELIVERBY.

DeliveryMode (*dx*)

Delivers in mode *x*. Legal modes are:

i

Delivers interactively (synchronously).

b
Delivers in background (asynchronously).

d
Deferred mode. Database lookups are deferred until the actual queue run.

q
Just queues the message (delivers during queue run).

Defaults to **b** if no option is specified, **i** if it is specified but given no argument (that is, **Od** is equivalent to **Odi**).

DHParameters

File containing the DH parameters.

DialDelay

If a connection fails, waits this many seconds and tries again. Zero means “do not retry”.

DontBlameSendmail

If set, overrides the file safety checks. This compromises system security and should not be used. See <http://www.sendmail.org/tips/dontBlameSendmail> for more information.

DontExpandCnames

If set, `$[ ... $]` lookups that do DNS-based lookups do not expand CNAME records.

DontInitGroups

If set, the `initgroups(3C)` routine is never invoked. If you set this, agents run on behalf of users only have their primary (`/etc/passwd`) group permissions.

DontProbeInterfaces

If set, `sendmail` does not insert the names and addresses of any local interfaces into the `=$w` class. If set, you must also include support for these addresses, otherwise mail to addresses in this list bounces with a configuration error.

DontPruneRoutes (R)

If set, does not prune route-addr syntax addresses to the minimum possible.

DoubleBounceAddress

If an error occurs when sending an error message, sends that “double bounce” error message to this address.

EightBitMode (8)

Uses 8-bit data handling. This option requires one of the following keys. The key can be selected by using just the first character, but using the full word is better for clarity.

mimify
Does any necessary conversion of 8BITIME to 7-bit.

pass
Passes unlabeled 8-bit input through as is.

strict

Rejects unlabeled 8-bit input.

ErrorHeader (*Efile/message*)

Appends error messages with the indicated message. If it begins with a slash, it is assumed to be the pathname of a file containing a message (this is the recommended setting). Otherwise, it is a literal message. The error file might contain the name, email address, and/or phone number of a local postmaster who could provide assistance to end users. If the option is missing or NULL, or if it names a file which does not exist or which is not readable, no message is printed.

ErrorMode (*ex*)

Disposes of errors using mode *x*. The values for *x* are:

e

Mails back errors and gives 0 exit status always.

m

Mails back errors.

p

Prints error messages (default).

q

No messages, just gives exit status.

w

Writes back errors (mail if user not logged in).

FallbackMXhost (*Vfallbackhost*)

If specified, the *fallbackhost* acts like a very low priority MX on every host. This is intended to be used by sites with poor network connectivity.

FallBackSmartHost

If specified, the *fallBackSmartHost* is used in a last-ditch effort for each host. This is intended to be used by sites with “fake internal DNS”. That is, a company whose DNS accurately reflects the world inside that company's domain but not outside.

FastSplit

If set to a value greater than zero (the default is one), it suppresses the MX lookups on addresses when they are initially sorted, that is, for the first delivery attempt. This usually results in faster envelope splitting unless the MX records are readily available in a local DNS cache. To enforce initial sorting based on MX records set *FastSplit* to zero. If the mail is submitted directly from the command line, then the value also limits the number of processes to deliver the envelopes; if more envelopes are created they are only queued up and must be taken care of by a queue run. Since the default submission method is by way of SMTP (either from a MUA or by way of the Message Submission Program [MSP]), the value of *FastSplit* is seldom used to limit the number of processes to deliver the envelopes.

ForkEachJob (Y)

If set, delivers each job that is run from the queue in a separate process. Use this option if you are short of memory, since the default tends to consume considerable amounts of memory while the queue is being processed.

ForwardPath (J*path*)

Sets the path for searching for users' .forward files. The default is \$z/.forward. Some sites that use the automounter may prefer to change this to /var/forward/\$u to search a file with the same name as the user in a system directory. It can also be set to a sequence of paths separated by colons; sendmail stops at the first file it can successfully and safely open. For example, /var/forward/\$u:\$z/.forward searches first in /var/forward/*username* and then in *~username*/.forward (but only if the first file does not exist). Refer to the NOTES section for more information.

HeloName=*name*

Sets the name to be used for HELO/EHLO (instead of \$j).

HelpFile (H*file*)

Specifies the help file for SMTP.

HoldExpensive (c)

If an outgoing mailer is marked as being expensive, does not connect immediately.

HostsFile

Sets the file to use when doing “file” type access of host names.

HostStatusDirectory

If set, host status is kept on disk between sendmail runs in the named directory tree. If a full path is not used, then the path is interpreted relative to the queue directory.

IgnoreDots (i)

Ignores dots in incoming messages. This is always disabled (that is, dots are always accepted) when reading SMTP mail.

LogLevel (Ln)

Sets the default log level to *n*. Defaults to 9.

(Mx *value*)

Sets the macro *x* to *value*. This is intended only for use from the command line.

MailboxDatabase

Type of lookup to find information about local mail boxes, defaults to pw which uses [getpwnam\(3C\)](#). Other types can be introduced by adding them to the source code, see `libsm/mbdb.c` for details.

MatchGECOS (G)

Tries to match recipient names using the GECOS field. This allows for mail to be delivered using names defined in the GECOS field in `/etc/passwd` as well as the login name.

MaxDaemonChildren

The maximum number of children the daemon permits. After this number, connections are rejected. If not set or ≤ 0 , there is no limit.

MaxHopCount (hN)

The maximum hop count. Messages that have been processed more than N times are assumed to be in a loop and are rejected. Defaults to 25.

MaxMessageSize

The maximum size of messages that are accepted (in bytes).

MaxMimeHeaderLength= $M[/N]$

Sets the maximum length of certain MIME header field values to M characters. For some of these headers which take parameters, the maximum length of each parameter is set to N if specified. If $/N$ is not specified, one half of M is used. By default, these values are 0, meaning no checks are done.

MaxNOOPCommands= N

Overrides the default of 20 for the number of useless commands.

MaxQueueChildren= N

When set, this limits the number of concurrent queue runner processes to N . This helps to control the amount of system resources used when processing the queue. When there are multiple queue groups defined and the total number of queue runners for these queue groups would exceed `MaxQueueChildren` then the queue groups are not all run concurrently. That is, some portion of the queue groups run concurrently such that `MaxQueueChildren` is not be exceeded, while the remaining queue groups are run later (in round robin order). See `MaxRunnersPerQueue`.

MaxQueueRunSize

If set, limits the maximum size of any given queue run to this number of entries. This stops reading the queue directory after this number of entries is reached; job priority is not used. If not set, there is no limit.

MaxRunnersPerQueue= N

This sets the default maximum number of queue runners for queue groups. Up to N queue runners work in parallel on a queue group's messages. This is useful where the processing of a message in the queue might delay the processing of subsequent messages. Such a delay can be the result of non-erroneous situations such as a low bandwidth connection. The can be overridden on a per queue group basis by setting the `Runners` option. The default is 1 when not set.

MeToo (m)

Sends to me too, even if I am in an alias expansion.

MaxRecipientsPerMessage

If set, allows no more than the specified number of recipients in an SMTP envelope. Further recipients receive a 452 error code and are deferred for the next delivery attempt.

MinFreeBlocks (bN/M)

Insists on at least N blocks free on the file system that holds the queue files before accepting email by way of SMTP. If there is insufficient space, sendmail gives a 452 response to the MAIL command. This invites the sender to try again later. The optional M is a maximum message size advertised in the ESMTP EHLO response. It is currently otherwise unused.

MinQueueAge

Specifies the amount of time a job must sit in the queue between queue runs. This allows you to set the queue run interval low for better responsiveness without trying all jobs in each run. The default value is 0.

MustQuoteChars

Specifies the characters to be quoted in a full name phrase. `&`, `;`, `:`, `\`, `()`, `[]` are quoted automatically.

NiceQueueRun

Specifies the priority of queue runners. See [nice\(1\)](#).

NoRecipientAction

Sets action if there are no legal recipient files in the message. The legal values are:

add-apparently-to

Adds an Apparently-to: header with all the known recipients (which may expose blind recipients).

add-bcc

Adds an empty Bcc: header.

add-to

Adds a To: header with all the known recipients (which may expose blind recipients).

add-to-undisclosed

Adds a To: undisclosed-recipients: header.

none

Does nothing, that is, leaves the message as it is.

OldStyleHeaders (o)

Assumes that the headers may be in old format, that is, spaces delimit names. This actually turns on an adaptive algorithm: if any recipient address contains a comma, parenthesis, or angle bracket, it is assumed that commas already exist. If this flag is not on, only commas delimit names. Headers are always output with commas between the names.

OperatorChars or \$o

Defines the list of characters that can be used to separate the components of an address into tokens.

PidFile

Specifies the filename of the pid file. The default is `/var/run/sendmail.pid`. The filename is macro-expanded before it is opened, and unlinked when sendmail exits.

PostmasterCopy (Ppostmaster)

If set, copies of error messages are sent to the named *postmaster*. Only the header of the failed message is sent. Since most errors are user problems, this is probably not a good idea on large sites, and arguably contains all sorts of privacy violations, but it seems to be popular with certain operating systems vendors.

PrivacyOptions (popt,opt,...)

Sets privacy options. Privacy is really a misnomer; many of these options are just a way of insisting on stricter adherence to the SMTP protocol.

The goaway pseudo-flag sets all flags except noreceipts, restrictmailq, restrictqrun, restrictexpand, noetrn, and nobodyreturn. If mailq is restricted, only people in the same group as the queue directory can print the queue. If queue runs are restricted, only root and the owner of the queue directory can run the queue. The restrict-expand pseudo-flag instructs sendmail to drop privileges when the -bv option is given by users who are neither root nor the TrustedUser so users cannot read private aliases, forwards, or :include: files. It adds the NonRootSafeAddr to the “DontBlame-Sendmail” option to prevent misleading unsafe address warnings. It also overrides the -v (verbose) command line option to prevent information leakage. Authentication Warnings add warnings about various conditions that may indicate attempts to fool the mail system, such as using an non-standard queue directory.

The options can be selected from:

authwarnings

Puts X-Authentication-Warning: headers in messages.

goaway

Disallows essentially all SMTP status queries.

needexpnhelo

Insists on HELO or EHLO command before EXPN.

needmailhelo

Insists on HELO or EHLO command before MAIL.

needvrfyhelo

Insists on HELO or EHLO command before VRFY.

noactualrecipient

Do not put an X-Actual-Recipient line in a DNS that reveals the actual account to which an address is mapped.

noetrn

Disallows ETRN entirely.

noexpn

Disallows EXPN entirely.

`noreceipts`

Prevents return receipts.

`nobodyreturn`

Does not return the body of a message with DSNs.

`novrfy`

Disallows VRFY entirely.

`public`

Allows open access.

`restrictexpand`

Restricts `-bv` and `-v` command line flags.

`restrictmailq`

Restricts `mailq` command.

`restrictqrun`

Restricts `-q` command line flag.

`ProcessTitlePrefix` *string*

Prefixes the process title shown on “`/usr/ucb/ps auxww`” listings with *string*. The string is macro processed.

`QueueDirectory` (*Qdir*)

Uses the named *dir* as the queue directory.

`QueueFactor` (*qfactor*)

Uses *factor* as the multiplier in the map function to decide when to just queue up jobs rather than run them. This value is divided by the difference between the current load average and the load average limit (`x` flag) to determine the maximum message priority to be sent. Defaults to 600000.

`QueueFileMode`=*mode*

Defaults permissions for queue files (octal). If not set, sendmail uses 0600 unless its real and effective uid are different in which case it uses 0644.

`QueueLA` (*xLA*)

When the system load average exceeds *LA*, just queues messages (that is, does not try to send them). Defaults to eight times the number of processors online when sendmail starts.

`QueueSortOrder`=*algorithm*

Sets the algorithm used for sorting the queue. Only the first character of the value is used. Legal values are `host` (to order by the name of the first host name of the first recipient), `filename` (to order by the name of the queue file name), `time` (to order by the submission/creation time), `random` (to order randomly), `modification` (to order by the modification time of the `qf` file (older entries first)), `none` (to not order), and `priority` (to order by message priority). Host ordering makes better use of the connection cache, but may tend to process low priority messages that go to a single host over high priority

messages that go to several hosts; it probably shouldn't be used on slow network links. Filename and modification time ordering saves the overhead of reading all of the queued items before starting the queue run. Creation (submission) time ordering is almost always a bad idea, since it allows large, bulk mail to go out before smaller, personal mail, but may have applicability on some hosts with very fast connections. Random is useful if several queue runners are started by hand which try to drain the same queue since odds are they are working on different parts of the queue at the same time. Priority ordering is the default.

QueueTimeout (*Trtime/wtime*)

Sets the queue timeout to *rtime*. After this interval, messages that have not been successfully sent are returned to the sender. Defaults to five days (5d). The optional *wtime* is the time after which a warning message is sent. If it is missing or 0, then no warning messages are sent.

RandFile

File containing random data (use prefix `file:`) or the name of the UNIX socket if EGD is used (use prefix `egd:`). Note that Solaris supports `random(7D)`, so this does not need to be specified.

RecipientFactor (*yfact*)

The indicated factor *fact* is added to the priority (thus *lowering* the priority of the job) for each recipient, that is, this value penalizes jobs with large numbers of recipients. Defaults to 30000.

RefuseLA (*XLA*)

When the system load average exceeds *LA*, refuses incoming SMTP connections. Defaults to 12 times the number of processors online when `sendmail` starts.

RejectLogInterval

Log interval when refusing connections for this long (default: 3h).

ResolverOptions (*I*)

Tunes DNS lookups.

RetryFactor (*Zfact*)

The indicated factor *fact* is added to the priority every time a job is processed. Thus, each time a job is processed, its priority is decreased by the indicated value. In most environments this should be positive, since hosts that are down are all too often down for a long time. Defaults to 90000.

RrtImpliesDsn

If this option is set, a Return-Receipt-To: header causes the request of a DSN, which is sent to the envelope sender as required by RFC 1891, not to the address given in the header.

RunAsUser

If set, becomes this user when reading and delivering mail. Intended for use of firewalls where users do not have accounts.

SafeFileEnvironment

If set, sendmail does a chroot into this directory before writing files.

SaveFromLine (f)

Saves Unix-style From lines at the front of headers. Normally they are assumed redundant and discarded.

SendMimeErrors (j)

If set, sends error messages in MIME format (see RFC 2045 and RFC 1344 for details). If disabled, sendmail does not return the DSN keyword in response to an EHLO and does not do Delivery Status Notification processing as described in RFC 1891.

ServerCertFile

File containing the cert of the server, that is, this cert is used when sendmail acts as server.

ServerKeyFile

File containing the private key belonging to the server cert.

ServiceSwitchFile

Defines the path to the service-switch file. Since the service-switch file is defined in the Solaris operating environment this option is ignored.

SevenBitInput (7)

Strips input to seven bits for compatibility with old systems. This should not be necessary.

SharedMemoryKey

Specifies key to use for shared memory segment. If not set (or 0), shared memory is not be used. If this option is set, sendmail can share some data between different instances. For example, the number of entries in a queue directory or the available space in a file system. This allows for more efficient program execution, since only one process needs to update the data instead of each individual process gathering the data each time it is required.

SharedMemoryKeyFile=*file*

If SharedMemoryKeyFile is set to -1, the automatically selected shared memory key will be stored in the specified file.

SingleLineFromHeader

If set, From: lines that have embedded newlines are unwrapped onto one line.

SingleThreadDelivery

If this option and the HostStatusDirectory option are both set, uses single thread deliveries to other hosts.

SmtgreetingMessage or \$e

Specifies the initial SMTP greeting message.

SoftBounce

If set, issue temporary errors (4xy) instead of permanent errors (5xy). This can be useful during testing of a new configuration to avoid erroneous bouncing of mail.

StatusFile (*Sfile*)

Logs statistics in the named *file*. By default, this is `/etc/mail/sendmail.st`. As root, you must `touch(1)` this file to enable `mailstats(1)`.

SuperSafe (*s*)

This option can be set to `True`, `False`, `Interactive`, or `PostMilter`. If set to `True`, `sendmail` is set to super-safe when running things, that is, always instantiate the queue file, even if you are going to attempt immediate delivery. `sendmail` always instantiates the queue file before returning control to the client under any circumstances. This should really always be set to `True`. The `Interactive` value has been introduced in 8.12 and can be used together with `DeliveryMode=i`. It skips some synchronization calls which are effectively doubled in the code execution path for this mode. If set to `PostMilter`, `sendmail` defers synchronizing the queue file until any milters have signaled acceptance of the message. `PostMilter` is useful only when `sendmail` is running as an SMTP server; in all other situations it acts the same as `True`.

TempFileMode (*Fmode*)

Specifies the file mode for queue files.

Timeout (*rtimeouts*)

Timeout reads after time interval. The *timeouts* argument is a list of *keyword=value* pairs. All but *command* apply to client SMTP. For backward compatibility, a timeout with no *keyword=* part is set all of the longer values. The recognized timeouts and their default values, and their minimum values specified in RFC 1123 section 5.3.2 are:

```

aconnect
    all connections for a single delivery attempt [0, unspecified]

command
    command read [1h, 5m]

connect
    initial connect [0, unspecified]

control
    complete control socket transaction [2m, none]

datablock
    data block read [1h, 3m]

datafinal
    reply to final . in data [1h, 10m]

datainit
    reply to DATA command [5m, 2m]

fileopen
    file open [60sec, none]

helo
    reply to HELO or EHLO command [5m, none]
```

`hoststatus`
 `host retry` [30m, unspecified]

`iconnect`
 first attempt to connect to a host [0, unspecified]

`ident`
 IDENT protocol timeout [5s, none]

`initial`
 wait for initial greeting message [5m, 5m]

`lhlo`
 wait for reply to an LMTP LHLO command [2m, unspecified]

`mail`
 reply to MAIL command [10m, 5m]

`misc`
 reply to NOOP and VERB commands [2m, none]

`queuereturn`
 undeliverable message returned [5d]

`queuewarn`
 deferred warning [4h]

`quit`
 reply to QUIT command [2m, none]

`rcpt`
 reply to RCPT command [1h, 5m]

`resolver.retrans`
 Resolver's retransmission time interval (in seconds) [varies]. Sets both
 `Timeout.resolver.retrans.first` and `Timeout.resolver.retrans.normal`.

`resolver.retrans.first`
 Resolver's retransmission time interval (in seconds) for the first attempt to deliver a
 message [varies].

`resolver.retrans.normal`
 Resolver's retransmission time interval (in seconds) for all look-ups except the first
 delivery attempt [varies].

`resolver.retry`
 Number of times to retransmit a resolver query [varies]. Sets both
 `Timeout.resolver.retry.first` and `Timeout.resolver.retry.normal`.

`resolver.retry.first`
 Number of times to retransmit a resolver query for the first attempt to deliver a message
 [varies].

`resolver.retry.normal`

Number of times to retransmit a resolver query for all look-ups except the first delivery attempt [varies].

`rset`

reply to RSET command [5m, none]

`starttls`

response to an SMTP STARTTLS command [1h]

`TimeZoneSpec (tzinfo)`

Sets the local time zone info to *tzinfo*, for example, “PST8PDT”. Actually, if this is not set, the TZ environment variable is cleared (so the system default is used); if set but null, the user's TZ variable is used, and if set and non-null, the TZ variable is set to this value.

`TLSSrvOptions`

If this option is 'V', then no client verification is performed, that is, the server does not ask for a certificate.

`TrustedUser`

The user parameter can be a user name (looked up in the passwd map) or a numeric user id. Trusted user for file ownership and starting the daemon. If set, generated alias databases and the control socket (if configured) are automatically owned by this user.

`TryNullMXList (w)`

If you are the “best” (that is, lowest preference) MX for a given host, you should normally detect this situation and treat that condition specially, by forwarding the mail to a UUCP feed, treating it as local, or whatever. However, in some cases (such as Internet firewalls) you may want to try to connect directly to that host as though it had no MX records at all. Setting this option causes `sendmail` to try this. The downside is that errors in your configuration are likely to be diagnosed as “host unknown” or “message timed out” instead of something more meaningful. This option is deprecated.

`UnixFromLine` or `$l`

The “From “ line used when sending to files or programs.

`UnsafeGroupWrites`

If set, group-writable `:include:` and `.forward` files are considered “unsafe”, that is, programs and files cannot be directly referenced from such files.

`UseErrorsTo (l)`

If there is an `Errors-To:` header, sends error messages to the addresses listed there. They normally go to the envelope sender. Use of this option causes `sendmail` to violate RFC 1123. This option is not recommended and deprecated.

`UseMSP`

Uses as mail submission program, that is, allows group writable queue files if the group is the same as that of a set-group-id `sendmail` binary.

UserDatabaseSpec (U)

Defines the name and location of the file containing User Database information.

Verbose (v)

Runs in verbose mode. If this is set, sendmail adjusts the `HoldExpensive` and `DeliveryMode` options so that all mail is delivered completely in a single job so that you can see the entire delivery process. The `Verbose` option should *never* be set in the configuration file; it is intended for command line use only.

XscriptFileBufferSize

Sets the threshold, in bytes, before a memory-based queue transcript file becomes disk-based. The default is 4096 bytes.

If the first character of the user name is a vertical bar, the rest of the user name is used as the name of a program to pipe the mail to. It may be necessary to quote the name of the user to keep sendmail from suppressing the blanks from between arguments.

If invoked as `newaliases`, sendmail rebuilds the alias database, so long as the `/etc/mail/aliases*` files are owned by root *and* root has exclusive write permission. If invoked as `mailq`, sendmail prints the contents of the mail queue.

操作数

address

address of an intended recipient of the message being sent.

用法

See [largefile\(5\)](#) for the description of the behavior of sendmail when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).

退出状态

sendmail returns an exit status describing what it did. The codes are defined in `/usr/include/sysexits.h`.

EX_OK

Successful completion on all addresses.

EX_NOUSER

User name not recognized.

EX_UNAVAILABLE

Catchall. Necessary resources were not available.

EX_SYNTAX

Syntax error in address.

EX_SOFTWARE

Internal software error, including bad arguments.

EX_OSERR

Temporary operating system error, such as “cannot fork”.

EX_NOHOST

Host name not recognized.

EX_TEMPFAIL

Message could not be sent immediately, but was queued.

环境变量

No environment variables are used. However, sendmail's start-up script, invoked by [svcadm\(1M\)](#), reads `/etc/default/sendmail`. In this file, if the variable `ETRN_HOSTS` is set, the start-up script parses this variable and invokes [etrn\(1M\)](#) appropriately. `ETRN_HOSTS` should be of the form:

```
"s1:c1.1,c1.2      s2:c2.1 s3:c3.1,c3.2,c3.3"
```

That is, white-space separated groups of *server:client* where *client* can be one or more comma-separated names. The *:client* part is optional. *server* is the name of the server to prod; a mail queue run is requested for each *client* name. This is comparable to running:

```
/usr/lib/sendmail -qR client
```

on the host *server*.

文件

`dead.letter`

Unmailable text

`/etc/default/sendmail`

Contains default settings. You can override some of the settings by command line options.

`/etc/mail/aliases`

Mail aliases file (ASCII)

`/etc/mail/aliases.db`

Database of mail aliases (binary)

`/etc/mail/aliases.dir`

Database of mail aliases (binary)

`/etc/mail/aliases.pag`

Database of mail aliases (binary)

`/etc/mail/sendmail.cf`

Defines environment for sendmail

`/etc/mail/submit.cf`

Defines environment for MSP

`/etc/mail/trusted-users`

Lists users that are “trusted”, that is, able to set their envelope from address using `-f` without generating a warning message. Note that this file is consulted by the default `sendmail.cf`, but not by the default `submit.cf`, in which the line referring to `/etc/mail/trusted-users` is commented out. See [sendmail\(4\)](#) for instructions on making changes to `submit.cf` and `sendmail.cf`.

`/var/spool/clientmqueue/*`

Temporary files and queued mail

- `/var/spool/mqueue/*`
Temporary files and queued mail
- `~/ . forward`
List of recipients for forwarding messages
- `/usr/include/libmilter/README`
Describes the steps needed to compile and run a filter

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/smtp/sendmail

另请参见

[svcs\(1\)](#), [biff\(1B\)](#), [mail\(1\)](#), [mailq\(1\)](#), [mailx\(1\)](#), [nice\(1\)](#), [check-hostname\(1M\)](#), [check-permissions\(1M\)](#), [etrn\(1M\)](#), [newaliases\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [fork\(2\)](#), [getpwnam\(3C\)](#), [getusershell\(3C\)](#), [resolver\(3RESOLV\)](#), [aliases\(4\)](#), [hosts\(4\)](#), [sendmail\(4\)](#), [shells\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#), [smf\(5\)](#), [random\(7D\)](#)

[tcpd\(1M\)](#), [hosts_access\(4\)](#) in the `security/tcp-wrapper` package.

RFC 2821 Simple Mail Transfer Protocol, John Klensin, April 2001.

RFC 2822 Internet Message Format, Pete Resnick, April 2001.

sendmail, Third Edition, Bryan Costales with Eric Allman, O'Reilly & Associates, Inc., 2003.

<http://www.sendmail.org>

<http://www.milter.org>

附注

The `sendmail` program requires a fully qualified host name when starting. A script has been included to help verify if the host name is defined properly (see [check-hostname\(1M\)](#)).

The permissions and the ownership of several directories have been changed in order to increase security. In particular, access to `/etc/mail` and `/var/spool/mqueue` has been restricted.

Security restrictions have been placed users using `. forward` files to pipe mail to a program or redirect mail to a file. The default shell (as listed in `/etc/passwd`) of these users must be listed in `/etc/shells`. This restriction does not affect mail that is being redirected to another alias.

Additional restrictions have been put in place on `. forward` and `: include:` files. These files and the directory structure that they are placed in cannot be group- or world-writable. See [check-permissions\(1M\)](#).

If you have interfaces that map to domains that have MX records that point to non-local destinations, you might need to enable the `DontProbeInterfaces` option to enable delivery to those destinations. In its default startup behavior, `sendmail` probes each interface and adds an

interface's IP addresses, as well as any domains that those addresses map to, to its list of domains that are considered local. For domains thus added, being on the list of local domains is equivalent to having a 0-preference MX record, with `localhost` as the MX value. If this is not the result you want, enable `DontProbeInterfaces`.

引用名 sf880drd – Sun Fire 880 Dynamic Reconfiguration daemon

用法概要 sf880drd

描述 The Sun Fire 880 Dynamic Reconfiguration daemon, sf880drd, is part of the PCI and system bus hotplug framework. sf880drd starts at boot time. It has no configuration options and does not report any system status.

sf880drd implements the Sun Fire 880 console-less system administration (per-slot pushbuttons and LED status indicators). It also manages various aspects of CPU/memory hotplug.

文件 /usr/platform/SUNW,Sun-Fire-880/lib/sf880drd

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/kernel/dynamic-reconfiguration/sun-fire-880

另请参见 [svcs\(1\)](#), [cfgadm\(1M\)](#), [cfgadm_pci\(1M\)](#), [cfgadm_sbd\(1M\)](#), [svcadm\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The sf880drd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/platform/sun4u/sf880drd

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	sftp-server – SFTP server subsystem
用法概要	<code>/usr/lib/ssh/sftp-server [-f <i>log_facility</i>] [-l <i>log_level</i>]</code>
描述	sftp-server implements the server side of the SSH File Transfer Protocol as defined in the IETF draft-ietf-secsh-filexfer. sftp-server is a subsystem for sshd(1M) and must not be run directly. Command-line flags to sftp-server should be specified in the Subsystem declaration. See sshd_config(4) for more information. To enable the sftp-server subsystem for sshd add the following to <code>/etc/ssh/sshd_config</code>: <pre>Subsystem sftp /usr/lib/ssh/sftp-server</pre> To run sftp-server in a chroot configuration, use <code>internal-sftp</code> instead of <code>/usr/lib/ssh/sftp-server</code>. Otherwise, the chroot directory must contain the necessary files and directories to support the user's session. See the ChrootDirectory and Subsystem options in sshd_config(4) for more information on how sshd and sftp-server work with chroot(2). See sshd_config(4) for a description of the format and contents of that file. There is no relationship between the protocol used by sftp-server and the FTP protocol (RFC 959) provided by <code>in.ftpd</code>. For logging to work, sftp-server must be able to access <code>/dev/log</code>. Use of sftp-server in a chroot configuration therefore requires that syslogd(1M) establish a logging socket inside the chroot directory.
选项	Valid options are listed below. As stated above, these options, if used, are specified in the Subsystem declaration of <code>sshd_config</code>. <code>-f <i>log_facility</i></code> Specifies the facility code that is used when logging messages from sftp-server. The possible values are: DAEMON, USER, AUTH, LOCAL0, LOCAL1, LOCAL2, LOCAL3, LOCAL4, LOCAL5, LOCAL6, LOCAL7. The default is AUTH. <code>-l <i>log_level</i></code> Specifies which messages will be logged by sftp-server. The possible values are: QUIET, FATAL, ERROR, INFO, VERBOSE, DEBUG, DEBUG1, DEBUG2, and DEBUG3. INFO and VERBOSE log transactions that sftp-server performs on behalf of the client. DEBUG and DEBUG1 are equivalent. DEBUG2 and DEBUG3 each specify higher levels of debugging output. The default is ERROR. <code>-u <i>umask</i></code> Sets an explicit umask(2) to be applied to newly-created files and directories, instead of the user's default mask.

- 退出状态

The following exit values are returned:

0

Successful completion.

>0

An error occurred.
- 文件

/usr/lib/ssh/sftp-server Server-side binary.
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/ssh
Interface Stability	Committed

另请参见

[sftp\(1\)](#), [ssh\(1\)](#), [ssh-add\(1\)](#), [ssh-keygen\(1\)](#), [sshd\(1M\)](#), [syslogd\(1M\)](#), [chroot\(2\)](#), [umask\(2\)](#), [sshd_config\(4\)](#), [attributes\(5\)](#)

引用名 shadowd – shadow migration daemon

用法概要 /usr/lib/fs/shadowd

svc:/system/filesystem/shadowd:default

描述 shadowd is a daemon that provides background worker threads to migrate data for a shadow migration. A shadow migration gradually moves data from a source file system into a new “shadow” file system. Users can access and change their data within the shadow file system while migration is occurring.

The shadowd service is managed by the service management facility, [smf\(5\)](#). Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

The [svccfg\(1M\)](#) command can be used to manage the following parameter related to shadowd:

config_params/shadow_threads

Number of threads to devote to background migration of data. These threads are global to the entire system, and increasing the number can increase concurrency and the overall speed of migration at the expense of increased resource consumption (network, I/O, and CPU).

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/zfs
Interface Stability	Unstable

另请参见 [shadowstat\(1M\)](#), [zfs\(1M\)](#), [attributes\(5\)](#)

引用名 shadowstat – report shadow migration statistics

用法概要 /usr/sbin/shadowstat [count]

描述 The shadowstat utility reports statistics on in-progress shadow migrations. A line of output is presented for each migrating file system; the output looks similar to the following:

	EST			
	BYTES	BYTES		ELAPSED
DATASET	XFRD	LEFT	ERRORS	TIME
pool/newfs	2.10M	28.8G	-	00:02:45

Each line identifies bytes transferred thus far, a rough estimate of bytes left to transfer, number of migration errors, and the time elapsed since the migration was started.

If *count* is unspecified, updated values are presented every few seconds until the program is terminated. If *count* is specified, updates only occur *count* times, after which the program exits.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/zfs
Interface Stability	Unstable

另请参见 [shadowd\(1M\)](#), [zfs\(1M\)](#), [attributes\(5\)](#)

引用名	share – display file system shares or make local file system available for mounting by remote systems
用法概要	<pre>share [-F <i>protocol</i>] -a share [-F <i>protocol</i>] [-o <i>options</i>] [-d <i>description</i>] <i>pathname</i> [<i>sharename</i>] share [-F <i>protocol</i>] [-A]</pre>
描述	The share command defines and publishes a file system share, which means the file system is available for mounting through a sharing protocol. If the <i>-F protocol</i> option is omitted, the first file sharing protocol listed in <i>/etc/dfs/fstypes</i> is used as the default. For a description of NFS-specific share options, see share_nfs(1M). For a description of SMB specific share options, see share_smb(1M). Using the share command to define and publish an NFS or SMB share of a ZFS file system is considered a legacy operation. Consider using the <i>zfs set share</i> command to define NFS or SMB share properties and then set the ZFS <i>sharenfs</i> or <i>sharesmb</i> property to publish the share. For more information, see zfs(1M) and share_smb(1M). In the third form of share command, as shown in the Synopsis above, share displays published shares or, with the <i>-A</i> option, displays all configured (defined) shares.
选项	<pre>-F <i>protocol</i> Specify the file sharing protocol. -o <i>specific_options</i> rw Share <i>pathname</i> is published with read and write access to all clients. This is the default behavior. rw=<i>client[:client]...</i> Share <i>pathname</i> is published with read and write access only to the listed clients. No other systems can access the share pathname. ro Share <i>pathname</i> is published with read-only access to all clients. ro=<i>client[:client]...</i> Share <i>pathname</i> is published with read-only access only to the listed clients. No other systems can access the share pathname. Separate multiple options with commas. Separate multiple operands for an option with colons. See EXAMPLES. -d Provide a comment that describes the file system share to be published.</pre>

- a
Publish all defined shares.
- A
Display all defined shares.

示例

示例 1 Publishing an NFS Share With Read-Only Access
The following command defines and publishes an NFS share of /ufsfs with read-only access.

```
# share -F nfs -o ro /ufsfs
```

示例 2 Publishing an NFS Share with Multiple Share Options
The following command defines and publishes an NFS share of the /export/manuals file system with a netgroup called users_nfs who have read-only access and users from specified hosts who have read and write access.

```
# share -F nfs -o ro=users_nfs,rw=host1:host2:host3 /export/manuals
```

文件

/etc/dfs/dfstab
This file is obsolete. An SMF service publishes NFS or SMB shares at boot time.

/etc/dfs/fstypes
List of file-sharing protocols. NFS is the default file sharing protocol.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	SUNWcs

另请参见

[mountd\(1M\)](#), [nfsd\(1M\)](#), [share_nfs\(1M\)](#), [share_smb\(1M\)](#), [shareall\(1M\)](#), [unshare\(1M\)](#), [zfs\(1M\)](#), [attributes\(5\)](#)

附注

If share commands are invoked multiple times on the same file system, the last share invocation supersedes the previous invocation. The options set by the last share command replace the old options. For example, if read-write permission was granted to usera on /somefs, then you want to grant read-write permission also to userb on /somefs, use the following syntax:

```
example% share -F nfs -o rw=usera:userb /somefs
```

引用名

shareall, unshareall – publish or unpublish multiple shares

用法概要

shareall [-F *FSType* [,*FSType*]...] [-| *file*]
unshareall [-F *FSType* [,*FSType*]...]

描述

When used with no arguments, `shareall` publishes shares from a specified file that contains a list of share command lines. If the operand is a hyphen (-), then the `share` command lines are obtained from standard input. Otherwise, if neither a file nor a hyphen is specified, then shares are published from the SMF shares repository.

Shares can be specified by identifying the file system type in a comma-separated list as an argument to `-F`.

`unshareall` unpublishes all published file system shares. Without an `-F` flag, it unpublishes all distributed file system shares.

选项

`-F FSType` Specify file system type. Defaults to the first entry in `/etc/dfs/fstypes`.

文件

`/etc/dfs/dfstab`

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWcs

另请参见

[share\(1M\)](#), [unshare\(1M\)](#), [attributes\(5\)](#)

引用名	sharectl – 配置和管理文件共享服务						
用法概要	<pre>sharectl [-h] sharectl status [-h] [protocol] sharectl get [-h] [-p property]... protocol sharectl set [-h] [-p property=value]... protocol</pre>						
描述	sharectl 命令对文件共享协议（例如 NFS 和 SMB）执行操作。此命令可设置客户机和服务器操作属性，获取和恢复配置快照，以及获取协议服务的状态。 get 和 set 子命令（请参见下文）要求使用 root 特权或者要求您承担主管理员角色。授权的用户可以使用 sharectl 设置 Solaris 服务器管理工具中的 NFS 和 SMB 属性的全局值。请参见 nfs(4) 和 smb(4)。						
选项	支持以下选项： -h 显示用法消息。 -p <i>property</i>[= <i>value</i>] 指定一个属性。请参见下文的“子命令”部分。						
子命令	sharectl 支持下面描述的子命令。sharectl 命令的格式是： <pre># sharectl subcommand [option]</pre> <table><tr><td>get [-p <i>property</i>] <i>protocol</i></td><td>获取指定协议的属性值。如果未提供 -p 选项，则将获取指定协议的所有属性。</td></tr><tr><td>set [-p <i>property</i>=<i>value</i>]... <i>protocol</i></td><td>设置指定的文件共享协议的属性。</td></tr><tr><td>status [<i>protocol</i>]</td><td>显示指定协议的状态，如果未指定协议，则显示所有文件共享协议的状态。</td></tr></table>	get [-p <i>property</i>] <i>protocol</i>	获取指定协议的属性值。如果未提供 -p 选项，则将获取指定协议的所有属性。	set [-p <i>property</i> = <i>value</i>]... <i>protocol</i>	设置指定的文件共享协议的属性。	status [<i>protocol</i>]	显示指定协议的状态，如果未指定协议，则显示所有文件共享协议的状态。
get [-p <i>property</i>] <i>protocol</i>	获取指定协议的属性值。如果未提供 -p 选项，则将获取指定协议的所有属性。						
set [-p <i>property</i> = <i>value</i>]... <i>protocol</i>	设置指定的文件共享协议的属性。						
status [<i>protocol</i>]	显示指定协议的状态，如果未指定协议，则显示所有文件共享协议的状态。						
示例	示例1 获取属性 以下命令获取 NFS 协议的属性。 <pre>% sharectl get nfs lockd_listen_backlog=32 lockd_servers=20 lockd_retransmit_timeout=5 grace_period=90 nfsmapid_domain=sun.com servers=16 server_versmin=2 server_versmax=4 max_connections=-1</pre> 以下命令获取 NFS 协议的 grace_period 属性的值。						

示例1 获取属性 (续)

```
% sharectl get -p grace_period nfs
grace_period=90
```

示例2 设置属性

请注意，在前面的示例中，服务器 NFS 协议的最低版本 (server_versmin) 设置为 2。以下命令将最低版本号设置为版本 3。

```
% sharectl set -p server_versmin=3 nfs
```

示例3 获取状态

以下命令获取系统上所有文件共享协议的状态。

```
% sharectl status
nfs      enabled
```

示例4 设置SMB服务器的属性

以下命令为 SMB 协议设置 server_signing_required 属性的值。

```
% sharectl set -p server_signing_required=true smb
```

示例5 设置SMB客户机的属性

以下命令为 SMB 协议设置 client_signing_required 属性的值。

```
% sharectl set -p client_signing_required=true smb
```

示例6 设置对 autofs 的 RPC (Remote Procedure Call, 远程过程调用) 调用的跟踪

以下命令扩展了对 autofs 的每个 RPC 调用，并将其记录到在 automountd(1M) 中为该服务指定的位置。

```
# sharectl set trace=1 autofs
```

退出状态

0 成功完成。

非零值 命令失败。

文件

/usr/include/libshare.h 用于退出状态的错误代码。

属性

有关以下属性的说明，请参见 attributes(5)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed (已确定)

另请参见

[automount\(1M\)](#)、[automountd\(1M\)](#)、[lockd\(1M\)](#)、[mountd\(1M\)](#)、[nfsd\(1M\)](#)、[nfsmapid\(1M\)](#)、[nfs\(4\)](#)、[s](#)

引用名	share_nfs – make NFS shares available for mounting by remote systems
用法概要	<pre>share -F nfs [-a [-o <i>specific_options</i>] [-d <i>description</i>] <i>pathname</i> [<i>sharename</i>] [-A]] zfs set share=name=<i>share-name</i>,path=<i>pathname</i>,[desc="<i>description</i>"], prot=nfs,[[<i>property=value</i>[,...]] <i>dataset</i></pre>
描述	The share utility defines and publishes a NFS share, which makes a local file system available for mounting by remote systems. It starts the nfsd(1M) and mountd(1M) daemons if they are not already running. You can use the share command to create and publish a ZFS file system share, but this is considered a legacy operation. See zfs(1M) for information about using the zfs set share command and the sharenfs property to create and publish NFS shares.
选项	The following options are supported: -F nfs Specify the NFS file sharing protocol. -a Publish all defined shares. -o <i>specific_options</i> Specify <i>specific_options</i> in a comma-separated list of keywords and attribute-value-assertions for interpretation by the NFS protocol. By default, a share is published with read-write access to all clients, unless a specific option overrides the default access. <i>specific_options</i> can be any combination of the following: aclok Allows the NFS server to do access control for NFS Version 2 clients. When aclok is set on the server, maximal access is given to all clients. For example, with aclok set, if anyone has read permissions, then everyone does. If aclok is not set, minimal access is given to all clients. anon=<i>uid</i> Set <i>uid</i> to be the effective user ID of unknown users. By default, unknown users are given the effective user ID UID_NOBODY. If <i>uid</i> is set to -1, access is denied. charset All clients will be assumed to be using the specified character set (see list in following description) and file and path names will be converted to UTF-8 for the server. charset=<i>access_list</i> Where <i>charset</i> is one of: euc-cn, euc-jp, euc-jpms, euc-kr, euc-tw, iso8859-1, iso8859-2, iso8859-5, iso8859-6, iso8859-7, iso8859-8, iso8859-9, iso8859-13, iso8859-15, koi8-r. Clients that match the <i>access_list</i> for one of these properties will be assumed to be using that character set and file and path names will be converted to UTF-8 for the server.

index=file

Load *file* rather than a listing of the directory containing this file when the directory is referenced by an NFS URL.

log=tag

Enables NFS server logging for the specified file system. The optional *tag* determines the location of the related log files. The *tag* is defined in `etc/nfs/nfslog.conf`. If no *tag* is specified, the default values associated with the *global tag* in `etc/nfs/nfslog.conf` is used. Support of NFS server logging is only available for NFS Version 2 and Version 3 requests.

noaclfab

Allows NFS servers to not return fabricated ACLs to NFS clients. The default behavior for NFS servers is to fabricate ACLs. If `noaclfab` is set, then the NFS server does not fabricate ACLs, which is the appropriate choice if the underlying filesystem does not support the POSIX Draft ACL.

none

Access is disallowed to all clients. The `ro` or `rw` options can override `none`.

none=access_list

Access is not allowed to any client that matches the access list. The exception is when the access list is an asterisk (*), in which case `ro` or `rw` can override `none`.

nosub

Prevents clients from mounting subdirectories of shared directories. For example, if `/export` is shared with the `nosub` option on server *foeey* then a NFS client cannot do:

```
mount -F nfs foeey:/export/home/mnt
```

NFS Version 4 does not use the MOUNT protocol. The `nosub` option only applies to NFS Version 2 and Version 3 requests.

nosuid

By default, clients are allowed to create files on the shared file system with the `setuid` or `setgid` mode enabled. Specifying `nosuid` causes the server file system to silently ignore any attempt to enable the `setuid` or `setgid` mode bits.

public

Moves the location of the public file handle from root (/) to the exported directory for WebNFS-enabled browsers and clients. This option does not enable WebNFS service. WebNFS is always on. Only one file system per server may use this option. Any other option, including the `-ro=list` and `-rw=list` options can be included with the `public` option.

ro

Share is published with read-only access to all clients.

ro=access_list

Share is published with read-only access to the clients listed in *access_list*; overrides the *rw* suboption for the clients specified. See *access_list* below.

root

Root users from all hosts have root access.

root=access_list

Only root users from the hosts specified in *access_list* have root access. See *access_list* below. By default, no host has root access, so root users are mapped to an anonymous user ID (see the *anon=uid* option described above). Netgroups can be used if the file system shared is using UNIX authentication (AUTH_SYS).

root_mapping=uid

For a client that is allowed root access, map the root UID to the specified user id.

rw

Share is published with read and write access to all clients.

rw=access_list

Share is published with read and write access to the clients listed in *access_list*; overrides the *ro* suboption for the clients specified. See *access_list* below.

sec=mode[:mode]. . .

Publishes a share by using one or more of the specified security modes. The *mode* in the *sec=mode* option must be a node name supported on the client. If the *sec=* option is not specified, the default security mode used is AUTH_SYS. Multiple *sec=* options can be specified on the command line, although each mode can appear only once. The security modes are defined in [nfssec\(5\)](#).

Each *sec=* option specifies modes that apply to any subsequent *window=*, *rw*, *ro*, *rw=*, *ro=* and *root=* options that are provided before another *sec=* option. Each additional *sec=* resets the security mode context, so that more *window=*, *rw*, *ro*, *rw=*, *ro=* and *root=* options can be supplied for additional modes.

sec=none

If the option *sec=none* is specified when the client uses AUTH_NONE, or if the client uses a security mode that is not one that the file system is shared with, then the credential of each NFS request is treated as unauthenticated. See the *anon=uid* option for a description of how unauthenticated requests are handled.

secure

This option has been deprecated in favor of the *sec=dh* option.

window=value

When a share is published with *sec=dh*, set the maximum life time (in seconds) of the RPC request's credential (in the authentication header) that the NFS server allows. If a credential arrives with a life time larger than what is allowed, the NFS server rejects the request. The default value is 30000 seconds (8.3 hours).

-d *description*

Provide a comment that describes the file system to be shared.

-A

Display all defined shares.

access_list

The *access_list* argument is either the string "*" to represent all hosts or a colon-separated list whose components may be any number of the following:

hostname

The name of a host. With a server configured for DNS or LDAP naming in the *nsswitch* "hosts" entry, any hostname must be represented as a fully qualified DNS or LDAP name. The hostname specified must be the canonical name for this host and must match the hostname returned on the reverse lookup of the incoming IP address of the NFS client.

netgroup

A netgroup contains a number of hostnames. With a server configured for DNS or LDAP naming in the *nsswitch* "hosts" entry, any hostname in a netgroup must be represented as a fully qualified DNS or LDAP name.

domain name suffix

To use domain membership the server must use DNS or LDAP to resolve hostnames to IP addresses; that is, the "hosts" entry in the */etc/nsswitch.conf* must specify "dns" or "ldap" ahead of "nis", since only DNS and LDAP return the full domain name of the host. Other name services like NIS cannot be used to resolve hostnames on the server because when mapping an IP address to a hostname they do not return domain information. For example,

```
NIS    172.16.45.9 --> "myhost"
```

and

```
DNS or LDAP  172.16.45.9 -->
              "myhost.mydomain.mycompany.com"
```

The domain name suffix is distinguished from hostnames and netgroups by a prefixed dot. For example,

```
rw=.mydomain.mycompany.com
```

A single dot can be used to match a hostname with no suffix. For example,

```
rw=.
```

matches "mydomain" but not "mydomain.mycompany.com". This feature can be used to match hosts resolved through NIS rather than DNS and LDAP.

network

The network or subnet component is preceded by an at-sign (@). It can be either a name or a dotted address. If a name, it is converted to a dotted address by [getnetbyname\(3SOCKET\)](#). For example,

```
=@mynet
```

would be equivalent to:

```
=@172.16 or =@172.16.0.0
```

The network prefix assumes an octet-aligned netmask determined from the zeroth octet in the low-order part of the address up to and including the high-order octet, if you want to specify a single IP address (see below). In the case where network prefixes are not byte-aligned, the syntax allows a mask length to be specified explicitly following a slash (/) delimiter. For example,

```
=@theothernet/17 or =@172.16.132/22
```

...where the mask is the number of leftmost contiguous significant bits in the corresponding IP address.

When specifying individual IP addresses, use the same @ notation described above, without a netmask specification. For example:

```
=@172.16.132.14
```

Multiple, individual IP addresses would be specified, for example, as:

```
root=@172.16.132.20:@172.16.134.20
```

A prefixed minus sign (–) denies access to that component of *access_list*. The list is searched sequentially until a match is found that either grants or denies access, or until the end of the list is reached. For example, if host "terra" is in the "engineering" netgroup, then

```
rw=-terra:engineering
```

denies access to terra but

```
rw=engineering:-terra
```

grants access to terra.

操作数

The following operands are supported:

pathname

The pathname of the file system to be shared.

示例

示例 1 Define and Publish an NFS Share

The following example shows how to define and publish the /export/manuals file system share.

```
# share -F NFS /export/manuals
```

退出状态 The following exit values are returned:

0
 Successful completion.

>0
 An error occurred.

文件

/etc/dfs/fstypes
 list of system types, NFS by default

/etc/dfs/sharetab
 system record of shared file systems

/etc/nfs/nfslogtab
 system record of logged file systems

/etc/nfs/nfslog.conf
 logging configuration file

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/file-system/nfs

另请参见 [mount\(1M\)](#), [mountd\(1M\)](#), [nfsd\(1M\)](#), [nfslogd\(1M\)](#), [share\(1M\)](#), [unshare\(1M\)](#), [zfs\(1M\)](#), [getnetbyname\(3SOCKET\)](#), [nfslog.conf\(4\)](#), [netgroup\(4\)](#), [attributes\(5\)](#), [nfssec\(5\)](#)

附注 Creating and publishing an NFS share with the share command is permanent until the share is unshared. Publishing NFS shares is managed by the following SMF service:

```
$ svcs | grep share
online           Mar_07   svc:/network/shares:default
```

If the file system being shared is a symbolic link to a valid pathname, the canonical path (the path which the symbolic link follows) are shared. For example, if /export/foo is a symbolic link to /export/bar (/export/foo -> /export/bar), the following share command results in /export/bar as the shared pathname (and not /export/foo).

```
# share -F nfs /export/foo
```

An NFS mount of server:/export/foo results in server:/export/bar really being mounted.

The [mountd\(1M\)](#) process allows the processing of a path name the contains a symbolic link. This allows the processing of paths that are not themselves explicitly shared with share_nfs. For example, /export/foo might be a symbolic link that refers to /export/bar which has been specifically shared. When the client mounts /export/foo the mountd processing follows the symbolic link and responds with the /export/bar. The NFS Version 4 protocol does not

use the mountd processing and the client's use of /export/foo does not work as it does with NFS Version 2 and Version 3 and the client receives an error when attempting to mount /export/foo.

引用名	share_smb – make SMB shares available for mounting by remote systems
用法概要	<pre>share -F smb [-a [-o <i>specific-options</i>] [-d <i>description</i>] <i>pathname</i> <i>share-name</i> [-A]] zfs set share=name=<i>share-name</i>,path=<i>pathname</i>,[desc="<i>description</i>"], prot=smb,[<i>property=value</i>[,...]] <i>dataset</i></pre>
描述	The share command defines and publishes a SMB share, which makes a local file system available for mounting by remote systems. You can modify the behavior of SMB shares by setting property values with either the share or zfs set share command. See the share(1M) and zfs(1M) man pages. The share command has the following options: -F smb Share SMB file sharing protocol. -a Publish all defined shares. -o <i>specific-options</i> Specify <i>specific-options</i> in a comma-separated list of keywords and attribute-value-assertions for interpretation by the SMB protocol. By default, a share is published with read-write access to all clients, unless a specific option overrides the default access. <i>specific-options</i> can be any combination of the properties supported by a given file system. -d <i>description</i> Provide a comment that describes the file system to be shared. -A Display all defined shares.
Share Properties	The following SMB share properties are supported and can be set by the zfs and share commands: <i>abe=boolean</i> Sets the access-based enumeration (ABE) policy for a share. When set to true, ABE filtering is enabled on this share and directory entries to which the requesting user has no access will be omitted from directory listings returned to the client. When set to false or not defined, ABE filtering will not be performed on this share. This property is not defined by default. false Disable ABE for this share. true Enable ABE for this share.

ad-container

Specifies the AD container in which to publish shares.

The AD container is specified as a comma-separated list of attribute name-value pairs using the LDAP distinguished name (DN) or relative distinguished name (RDN) format.

The following example uses the `share` command to specify the AD container:

```
$ share -F smb -o name=sh1,abe=true,\
ad-container=cn=sales,ou=mycompany,dc=com /export/home
```

The following example uses the `zfs` command to specify the AD container:

```
$ zfs set share=name=sh1,path=/export/home,prot=smb,\
ad-container=cn=sales,ou=mycompany,dc=com,abe=true rpool/export/home
```

The DN or RDN must be specified in LDAP format using the `cn=`, `ou=`, and `dc=` prefixes:

- `cn` represents the common name
- `ou` represents the organizational unit
- `dc` represents the domain component

`cn=`, `ou=` and `dc=` are attribute types. The attribute type used to describe an object's RDN is called the naming attribute, which, for ADS, includes the following object classes:

- `cn` for the user object class
- `ou` for the organizational unit (OU) object class
- `dc` for the domainDns object class

catia=boolean

Specifies whether to perform CATIA character substitution. CATIA V4 uses characters in file names that are considered to be invalid by Windows. A CATIA V4 file could be inaccessible to Windows clients if the file name contains any of the characters that are considered illegal in Windows. By default, CATIA character substitution is not performed. See [《Oracle Solaris Administration: SMB and Windows Interoperability》](#) 中的“Enabling CATIA V4/V5 Character Translations”.

If the `catia` property is set to `true`, the following character substitution is applied to file names.

CATIA V4 UNIX	CATIA V5 Windows	
"	\250 0x00a8	Dieresis
*	\244 0x00a4	Currency Sign
/	\370 0x00f8	Latin Small Letter O with Stroke
:	\367 0x00f7	Division Sign
<	\253 0x00ab	Left-Pointing Double Angle Quotation Mark
>	\273 0x00bb	Right-Pointing Double Angle Quotation Mark
?	\277 0x00bf	Inverted Question Mark
\	\377 0x00ff	Latin Small Letter Y with Dieresis
	\246 0x00a6	Broken Bar

csc=value

Sets the client-side caching policy for a share. Client-side caching is a client feature and offline files are managed entirely by the clients.

The following are valid values for the `csc` property:

- `manual` – Clients are permitted to cache files from the specified share for offline use as requested by users. However, automatic file-by-file reintegration is not permitted. `manual` is the default value.
- `auto` – Clients are permitted to automatically cache files from the specified share for offline use and file-by-file reintegration is permitted.
- `vdo` – Clients are permitted to automatically cache files from the specified share for offline use, file-by-file reintegration is permitted, and clients are permitted to work from their local cache even while offline.
- `disabled` – Client-side caching is not permitted for this share.

dfsroot=boolean

Marks a share as a distributed file system (DFS) root share to distinguish it from a regular share. By default, `dfsroot` is not defined. If `dfsroot` is `false` or not defined, the share is not a DFS root share.

guestok=boolean

Sets the guest access policy for the share. When set to `true` guest access is allowed on this share. When set to `false` or not defined guest access is not allowed on this share. This property is not defined by default.

An `idmap(1M)` name-based rule can be used to map `guest` to any local user name, such as `guest` or `nobody`. If the local account has a password in `/var/smb/smbpasswd` the guest connection will be authenticated against that password. Any connection made using an account that maps to the local guest account will be treated as a guest connection.

The following name-based rule maps the Windows Guest user to the UNIX guest user:

```
# idmap add winname:Guest unixuser:guest
```

none=access-list

Specifies that access is not allowed to any client that matches the access list. The exception is when the access list is an asterisk (*), in which case `ro` or `rw` can override none.

ro=access-list

Specifies that sharing is read-only to the clients listed in *access-list*. Overrides the `rw` suboption for the clients specified. See *access-list*.

rw=access-list

Specifies that sharing is read-write to the clients listed in *access-list*. Overrides the `ro` suboption for the clients specified. See *access-list*.

Access List Argument The *access-list* argument is either the string "*" to represent all hosts or a colon-separated list whose components may be any number of the following:

hostname

Specifies the name of a host. *hostname* must be a fully qualified DNS or LDAP name when the host specifies these naming schemes in the `hosts` portion of the `nsswitch.conf` file.

netgroup

A netgroup contains a number of host names. Any *hostname* in a netgroup must be a fully qualified DNS or LDAP name when the host specifies these naming schemes in the `hosts` portion of the `nsswitch.conf` file.

domainname.suffix

To use domain membership, the server must use DNS or LDAP to resolve host names to IP addresses. This means that the `hosts` entry of the `/etc/nsswitch.conf` file must specify `dns` or `ldap` before `nis`. You must do this because only DNS and LDAP return the full domain name of the host.

Other naming services, such as NIS, cannot be used to resolve host names on the server because these naming services do not return domain information. For example, the following shows how NIS, DNS, and LDAP return host name information for the 172.16.45.9 IP address:

NIS Returns: myhost

DNS or LDAP Returns: myhost.mydomain.mycompany.com

The domain name suffix is distinguished from host names and netgroups by a prefixed dot. For example, `rw=.mydomain.mycompany.com` matches all host names in `mydomain.mycompany.com`.

The `rw=.` notation uses a single dot to match a host name that has no suffix. This notation matches `mydomain` but not `mydomain.mycompany.com`. This feature can be used to match hosts that are resolved by NIS rather than by DNS and LDAP.

network

The network or subnet component is preceded by an at-sign character (@). It can be either a network name or a dotted address.

A network name is converted to a dotted address by using `getnetbyname(3SOCKET)`. For example, `=@myinet` is equivalent to `=@172.16` or `=@172.16.0.0`.

The network prefix assumes an octet-aligned netmask. The netmask is determined from the zeroth octet in the low-order part of the address up to and including the high-order octet. If network prefixes are not byte-aligned, the syntax permits a mask length to be explicitly specified following a slash delimiter (/). For example, `=@theothernet/17` or `=@172.16.132/22` where the mask is the number of leftmost contiguous significant bits in the corresponding IP address.

When specifying individual IP addresses, use the same @ notation described previously, but do not use a netmask specification. For example, `=@172.16.132.14`.

You can use a colon character (:) to separate multiple, individual IP addresses. For example, `root=@172.16.132.20:@172.16.134.20`.

A prefixed minus sign (–) denies access to that component of *access-list*. The list is searched sequentially until a match is found that either grants or denies access, or until the end of the list is reached. For example, if host "terra" is in the "engineering" netgroup, specifying `rw=- terra:engineering` denies access to terra. However, specifying `rw=engineering: - terra` grants access to terra.

示例

示例 1 Setting a Share Property

The following examples use the `zfs share` and `share` commands to create and publish an SMB share.

- The following example shows how to use the `zfs share` command to create and publish an SMB share that also enables guest access:

```
# zfs set share=name=smb_share,path=/mntpnt/dir2,prot=smb,\
guestok=true tank/home
```

- The following example shows how to use the `share` command to enable guest access on a share:

```
# share -F smb -o guestok=true /mntpnt/dir2
```

示例 2 Viewing the Share Properties on a Share

The following examples show how to use the `zfs get` command and the `/etc/dfs/sharetab` file to view share properties:

- The `zfs get` command enables you to view share properties on the `tank/home` dataset:

```
# zfs get share tank/home
NAME PROPERTY VALUE SOURCE
tank/home share name=smb_share,path=/mntpnt/dir2,guestok=true \
local
```

- The `/etc/dfs/sharetab` file shows all the active shares on the system. The entry for each share shows the properties set and their values:

```
# cat /etc/dfs/sharetab
/mntpnt/dir2 smb_share smb csc=auto,guestok=true
```

文件

`/etc/dfs/sharetab`
System record of shared file systems

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/file-system/smb
Interface Stability	Committed

另请参见 [idmap\(1M\)](#), [share\(1M\)](#), [zfs\(1M\)](#), [zfs\(1M\)](#), [getnetbyname\(3SOCKET\)](#), [netgroup\(4\)](#), [attributes\(5\)](#)

引用名	showmount – show remote mounts				
用法概要	<code>/usr/sbin/showmount [-ade] [hostname]</code>				
描述	showmount lists the clients that have remotely mounted a filesystem from <i>host</i>. This information is maintained by the mountd(1M) server on <i>host</i>, and is saved across crashes in the file <code>/etc/rmtab</code>. The default value for <i>host</i> is the value returned by hostname(1). The showmount command does not display the names of NFS Version 4 clients.				
选项	-a Print all remote mounts in the format: <code>hostname : directory</code> where <i>hostname</i> is the name of the client, and <i>directory</i> is the root of the file system that has been mounted. -d List directories that have been remotely mounted by clients. -e Print the list of shared file systems.				
文件	<code>/etc/rmtab</code>				
属性	See attributes(5) for descriptions of the following attributes:				
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>system/file-system/nfs</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	system/file-system/nfs
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	system/file-system/nfs				
另请参见	hostname(1) , mountd(1M) , attributes(5) 《安装 Oracle Solaris 11 系统》				
已知问题	If a client crashes, its entry is removed from the list of remote mounts on the server.				

引用名	shutdown – shut down system, change system state										
用法概要	<pre>/usr/sbin/shutdown [-y] [-g <i>grace-period</i>] [-r -i <i>init-state</i>] [<i>message</i>]</pre>										
描述	shutdown is executed by the super user to change the state of the machine. In most cases, it is used to change from the multi-user state (state 2) to another state. By default, shutdown brings the system to a state where only the console has access to the operating system. This state is called single-user. Before starting to shut down daemons and killing processes, shutdown sends a warning message and, by default, a final message asking for confirmation. <i>message</i> is a string that is sent out following the standard warning message: The system will be shut down in . . . If the string contains more than one word, it should be contained within single (') or double (") quotation marks. The warning message and the user provided <i>message</i> are output when there are 7200, 3600, 1800, 1200, 600, 300, 120, 60, and 30 seconds remaining before shutdown begins. See EXAMPLES. System state definitions are: <table><tr><td>state 0</td><td>Stop the operating system.</td></tr><tr><td>state 1</td><td>State 1 is referred to as the administrative state. In state 1 file systems required for multi-user operations are mounted, and logins requiring access to multi-user file systems can be used. When the system comes up from firmware mode into state 1, only the console is active and other multi-user (state 2) services are unavailable. Note that not all user processes are stopped when transitioning from multi-user state to state 1.</td></tr><tr><td>state s, S</td><td>State s (or S) is referred to as the single-user state. All user processes are stopped on transitions to this state. In the single-user state, file systems required for multi-user logins are unmounted and the system can only be accessed through the console. Logins requiring access to multi-user file systems cannot be used.</td></tr><tr><td>state 5</td><td>Shut the machine down so that it is safe to remove the power. Have the machine remove power, if possible. The rc0 procedure is called to perform this task.</td></tr><tr><td>state 6</td><td>Stop the operating system and reboot to the state defined by the <code>initdefault</code> entry in <code>/etc/inittab</code>. The rc6 procedure is called to perform this task.</td></tr></table>	state 0	Stop the operating system.	state 1	State 1 is referred to as the administrative state. In state 1 file systems required for multi-user operations are mounted, and logins requiring access to multi-user file systems can be used. When the system comes up from firmware mode into state 1, only the console is active and other multi-user (state 2) services are unavailable. Note that not all user processes are stopped when transitioning from multi-user state to state 1.	state s, S	State s (or S) is referred to as the single-user state. All user processes are stopped on transitions to this state. In the single-user state, file systems required for multi-user logins are unmounted and the system can only be accessed through the console. Logins requiring access to multi-user file systems cannot be used.	state 5	Shut the machine down so that it is safe to remove the power. Have the machine remove power, if possible. The rc0 procedure is called to perform this task.	state 6	Stop the operating system and reboot to the state defined by the <code>initdefault</code> entry in <code>/etc/inittab</code> . The rc6 procedure is called to perform this task.
state 0	Stop the operating system.										
state 1	State 1 is referred to as the administrative state. In state 1 file systems required for multi-user operations are mounted, and logins requiring access to multi-user file systems can be used. When the system comes up from firmware mode into state 1, only the console is active and other multi-user (state 2) services are unavailable. Note that not all user processes are stopped when transitioning from multi-user state to state 1.										
state s, S	State s (or S) is referred to as the single-user state. All user processes are stopped on transitions to this state. In the single-user state, file systems required for multi-user logins are unmounted and the system can only be accessed through the console. Logins requiring access to multi-user file systems cannot be used.										
state 5	Shut the machine down so that it is safe to remove the power. Have the machine remove power, if possible. The rc0 procedure is called to perform this task.										
state 6	Stop the operating system and reboot to the state defined by the <code>initdefault</code> entry in <code>/etc/inittab</code> . The rc6 procedure is called to perform this task.										
选项	-y Pre-answer the confirmation question so the command can be run without user intervention.										

- g *grace-period*
Allow the super user to change the number of seconds from the 60-second default.
- i *init-state*
If there are warnings, *init-state* specifies the state *init* is to be in. By default, system state ‘s’ is used.
- r
Equivalent to specifying -i6.

示例

```
示例1 Using shutdown

In the following example, shutdown is being executed on host foo and is scheduled in 120
seconds. The warning message is output 2 minutes, 1 minute, and 30 seconds before the final
confirmation message.

example# shutdown -i S -g 120 "==== disk replacement ====="
Shutdown started.   Tue Jun  7 14:51:40 PDT 1994

Broadcast Message from root (pts/1) on foo Tue Jun  7 14:51:41. . .
The system will be shut down in 2 minutes
==== disk replacement =====
Broadcast Message from root (pts/1) on foo Tue Jun  7 14:52:41. . .
The system will be shut down in 1 minutes
==== disk replacement =====
Broadcast Message from root (pts/1) on foo Tue Jun  7 14:53:41. . .
The system will be shut down in 30 seconds
==== disk replacement =====
Do you want to continue? (y or n):
```

文件

/etc/inittab controls process dispatching by *init*

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[boot\(1M\)](#), [halt\(1M\)](#), [init\(1M\)](#), [killall\(1M\)](#), [reboot\(1M\)](#), [ufsdump\(1M\)](#), [init.d\(4\)](#), [inittab\(4\)](#), [nologin\(4\)](#), [attributes\(5\)](#)

附注

When a system transitions down to the S or s state, the /etc/nologin file (see [nologin\(4\)](#)) is created. Upon subsequent transition to state 2 (multi-user state), this file is removed by a script in the /etc/rc2.d directory.

引用名	slpd – Service Location Protocol Daemon								
用法概要	<code>/usr/lib/inet/slpd [-f <i>configuration-file</i>]</code>								
描述	The <code>slpd</code> daemon provides common server functionality for the Service Location Protocol (“SLP”) versions 1 and 2, as defined by IETF in RFC 2165 and RFC 2608. SLP provides a scalable framework for the discovery and selection of network services. <code>slpd</code> provides the following framework services: <table><tr><td>Directory Agent</td><td>This service automatically caches service advertisements from service agents to provide them to user agents, and makes directory agent advertisements of its services. This service is optional. <code>slpd</code> does not provide directory agent service by default. Directory agents are not databases, and they do not need to be maintained.</td></tr><tr><td>Service Agent Server</td><td>All service agents on the local host register and deregister with this server. This service responds to all requests for services, and forwards registrations to directory agents. By default, <code>slpd</code> is a service agent server.</td></tr><tr><td>Passive Directory Agent Discovery</td><td>This service listens for directory agent advertisements and maintains a table of active directory agents. When a user agent wishes to discover a directory agent, it can simply query <code>slpd</code>, obviating the need to perform discovery by means of multicast. By default, <code>slpd</code> performs this service.</td></tr><tr><td>Proxy Registration</td><td>This service can act as a proxy service agent for services that cannot register themselves. <code>slpd</code> reads the proxy registration file for information on services it is to proxy. By default, no services are registered by proxy.</td></tr></table> All configuration options are available from the configuration file. <code>slpd</code> reads its configuration file upon startup. Stop and start the <code>slpd</code> daemon using <code>svcadm(1M)</code>. Use the command <code>svcadm enable network/slp</code> to start the <code>slpd</code> daemon. Use the command <code>svcadm disable network/slp</code> to stop it. The file <code>/etc/inet/slp.conf</code> must exist before the <code>slp</code> service can start the daemon. Only the example file <code>/etc/inet/slp.conf.example</code> is present by default. To enable SLP, copy <code>/etc/inet/slp.conf.example</code> to <code>/etc/inet/slp.conf</code>.	Directory Agent	This service automatically caches service advertisements from service agents to provide them to user agents, and makes directory agent advertisements of its services. This service is optional. <code>slpd</code> does not provide directory agent service by default. Directory agents are not databases, and they do not need to be maintained.	Service Agent Server	All service agents on the local host register and deregister with this server. This service responds to all requests for services, and forwards registrations to directory agents. By default, <code>slpd</code> is a service agent server.	Passive Directory Agent Discovery	This service listens for directory agent advertisements and maintains a table of active directory agents. When a user agent wishes to discover a directory agent, it can simply query <code>slpd</code> , obviating the need to perform discovery by means of multicast. By default, <code>slpd</code> performs this service.	Proxy Registration	This service can act as a proxy service agent for services that cannot register themselves. <code>slpd</code> reads the proxy registration file for information on services it is to proxy. By default, no services are registered by proxy.
Directory Agent	This service automatically caches service advertisements from service agents to provide them to user agents, and makes directory agent advertisements of its services. This service is optional. <code>slpd</code> does not provide directory agent service by default. Directory agents are not databases, and they do not need to be maintained.								
Service Agent Server	All service agents on the local host register and deregister with this server. This service responds to all requests for services, and forwards registrations to directory agents. By default, <code>slpd</code> is a service agent server.								
Passive Directory Agent Discovery	This service listens for directory agent advertisements and maintains a table of active directory agents. When a user agent wishes to discover a directory agent, it can simply query <code>slpd</code> , obviating the need to perform discovery by means of multicast. By default, <code>slpd</code> performs this service.								
Proxy Registration	This service can act as a proxy service agent for services that cannot register themselves. <code>slpd</code> reads the proxy registration file for information on services it is to proxy. By default, no services are registered by proxy.								

选项 The following options are supported:

 -f *configuration-file* Specify an alternate configuration file

示例 示例1 Stopping the slpd daemon

 The following command stops the slpd daemon:

 example# svcadm disable network/slp

 示例2 Restarting the slpd daemon

 The following command restarts the slpd daemon:

 example# svcadm restart network/slp

文件 /etc/inet/slp.conf The default configuration file

 slpd.reg The proxy registration file

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/slp, service/network/slp
CSI	Enabled
Interface Stability	Committed

另请参见 [svcs\(1\)](#), [svcadm\(1M\)](#), [slp_api\(3SLP\)](#), [slp.conf\(4\)](#), [slpd.reg\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#), [slp\(7P\)](#)

 《Oracle Solaris 管理：网络服务》

 Guttman, E., Perkins, C., Veizades, J., and Day, M., RFC 2608, Service Location Protocol, Version 2, The Internet Society, June 1999.

附注 The slpd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

 svc:/network/slp

 Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	smatrrpop – populate security attribute databases in a name service
用法概要	<pre>smatrrpop [-c] [-f] [-m] [-p <i>policy</i>] [-r] -s <i>scope</i> -t <i>scope</i> [-v] <i>database</i></pre>
描述	The smatrrpop command updates the auth_attr(4), exec_attr(4), prof_attr(4), and user_attr(4) role-based access control databases in a target NIS, LDAP, or local /etc files name service from the corresponding databases in a source name service or files. This command processes the table entries from the source database and merges each source entry field into the same field in the corresponding table entry in the target database. If a source entry does not exist in the target database, the entry is created. If the source entry exists in the target database, the fields are merged or replaced according to the command options. Any errors encountered while updating the target entry are reported to stdout, and the command continues with the next source database entry.
选项	The following options are supported: -c Performs cross-table checking. If you specify this option and a check error occurs, a message identifying the check error is written to stdout. The target entry values are checked against entries in related databases: ▪ auths values — Each value must exist as the name of an authorization in the auth_attr(4) database. ▪ profiles values — Each value must exist as a name of a profile in the prof_attr(4) database. ▪ roles values — Each value must exist as the name of a role identity in the user_attr(4) database. ▪ For each exec_attr(4) entry in the source database, the name must exist as the name of a profile in the prof_attr(4) database. -f Specifies that the value in each field in the source entry replaces the value in the corresponding field in the target entry, if the source entry field has a non-empty value. -m For the auths, profiles, and roles attributes, specifies that the values in each field in the source entry are merged with the values in the corresponding target entry field. If a source value does not exist in the target field, the value is appended to the set of target values. If the target field is empty, the source values replace the target field. The attribute values that merge depend on the database being updated: ▪ prof_attr(4) — the auths and profiles attribute values are merged. ▪ user_attr(4) — the auths, profiles, and roles attribute values are merged. ▪ exec_attr(4) — the uid, gid, euid, and egid values are merged.

- p *policy* Specifies the value of the policy field in the `exec_attr(4)` database. Valid values are `suser` (standard Solaris superuser) and `tsol` (Trusted Solaris). If you specify this option, only the entries in the source `exec_attr` database with the specified policy are processed. If you omit this option, all entries in the source `exec_attr` database are processed.
- r Specifies that role identities in the `user_attr(4)` database in the source name service are processed. If you omit this option, only the normal user entries in the `user_attr` source database are processed.
- s *scope* Specifies the source name service or local file directory for database updates, using the following syntax:
- type:/server/domain*
- where *type* indicates the type of name service. Valid values for *type* are:
- `file` — local files
 - `nis` — NIS name service
 - `ldap` — LDAP name service
- server* indicates the local host name of the Solaris system on which the `smatrrpop` command is executed, and on which both the source and target databases exist.
- domain* specifies the management domain name for the name service.
- You can use two special cases of *scope* values:
- To indicate the databases in the `/etc/security` local system directory, use the scope `file:/server`, where *server* is the name of the local system.
 - To load from databases in an arbitrary directory on the Solaris server, use the scope `file:/server/pathname`, where *server* is the name of the local system and *pathname* is the fully-qualified directory path name to the database files.
- t *scope* Specifies the target name service or local file directory for database updates, using the following syntax:
- type:/server/domain*
- where *type* indicates the type of name service. Valid values for *type* are:
- `file` — local files
 - `nis` — NIS name service
 - `ldap` — LDAP name service
- server* indicates the local host name of the Solaris system on which the `smatrrpop` command is executed, and on which both the source and target databases exist.
- domain* specifies the management domain name for the name service.

You can use two special cases of *scope* values:

- To indicate the databases in the `/etc/security` local system directory, use the scope `file:/server`, where *server* is the name of the local system.
- To update to databases in an arbitrary directory on the Solaris server, use the scope `file:/server/pathname`, where *server* is the name of the local system and *pathname* is the fully-qualified directory path name to the database files.

-v Specifies that verbose messages are written. A message is written to `stdout` for each entry processed.

操作数

The following operands are supported:

database Populates one or all databases. You can specify either the name of the database you want to process (for example, `auth_attr`), or `all` to process all databases. If you specify `all`, the databases are processed in the following order:

1. `auth_attr(4)`
2. `prof_attr(4)`
3. `exec_attr(4)`
4. `user_attr(4)`

示例

示例 1 Populating all tables in the NIS name service

The following example merges the values from all four attribute databases in the `/etc/security` directory of the local system into the corresponding tables in the NIS domain, `east.example.com`. The command is executed on the master server, `hoosier`, for the NIS domain and the source files are in the `/etc` and `/etc/security` directories on the NIS master server. No cross-table checking is performed. A summary message indicating the number of entries processed and updated for each table is written to `stdout`.

```
/usr/sadm/bin/smatrrpop -s file:/hoosier \
-t nis:/hoosier/east.example.com all
```

环境变量

See [environ\(5\)](#) for a description of the `JAVA_HOME` environment variable, which affects the execution of the `smatrrpop` command. If this environment variable is not specified, the `/usr/java` location is used.

退出状态

Any errors encountered while updating the target entry are reported to `stdout`. The following exit values are returned:

- 0 The specified tables were updated. Individual entries may have encountered checking errors.
- 1 A syntax error occurred in the command line.
- 2 A fatal error occurred and the tables were not completely processed. Some entries may have been updated before the failure.

文件	<code>/etc/security/auth_attr</code>	Authorization description database. See auth_attr(4) .
	<code>/etc/security/exec_attr</code>	Execution profiles database. See exec_attr(4) .
	<code>/etc/security/prof_attr</code>	Profile description database. See prof_attr(4) .
	<code>/etc/user_attr</code>	Extended user attribute database. See user_attr(4) .

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmga

另请参见 [auth_attr\(4\)](#), [exec_attr\(4\)](#), [prof_attr\(4\)](#), [user_attr\(4\)](#), [attributes\(5\)](#), [environ\(5\)](#)

引用名 `smbadm` – 配置和管理 SMB 本地组和用户，管理域成员身份，管理持久性口令信息，以及发出各种命令

用法概要

```
smbadm add-key [-u username]

smbadm add-member -m member [[-m member] ...] group

smbadm create-group [-d description] group

smbadm crypt

smbadm delete-group group

smbadm disable-user username

smbadm enable-user username

smbadm get-group [[-p property] ...] group

smbadm join -u username domain

smbadm join -w workgroup

smbadm lookup-server //server

smbadm lookup-user [-u username] name | SID

smbadm print [-u username] //server/share {print_file|-}

smbadm remove-key [-u username]

smbadm remove-member -m member [[-m member] ...] group

smbadm rename-group group new-group

smbadm set-group -p property=value [[-p property=value] ...] group

smbadm show-connections [-t] [-u username] [-c computername | -s sharename]
    server

smbadm show-domains

smbadm show-files [-t] [-u username] server

smbadm show-groups [-m] [-p] [group]

smbadm show-sessions [-t] [-u username] server

smbadm show-shares [-t] [-A | -u username] server
```

描述

`smbadm` 命令用于配置 SMB 本地组、管理域成员身份、管理持久性口令信息，以及发出各种命令。您还可以使用 `smbadm` 命令启用或禁用各个本地用户的 SMB 口令生成。

当 Windows 帐户必须是某些本地组的成员时以及必须授予 Windows 式特权时，可以使用 SMB 本地组。Solaris 本地组无法提供这些功能。

有两种类型的本地组：用户定义的本地组和内置本地组。内置本地组是用于支持常用管理任务的预定义本地组。

为了在 SMB 本地组与 Solaris 组之间提供正确的标识映射，SMB 本地组必须具有对应的 Solaris 组。此项要求有两个推论：首先，组名必须同时符合 Windows 和 Solaris 组名规则。因此，SMB 本地组名称可包含最多八 (8) 个字符，并且只包含小写字母和数字。其次，创建 SMB 本地组之前，必须先创建一个 Solaris 本地组。

内置组是由 SMB 服务预定义的标准 Windows 组。无法添加、删除或重命名内置组，而且这些组不遵循 SMB 本地组命名约定。

当 SMB 服务器启动后，下列内置组可用：

Administrators

组成员可以管理系统。

Backup Operators

组成员可以绕过文件访问控制来备份和恢复文件。

Power Users

组成员可以共享目录。

Solaris 本地用户必须具有 SMB 口令才能进行验证，才能访问 SMB 资源。该口令是在系统的 PAM 配置中添加有 `pam_smb_password` 模块的情况下使用 `passwd(1)` 命令创建的。请参见 `pam_smb_passwd(5)` 手册页。

`disable-user` 和 `enable-user` 子命令控制特定本地用户的 SMB 口令生成。被禁用时，用户无法连接到 Solaris SMB 服务。缺省情况下，会为所有本地用户启用 SMB 口令生成。

要重新启用已禁用的用户，必须使用 `enable-user` 子命令，然后使用 `passwd` 命令重置用户的口令。必须将 `pam_smb_passwd.so.1` 模块添加到系统的 PAM 配置中才能生成 SMB 口令。

转义反斜杠字符

对于 `add-member`、`remove-member` 以及 `join`（带 `-u`）子命令，转义字符（\）是用于成员或用户名与域名之间的一个有效分隔符。反斜杠字符是一个 shell 特殊字符，而且必须括起。例如，可以通过一个反斜杠字符对另一个反斜杠字符进行转义：`domain\username`。有关处理 shell 特殊字符的更多信息，请参阅适用于您的 shell 的手册页。

操作数

`smbadm` 命令使用以下操作数：

域

指定要加入的现有 Windows 域的名称。

group

指定 SMB 本地组的名称。

username

指定 Windows 用户名。`username` 可以指定为下列任一格式：

```
domain\username[+password]
domain/username[+password]
username@domain
username
```


...其中 *domain* 可以是 NetBIOS 或 DNS 域名。

server

指定本地主机的名称或 IP 地址。

子命令

smbadm 命令包括以下子命令：

add-key [-k *username*]

指定用于 SMB 服务器用户帐户的持久性口令信息。指定此信息后，在非 Kerberos 配置中无需提示口令即可完成挂载。Kerberos 站点应该会自动使用 Kerberos，不会提示输入口令。如果缺省域在 SMF 中可用，则可以省略域。如果未指定用户名，则会使用 Solaris 用户帐户名称。也可以使用加密（散列）的口令，请参见 **crypt** 子命令。该命令还可以从标准输入读取口令，提示标准输入是否为 TTY。

还可以通过用服务器名称替代域名称来为特定服务器存储口令。

持久性口令信息还为运行命令的用户存储在 `/var/smb/smbfspasswd` 中。

add-member -m *member* [[-m *member*] ...] *group*

将指定的成员添加到指定的 SMB 本地组。-m *member* 选项指定 SMB 本地组成员的名称。成员名称必须包括一个现有的用户名和一个可选的域名。

以下列格式之一指定成员名称：

`[domain\]username`

`[domain/]username`

例如，有效的成员名称可以是 `sales\terry` 或 `sales/terry`，其中 `sales` 是 Windows 域名，`terry` 是 `sales` 域中的一个用户的名称。

create-group [-d *description*] *group*

使用指定的名称创建一个 SMB 本地组。您可以通过使用 -d 选项来指定组的描述。

crypt

创建口令散列。此子命令会提示输入口令，并将散列写入标准输出。该散列值适合用作 `smbfs` 挂载和各种 `smbadm` 子命令的加密口令选项的值。

delete-group *group*

删除指定的 SMB 本地组。内置组无法删除。

disable-user *username*

为指定的本地用户禁用 SMB 口令生成功能。被禁用的本地用户无法通过 SMB 服务访问系统。如果某个本地用户帐户被禁用，在重新启用该用户帐户前，您无法使用 `passwd` 命令来修改该用户的 SMB 口令。

enable-user *username*

为指定的本地用户启用 SMB 口令生成功能。重新启用口令生成功能后，必须使用 `passwd` 命令来为本地用户生成 SMB 口令，然后才能连接到 SMB 服务。

如果已将 `pam_smb_passwd` 模块添加到系统的 PAM 配置，则 `passwd` 命令可以同时管理用户的 Solaris 口令和 SMB 口令。

`get-group [[-p property=value] ...] group`

检索指定组的属性值。如果未指定任何属性，将显示所有属性值。

`join -u username domain`

加入一个 Windows 域或工作组。

SMB 服务的缺省模式是工作组模式，该模式使用缺省的工作组名称 `WORKGROUP`。

要加入域，必须使用经过验证的用户帐户，因此必须使用 `-u` 选项来指定 Windows 管理用户名称。如果未在命令行中指定口令，将提示用户输入口令。该用户应当是域管理员或对目标域具有管理特权的用户。

username 和 *domain* 可以采用下面的任一格式输入：

```
username[+password] domain
domain\username[+password]
domain/username[+password]
username@domain
```

...其中 *domain* 可以是 NetBIOS 或 DNS 域名。

如果系统的计算机信任帐户已存在于域控制器中，在加入域时，可以使用经过验证的任意用户帐户。不过，如果计算机信任帐户不存在，则必须使用对域具有管理特权的帐户才能加入域。

`join -w workgroup`

加入一个 Windows 域或工作组。

使用 `join` 子命令时，`-w workgroup` 选项指定要加入的工作组的名称。

`lookup-server //server`

将指定的服务器解析为 IP 地址、NetBIOS 域和 NetBIOS 服务器名称。*server* 可以为以下值之一：

- NetBIOS 主机名
- DNS 主机名
- IP 地址

`lookup-user [-u username] name | SID`

解析当前域或任意受信任域中某个帐户的名称或 SID 的相关信息。

`print [-u username] //server/share {print_file|-}`

将文件输出到指定的远程打印机。如果 *print_file* 是连字符 (-)，则读取标准输入。如果缺省域在 SMF 中可用，则可以省略域。如果未指定用户名，则会使用 Solaris 用户帐户名称。也可以使用加密（散列）的口令，请参见 `crypt` 子命令。该命令也可以通过重定向获取口令。

`remove-key [-u username]`

为正在运行命令的用户清除口令。还将为正在运行命令的用户清除 `/var/smb/smbpasswd` 中的口令。

名称的用户名和域名部分是可选的。如果缺省域在 SMF 中可用，则可以省略域。如果未指定 *username*，将删除为正在运行命令的用户存储的所有密钥。

如果为特定服务器存储用户口令，则应指定服务器名称替代域名称。

remove-member -m member *[[-m member] ...] group*

从指定的 SMB 本地组中删除指定成员。*-m member* 选项指定 SMB 本地组成员的名称。成员名称必须包括一个现有的用户名和一个可选的域名。

以下列格式之一指定成员名称：

[domain\]username

[domain/]username

例如，有效的成员名称可以是 *sales\terry* 或 *sales/terry*，其中 *sales* 是 Windows 域名，*terry* 是 *sales* 域中的一个用户的名称。

rename-group group new-group

重命名指定的 SMB 本地组。该组必须已存在。内置组无法重命名。

set-group -p property=value *[[-p property=value] ...] group*

设置 SMB 本地组的配置属性。无法为内置组更改描述和特权。

-p property=value 选项指定要在指定组上设置的属性的列表。

与组相关的属性如下所示：

backup=[on|off]

指定 SMB 本地组的成员能否绕过文件访问控制来备份文件系统对象。

description=description-text

指定 SMB 本地组的文本描述。

restore=[on|off]

指定 SMB 本地组的成员能否绕过文件访问控制来恢复文件系统对象。

take-ownership=[on|off]

指定 SMB 本地组的成员能否取得文件系统对象的所有权。

show-domains

显示当前工作组或域的相关信息。该信息通常包括工作组名称或主域名称。在域模式下，该信息包括域控制器名称和受信任域的名称。

输出中的每个条目由以下标记之一予以标识：

- [*] 主域
- [.] 本地域
- [-] 其他域
- [+] 选定的域控制器

`show-groups [-m] [-p] [group]`
显示指定的 SMB 本地组的相关信息。如果未指定任何组，则显示所有组的信息。如果指定了 `-m` 选项，还将显示组成员。如果指定了 `-p` 选项，还将显示组特权。

下面的子命令集显示有关用户共享、会话、连接以及在本地或远程服务器上打开的文件的信息。

要显示上面列出的那些类型的数据，需要使用经过验证的用户帐户，因此您必须使用 `-u` 选项来指定 Windows 管理用户名。如果未在命令行中指定口令，将提示用户输入口令。该用户应当是域管理员或对目标域具有管理特权的用户。如果未指定用户名，则会使用 Solaris 用户帐户名称。

`username` 可以采用“操作数”部分中描述的任意格式。

`show-connections [-t] [-u username] [-c computername] [-s sharename] server`
显示在服务器上创建的 SMB 树连接的信息。`-c` 选项指定客户机关注的连接的计算机名称。`-s` 选项指定客户机关注的连接的共享名称。`-t` 选项指定要显示的命令头。

`show-files [-t] [-u username] server`
显示有关在服务器上通过 SMB 打开的文件的信息。`-t` 选项指定要显示的命令头。

`show-sessions [-t] [-u username] server`
显示有关服务器上的 SMB 用户会话的信息。`-t` 选项指定要显示的命令头。

`show-shares [-t] [-A] [-u username] server`
显示有关服务器上的 SMB 共享的信息。`-t` 选项指定要显示的命令头。`-A` 选项指定匿名用户。

退出状态 将返回下列退出值：

- 0 成功完成。
- >0 出现错误。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#) 手册页：

属性类型	属性值
可用性	service/file-system/smb
实用程序名称和选项	Uncommitted（未确定）
实用程序输出格式	Not-An-Interface（不是接口）
smbadm join	Obsolete（已过时）

另请参见

[passwd\(1\)](#)、[groupadd\(1M\)](#)、[idmap\(1M\)](#)、[idmapd\(1M\)](#)、[kclient\(1M\)](#)、[mount_smbfs\(1M\)](#)、[share](#)

引用名	smbiod, smbiod-svc – SMB 客户机 I/O 守护进程
用法概要	<pre>/usr/lib/smbfs/smbiod /usr/lib/smbfs/smbiod-svc</pre>
描述	smbiod 和 smbiod-svc 是 SMB 客户机的内部组件。它们没有可供用户访问的外部接口。有关 SMB 客户机的更多信息，请参见 mount_smbfs(1M) 手册页。 smbiod-svc 是由 SMF 服务 <code>svc:/network/smb/client:default</code> 运行的，用于启动和管理 SMB 客户机所需的 smbiod 进程。 smbiod 是由 smbiod-svc 启动的，后者针对请求 SMB 客户机连接的每个用户作为一个单独的进程运行。这些 smbiod 进程会一直运行，直至所属用户不再具有任何 SMB 客户机连接。
文件	<pre>/usr/lib/smbfs/smbiod SMB I/O 守护进程。 /usr/lib/smbfs/smbiod-svc SMF 服务程序。 /var/run/smbiod/.svc smbiod-svc 在其上对门进行实例化的文件。 /var/run/smbiod/\$UID 用户的 smbiod 进程在其上对门进行实例化的文件。\$UID 是所有者的用户 ID 的数字表示。</pre>
属性	有关以下属性的说明，请参见 attributes(5) 手册页：

属性类型	属性值
可用性	system/file-system/ntfsprogs
接口稳定性	Uncommitted（未确定）

另请参见 [ps\(1\)](#)、[mount_smbfs\(1M\)](#)、[attributes\(5\)](#)、[smbfs\(7FS\)](#)

引用名 smbios – 显示系统管理 BIOS 映像的内容

用法概要 smbios [-BeOsx] [-i *id*] [-t *type*] [-w *file*] [*file*]

描述 smbios 实用程序显示当前系统导出的或文件中存储的系统管理 BIOS (System Management BIOS, SMBIOS) 映像的内容。SMBIOS 是一种行业标准机制，低层次系统软件使用该机制将硬件配置信息导入较高层次的系统管理软件。SMBIOS 数据格式自身是由分布式管理任务组 (Distributed Management Task Force, DMTF) 定义的。有关 SMBIOS 以及如何获取 SMBIOS 规范和实施指南副本的更多信息，请参阅 <http://www.dmtf.org>。

SMBIOS 映像包括一个结构表，每个结构描述系统软件或硬件配置的某个方面。缺省情况下，smbios 显示当前 SMBIOS 映像的完整内容。如果指定了 -s 选项，smbios 将显示映像中存在的结构的摘要。如果指定了 -w 选项，smbios 会将 SMBIOS 映像的副本写入到指定文件。然后，smbios 可以应用于生成的文件以显示其内容。

smbios 尝试以用户可阅读的方式显示每个结构及其内容。如果 smbios 不能识别某个结构的类型或内容，则会显示该结构的原始十六进制数据。

选项 所支持的选项如下：

- B 为损坏的 BIOS 禁用头验证。

缺省情况下，smbios 尝试通过验证锚字符串、头校验和以及版本号来验证 SMBIOS 头。当 BIOS 具有不兼容的头时可能需要此选项。
- e 显示 SMBIOS 入口点的内容，而不是 SMBIOS 结构表的内容。
- i *id* 仅显示指定的结构（由其整数 ID 指定）。
- O 显示过时的结构类型。

缺省情况下，smbios 将省略 DMTF SMBIOS 规范中其类型被标记为已过时的结构的输出。
- s 仅显示结构标识符和类型的摘要列表，而不是每个选定结构的内容。
- t *type* 仅显示其类型与指定的整数类型匹配的那些结构，如 DMTF SMBIOS 规范所定义的那样。
- w *file* 将 SMBIOS 映像的副本写入指定文件并退出。

将 SMBIOS 入口点写入到其结构表地址设置为结构表的文件偏移量的文件的开头，并计算新的入口点校验和。
- x 除适合用户阅读的输出外，还显示所选结构的原始十六进制数据。

缺省情况下，如果 smbios 无法为所选结构显示适合用户阅读的输出，则仅显示十六进制数据。

- 操作数

支持下列操作数：
file 指定要显示的其他 SMBIOS 映像，而不是当前系统的 SMBIOS 映像。
- 退出状态

将返回下列退出值：

0

成功完成。成功检查了 SMBIOS 映像中的所有结构。

1

出现致命错误，如未能打开指定的文件或设备，或映像损坏。

2

指定的命令行选项无效。
- 文件

/dev/smbios 内核 SMBIOS 映像设备。该设备特殊文件用于导出当前系统 SMBIOS 映像的快照。
- 属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	请参见下文。

- 命令行选项为 Committed（已确定）。人类可阅读的输出为 Uncommitted（未确定）。
- 另请参见

[prtdiag\(1M\)](#)、[attributes\(5\)](#)、[smbios\(7D\)](#)
系统管理 BIOS 参考规范（请参见 <http://www.dmtf.org>）
- 附注

系统管理 BIOS 映像的实现完全是由系统和 BIOS 供应商决定的。并不是所有系统都导出 SMBIOS。在不同的系统和 BIOS 供应商之间，SMBIOS 结构内容变化很大，并且经常不遵循规范中包含的准则。BIOS 可能根本不填充某些结构字段，而某些字段可能是由不合适的值填充的。

引用名	smbstat – 显示 Solaris SMB 文件服务器统计信息
用法概要	smbstat [-r [-n [-a -z]] [-t] [-u] [-c] <i>interval</i>
描述	smbstat 命令显示 smbd(1M) 服务器的统计信息。smbstat 具有多个选项（如下所述）和一个操作数 <i>interval</i>。如果指定了 <i>interval</i>，首先将显示自服务器启动以来直至输入命令时捕获的统计信息。接下来显示最后一个 <i>interval</i> 中捕获的统计信息。 缺省情况下，smbstat 命令显示所有统计信息。
选项	smbstat 命令具有以下选项： -c 显示计数器。 -r 显示请求的统计信息。-r 选项可以与以下选项组合使用。 -a 显示所有类型的服务器请求（包括有效的和无效的）的统计信息。请注意，有 256 种类型的服务器请求。-a 和 -z 选项是互斥的。 -n 按字母顺序显示。 -z 显示实际接收的请求的统计信息。-a 选项和 -z 选项是互斥的。 -t 显示 SMB 服务器的吞吐量。 -u 显示 SMB 服务器的使用率。
扩展描述	smbstat 命令显示下面所示的标题。为给定命令显示的标题根据所指定的选项不同而不同。有四种类别的标题，分别对应于四种类别的统计信息：计数器、吞吐量、使用率和（服务器）请求。
计数器	nbt SMB NetBIOS-over-TCP (NBT) 会话的数目。 tcp SMB TCP 会话的数目。 users 登录的用户的数目。 trees 连接的树的数目。 files 打开的文件的数目。

吞吐量

`pipes`

打开的管道的数目。

`rbytes/s`

每秒接收的字节数。

`tbytes/s`

每秒传输的字节数。

`reqs/s`

每秒处理的请求数。

`reads/s`

每秒的读取请求数。这将是下列请求的聚

合：`SMB_COM_READ`、`SMB_COM_LOCK_AND_READ`、`SMB_COM_READ_RAW`和
`SMB_COM_READ_ANDX`。

`writes/s`

每秒的写入请求数。这将是下列请求的聚

合：`SMB_COM_WRITE`、`SMB_COM_WRITE_AND_UNLOCK`、`SMB_COM_WRITE_RAW`和
`SMB_COM_WRITE_AND_CLOSE`。

使用率

`wcnt`

已被接收但在等待 SMB 工作线程执行的平均请求数。

`rcnt`

SMB 工作线程同时执行的平均请求数。

`wtime`

请求在 SMB 工作线程开始执行它之前等待的平均时间。

`rtime`

请求的平均执行时间。

`w%`

至少有一个请求处于等待状态的时间所占的百分比。

`r%`

至少有一个请求处于执行状态的时间所占的百分比。

`u%`

SMB 服务器使用率百分比。此数字的定义为： $rcnt / (max\_worker\_threads)$ 。

`sat`

表示服务器过去是否已经饱和的标志。饱和的定义是 `u% == 100%`。

`usr%`

处理器在用户空间中花费的时间所占的百分比。

`sys%`

处理器在内核空间中花费的时间所占的百分比。

请求

idle%

处理空闲时间所占的百分比。

为每种类型的请求显示以下标题。

code

以十六进制表示的请求代码。

%

给定类型的请求所占的百分比。

count

已接收的请求数。

rbytes/s

每秒接收的字节数。

tbytes/s

每秒接收的字节数。

req/s

每秒处理的请求数。

rt-mean

平均响应时间，以秒为单位。

rt-stddev

响应时间的标准偏差。

示例

示例1 组合选项

以下命令将 -c、-t 和 -u 选项组合使用。

```
% smbstat -ctu
nbt  tcp users trees files pipes
0    1    1    2    20    0

rbytes/s  tbytes/s  reqs/s  reads/s  writes/s
1.036e+02 1.298e+00 0.000e+00 0.000e+00 0.000e+00

wcnt      rcnt      wtime     rtime     w%    r%  u% sat usr% sys% idle%
4.317e-01 7.410e+00 2.461e-05 4.224e-04 31   100 0  no  0    76   24
```

退出状态

将返回下列退出值：

0

成功完成。

>0

出现错误。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#) 手册页：

属性类型	属性值
可用性	service/file-system/smb
接口稳定性	Uncommitted（未确定）
实用程序输出格式	Not-an-Interface（不是接口）

另请参见 [sharectl\(1M\)](#)、[smbadm\(1M\)](#)、[smbd\(1M\)](#)、[attributes\(5\)](#)

引用名 smrsh – restricted shell for sendmail

用法概要 smrsh -c *command*

描述 The smrsh program is intended as a replacement for the sh command in the prog mailer in [sendmail\(1M\)](#) configuration files. The smrsh program sharply limits commands that can be run using the |program syntax of sendmail. This improves overall system security. smrsh limits the set of programs that a programmer can execute, even if sendmail runs a program without going through an alias or forward file.

Briefly, smrsh limits programs to be in the directory /var/adm/sm.bin, allowing system administrators to choose the set of acceptable commands. It also rejects any commands with the characters: , , < , > , | , ; , & , \$, \r (RETURN), or \n (NEWLINE) on the command line to prevent end run attacks.

Initial pathnames on programs are stripped, so forwarding to /usr/ucb/vacation, /usr/bin/vacation, /home/server/mydir/bin/vacation, and vacation all actually forward to /var/adm/sm.bin/vacation.

System administrators should be conservative about populating /var/adm/sm.bin. Reasonable additions are utilities such as [vacation\(1\)](#) and procmail. Never include any shell or shell-like program (for example, perl) in the sm.bin directory. This does not restrict the use of shell or perl scrips in the sm.bin directory (using the #! syntax); it simply disallows the execution of arbitrary programs.

选项 The following options are supported:

-c *command* Where *command* is a valid command, executes *command*.

文件 /var/adm/sm.bin directory for restricted programs

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [sendmail\(1M\)](#), [attributes\(5\)](#)

引用名	smtnrhdb – manage entries in the tn rhdb database
用法概要	<code>/usr/sadm/bin/smtnrhdb subcommand [auth_args] -- subcommand_args</code>
描述	The smtnrhdb command adds, modifies, deletes, and lists entries in the tn rhdb database. The tn rhdb database specifies which remote-host template to use for each host, including the local host, in the distributed system. If a host's IP address cannot be matched to some entry in the tn rhdb database, communication with the host is not permitted.
Valid Host Addresses and Wildcards	The trusted network software uses a network “longest prefix of matching bits” mechanism when looking for a host. The software looks first for the IP address of the host. If the software does not find this address, then the software falls back to searching for an IP address with the longest prefix of a matching bit pattern, and so on. 注 – The actual numeric value of the subnet address or other subnetting information on the system (for example, from the <code>netmasks(4)</code> file) are not considered by this mechanism. Using the “longest prefix of matching bits” mechanism, an IPv4 address of 0.0.0.0 is a wildcard address with a prefix length of 0 and hence matches any IPv4 address. For more information about prefix lengths in IPv4 and IPv6 addresses, see 《Oracle Solaris 管理：IP 服务》. The smtnrhdb command accepts a hostname, IP address, and wildcard address with as optional prefix as valid addresses. See <code>subcommand_args</code>, below, for the format of valid addresses.
子命令	smtnrhdb <i>subcommands</i> are: add Adds a new entry to the tn rhdb database. To add an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations. delete Deletes an entry from the tn rhdb database. To delete an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations. list Lists all entries in the tn rhdb database. To list an entry, the administrator must have the <code>solaris.network.host.read</code> and <code>solaris.network.security.read</code> authorizations. modify Modifies an entry in the tn rhdb database. To modify an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.

选项

The `smtnrhdb` authentication arguments, *auth_args*, are derived from the `smc` arg set. These arguments are the same regardless of which subcommand you use.

The subcommand-specific options, *subcommand_args*, must be preceded by the `--` option.

auth_args

The valid *auth_args* are `-D`, `-H`, `-l`, `-p`, `-r`, and `-u`; they are all optional. If no *auth_args* are specified, certain defaults will be assumed and the user might be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either `-D` or `--domain`.

`-D | --domain domain`

Specifies the default domain that you want to manage. The syntax of *domain=type:/host_name/domain_name*, where *type* is `dns`, `ldap`, or `file`; *host_name* is the name of the server; and *domain_name* is the name of the domain you want to manage.

If you do not specify this option, the system assumes the `file` default domain on whatever server you choose to manage, meaning that changes are local to the server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools.

`-H | --hostname host_name:port`

Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the system connects to the local host on port 898.

`-l | --rolepassword role_password`

Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

`-p | --password password`

Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

`-r | --rolename role_name`

Specifies a role name for authentication. If you do not specify this option, no role is assumed.

`-u | --username user_name`

Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.

`--`

This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the `--` option.

subcommand_args

Note: Descriptions and other arg options that contain white spaces must be enclosed in double quotes.

-h

Displays the command's usage statement.

-H *hostname*

Specifies the name of the host. For the list subcommand, the hostname argument is not specified. This is not required if the ipaddress subcommand argument is specified.

-i *ipaddress*

Specifies the IP address of the host. This is not required if the hostname subcommand argument is specified. This option is not valid with the -w option.

-n *templatename*

Specifies the name of an existing template.

-p *prefixlen*

Specifies the prefix length (in bits) of a wildcard representation of the IP address. The prefix is the left-most portion of the IP address. This option is valid only with the -w option. For example, when the value of -w *ipaddress-wildcard* is 192.168.0.0, a *prefixlen* value of 24 indicates that the wildcard matches all addresses on the 192.168.0 network. With a *prefixlen* of 32, the wildcard 192.168.0.0 matches all addresses on the 192.168.0.0 network.

-w *ipaddress-wildcard*

Specifies the IP address of the subnet using a wildcard.

- One of the following sets of arguments must be specified for subcommand `add`:

```
-H hostname -n templatename |
-i ipaddress -n templatename |
-w ipaddress-wildcard -n templatename [ -p prefixlen ] |
-h
```

- One of the following sets of arguments must be specified for subcommand `modify`:

```
-H hostname -n templatename |
-i ipaddress -n templatename |
-w ipaddress-wildcard -n templatename [ -p prefixlen ] |
-h
```

- One of the following sets of arguments must be specified for subcommand `delete`:

```
-H hostname |
-i ipaddress |
-w ipaddress-wildcard [ -p prefixlen ] |
-h
```

- The subcommand `list` takes the following argument:

```
-h
```

示例

示例 1 Specifying the Template Name for a Wildcard IP Address

The admin role specifies the template name, `cipso_lan`, for a series of hosts that use the IP address wildcard `192.168.113.0` on the local file system. Since no authorization arguments were specified, the administrator connects to port 898 of the local host on the local server with the file domain type, which are the defaults. The administrator is prompted for the admin password.

```
$ usr/sadm/bin/smtnrhdb add -- -w 192.168.113.0 -n cipso_lan
```

示例 2 Deleting an Entry in the tnrhdb Database

The admin role connects to port 898 (which happens to be the default) of the LDAP server and deletes a host entry from the database by specifying its IP address, `192.168.113.8`. Since the domain was not specified, the file domain type and local server are used by default. The administrator is prompted for the admin password.

```
# /usr/sadm/bin/smtnrhdb delete \  
-D ldap:/example.domain -i 192.168.113.8
```

示例 3 Adding a Subnet to the tnrhdb Database

The following command adds all the addresses on the 192.168.55.0 subnet, from 192.168.55.1 to 192.168.55.255, to the tnrhdb database:

```
# /usr/sadm/bin/smtnrhdb add \  
-D file:/machine1.ExampleCo.COM/machine1.ExampleCo.COM \  
-- -w 192.168.55.0 -n cipso  
Authenticating as user: root  
Type /? for help, pressing <enter> accepts the default denoted by [ ]  
Please enter a string value for: password ::  
Loading Tool: com.exampleco.admin.hostmgr.cli.smtnrhdb.HostMgrTnrhdbCli  
from machine1.ExampleCo.COM  
Login to machine1.ExampleCo.COM as user root was successful.  
Download of com.exampleco.admin.hostmgr.cli.smtnrhdb.HostMgrTnrhdbCli  
from machine1.ExampleCo.COM  
was successful.
```

示例 4 Adding Subnet 192.168.0 to the tnrhdb Database

The following command adds all the addresses on the 192.168.0 subnet, from 192.168.0.1 to 192.168.0.255 to the tnrhdb database. The prefix, 24, indicates that the first 24 bits (192.168.0) are fixed. Only the final zero is a wildcard.

```
# /usr/sadm/bin/smtnrhdb add \  
-D file:/machine1.ExampleCo.COM/machine1.ExampleCo.COM \  
-- -w 192.168.0.0 -p 24 -n cipso
```

```
Login to machine1.ExampleCo.COM as user root was successful.  
Download of com.exampleco.admin.hostmgr.cli.smtnrhdb.HostMgrTnrhdbCli  
from machine1.ExampleCo.COM was successful.
```

退出状态 The following exit values are returned:

0
 Successful completion.

1
 Invalid command syntax. A usage message displays.

2
 An error occurred while executing the command. An error message displays.

文件 The following files are used by the smtnrhdb command:

/etc/security/tsol/tnrhdb
 Trusted network remote-host database.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmgts
Interface Stability	Committed

另请参见 [netmasks\(4\)](#), [attributes\(5\)](#)

 《Oracle Solaris 管理：安全服务》

附注 The functionality described on this manual page is available only if the system is configured with Trusted Extensions.

引用名	smtnrhtp – manage entries in the trusted network template database								
用法概要	<code>/usr/sadm/bin/smtnrhtp subcommand [auth_args] -- [subcommand_args]</code>								
描述	The <code>smtnrhtp</code> command adds, modifies, deletes, and lists entries in the <code>tnrhtp</code> database. <code>smtnrhtp</code> subcommands are: <table><tr><td>add</td><td>Adds a new entry to the <code>tnrhtp</code> database. To add an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.</td></tr><tr><td>modify</td><td>Modifies an entry in the <code>tnrhtp</code> database. To modify an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.</td></tr><tr><td>delete</td><td>Deletes an entry from <code>tnrhtp</code> database. To delete an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.</td></tr><tr><td>list</td><td>Lists entries in the <code>tnrhtp</code> database. To list an entry, the administrator must have the <code>solaris.network.security.read</code> authorizations.</td></tr></table>	add	Adds a new entry to the <code>tnrhtp</code> database. To add an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.	modify	Modifies an entry in the <code>tnrhtp</code> database. To modify an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.	delete	Deletes an entry from <code>tnrhtp</code> database. To delete an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.	list	Lists entries in the <code>tnrhtp</code> database. To list an entry, the administrator must have the <code>solaris.network.security.read</code> authorizations.
add	Adds a new entry to the <code>tnrhtp</code> database. To add an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.								
modify	Modifies an entry in the <code>tnrhtp</code> database. To modify an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.								
delete	Deletes an entry from <code>tnrhtp</code> database. To delete an entry, the administrator must have the <code>solaris.network.security.read</code> and <code>solaris.network.security.write</code> authorizations.								
list	Lists entries in the <code>tnrhtp</code> database. To list an entry, the administrator must have the <code>solaris.network.security.read</code> authorizations.								
选项	The subcommand-specific options, <code>subcommand_args</code>, must be <i>preceded</i> by the <code>--</code> option.								
auth_args	The valid <code>auth_args</code> are <code>-D</code>, <code>-H</code>, <code>-l</code>, <code>-p</code>, <code>-r</code>, and <code>-u</code>; they are all optional. If no <code>auth_args</code> are specified, certain defaults will be assumed and the user might be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either <code>-D</code> or <code>--domain</code>. <code>-D --domain domain</code> Specifies the default domain that you want to manage. The syntax of <code>domain=type:/host_name/domain_name</code>, where <code>type</code> is <code>dns</code>, <code>ldap</code>, or <code>file</code>; <code>host_name</code> is the name of the server; and <code>domain_name</code> is the name of the domain you want to manage. If you do not specify this option, the system assumes the <code>file</code> default domain on whatever server you choose to manage, meaning that changes are local to the server. Toolboxes can change the domain on a tool-by-tool basis; this option specifies the domain for all other tools. <code>-H --hostname host_name:port</code> Specifies the <code>host_name</code> and <code>port</code> to which you want to connect. If you do not specify a <code>port</code>, the system connects to the default port, 898. If you do not specify <code>host_name:port</code>, the system connects to the local host on port 898.								

- l | --rolepassword *role_password*
Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.
- p | --password *password*
Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.
- r | --rolename *role_name*
Specifies a role name for authentication. If you do not specify this option, no role is assumed.
- u | --username *user_name*
Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.
- This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the -- option.

subcommand_args

Descriptions and other argument options that contain white spaces must be enclosed in double quotes.

- h Displays the command's usage statement.
 - n *templatename* Specifies the name of the template.
 - t *hosttype* Specifies the host type of the new host. Valid values are unlabeled and cipso. The cipso host type is for hosts that use CIPSO (Common IP Security Options - Tag Type 1 only) to label packets.
 - x doi=*doi-value* Specifies the DOI value (the domain of interpretation). In the case of the unlabeled host type, this is the domain of interpretation for the def_label.
- The domain of interpretation defines the set of rules for translating between the external or internal representation of the security attributes and their network representation. When systems that are configured with Trusted Extensions software have the same doi, they share that set of rules. In the case of the unlabeled host type, these systems also share the same interpretation for the default attributes that are assigned to the unlabeled templates that have that same doi.

- x *max=maximum-label* Specifies the maximum label. Together with *min*, this value specifies the label accreditation range for the remote hosts that use this template. Values can be a hex value or string (such as *admin_high*).
- x *min=minimum-label* Specifies the minimum label. Together with *max*, this value specifies the label accreditation range for the remote hosts that use this template. For gateway systems, *min* and *max* define the default range for forwarding labeled packets. The label range for routes is typically set by using a [route\(1M\)](#) subcommand with the *-secattr* option. When the label range for routes is not specified, the *min* to *max* range in the security template is used. Values can be a hex value or string (such as *admin_low*).
- x *label=default-label* Specifies the default label to be applied to incoming data from remote hosts that do not support these attributes. This option does not apply if *hosttype* is *cipso*. Values can be a hex value or string (such as *admin_low*).
- x *slset=l1,l2,l3,l4* Specifies a set of sensitivity labels. For gateway systems, the labels in *slset* are used for forwarding labeled packets. *slset* is optional. You can specify up to four label values, separated by commas. Values can be a hex value or string (such as *admin_low*).

- One of the following sets of arguments must be specified for subcommand *add*:

```
-n template name (
    -t cipso [ -x doi=doi-value -x min=minimum-label -x max=maximum-label -x
    slset=l1,l2,l3,l4 ] |
    -t unlabeled [ -x doi=doi-value -x min=minimum-label -x max=maximum-label -x
    label=default-label -x slset=l1,l2,l3,l4 ] |
    -h
)
```

- One of the following sets of arguments must be specified for subcommand *modify*:

```
-n template name (
    -t cipso [ -x doi=doi-value -x min=minimum-label -x max=maximum-label -x
    slset=l1,l2,l3,l4 ] |
    -t unlabeled [ -x doi=doi-value -x min=minimum-label -x max=maximum-label -x
    label=default-label -x slset=l1,l2,l3,l4 ] |
    -h
)
```

If the host type is changed, all options for the new host type must be specified.

- One of the following sets of arguments must be specified for subcommand `delete`:

`-n templatename |`
`-h`

- The following argument can be specified for subcommand `list`:

`-n templatename |`
`-h`

示例

示例 1 Adding a New Entry to the Network Template Database

The admin role connects to port 898 of the LDAP server and creates the `unlabeled_ntk` entry in the `tnrhttp` database. The new template is assigned a host type of `unlabeled`, a domain of interpretation of 1, minimum label of `public`, maximum label of `restricted`, and a default label of `needtoknow`. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smtnrhttp \  
add -D ldap:directoryname -H servername:898 -- \  
-n unlabeled_ntk -t unlabeled -x DOI=1 \  
-x min=public -x max=restricted -x label="need to know"
```

退出状态

The following exit values are returned:

- 0 Successful completion.
- 1 Invalid command syntax. A usage message displays.
- 2 An error occurred while executing the command. An error message displays.

文件

The following files are used by the `smtnrhttp` command:

`/etc/security/tso1/tnrhttp` Trusted network remote-host templates.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmgts
Interface Stability	Committed

另请参见

[attributes\(5\)](#)

附注

The functionality described on this manual page is available only if the system is configured with Trusted Extensions.

警告

Changing a template while the network is up can change the security view of an undetermined number of hosts.

Allowing unlabeled hosts onto a Solaris Trusted Extensions network is a security risk. To avoid compromising the rest of your network, such hosts must be *trusted* in the sense that the administrator is certain that these unlabeled hosts will not be used to compromise the distributed system. These hosts should also be physically protected to restrict access to authorized individuals. If you cannot guarantee that an unlabeled host is physically secure from tampering, it and similar hosts should be isolated on a separate branch of the network.

If the security template is modified while the network is up, the changes do not take effect immediately unless `tnctl(1M)` is used to update the template entries. Otherwise, the changes take effect when next polled by the trusted network daemon, `tnd(1M)`. Administrators are allowed to add new templates and modify attributes of existing templates while the network is up.

引用名	smtntzonecfg – manage entries in the zone configuration database for Trusted Extensions networking								
用法概要	<code>/usr/sadm/bin/smtntzonecfg subcommand [auth_args] -- [subcommand_args]</code>								
描述	The <code>smtntzonecfg</code> command adds, modifies, deletes, and lists entries in the <code>tnzonecfg</code> database. <code>smtntzonecfg</code> <i>subcommands</i> are: <table><tr><td>add</td><td>Adds a new entry to the <code>tnzonecfg</code> database. To add an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.</td></tr><tr><td>modify</td><td>Modifies an entry in the <code>tnzonecfg</code> database. To modify an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.</td></tr><tr><td>delete</td><td>Deletes an entry from the <code>tnzonecfg</code> database. To delete an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.</td></tr><tr><td>list</td><td>Lists entries in the <code>tnzonecfg</code> database. To list an entry, the administrator must have the <code>solaris.network.host.read</code> and <code>solaris.network.security.read</code> authorizations.</td></tr></table>	add	Adds a new entry to the <code>tnzonecfg</code> database. To add an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.	modify	Modifies an entry in the <code>tnzonecfg</code> database. To modify an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.	delete	Deletes an entry from the <code>tnzonecfg</code> database. To delete an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.	list	Lists entries in the <code>tnzonecfg</code> database. To list an entry, the administrator must have the <code>solaris.network.host.read</code> and <code>solaris.network.security.read</code> authorizations.
add	Adds a new entry to the <code>tnzonecfg</code> database. To add an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.								
modify	Modifies an entry in the <code>tnzonecfg</code> database. To modify an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.								
delete	Deletes an entry from the <code>tnzonecfg</code> database. To delete an entry, the administrator must have the <code>solaris.network.host.write</code> and <code>solaris.network.security.write</code> authorizations.								
list	Lists entries in the <code>tnzonecfg</code> database. To list an entry, the administrator must have the <code>solaris.network.host.read</code> and <code>solaris.network.security.read</code> authorizations.								
选项	The <code>smtntzonecfg</code> authentication arguments, <i>auth_args</i>, are derived from the <code>smc</code> argument set and are the same regardless of which subcommand you use. The subcommand-specific options, <i>subcommand_args</i>, must be <i>preceded</i> by the <code>--</code> option.								
auth_args	The valid <i>auth_args</i> are <code>-D</code>, <code>-H</code>, <code>-l</code>, <code>-p</code>, <code>-r</code>, and <code>-u</code>; they are all optional. If no <i>auth_args</i> are specified, certain defaults will be assumed and the user can be prompted for additional information, such as a password for authentication purposes. These letter options can also be specified by their equivalent option words preceded by a double dash. For example, you can use either <code>-D</code> or <code>--domain</code>. <code>-D --domain domain</code> Specifies the default domain that you want to manage. The syntax of <i>domain=type:/host_name/domain_name</i>, where <i>type</i> is <code>dns</code>, <code>ldap</code>, or <code>file</code>; <i>host_name</i> is the name of the server; and <i>domain_name</i> is the name of the domain you want to manage. If you do not specify this option, the system assumes the <code>file</code> default domain on whatever server you choose to manage, meaning that changes are local to the server. Toolboxes can change the domain on a tool-by-tool basis. This option specifies the domain for all other tools.								

-H | --hostname *host_name:port*

Specifies the *host_name* and *port* to which you want to connect. If you do not specify a *port*, the system connects to the default port, 898. If you do not specify *host_name:port*, the system connects to the local host on port 898.

-l | --rolepassword *role_password*

Specifies the password for the *role_name*. If you specify a *role_name* but do not specify a *role_password*, the system prompts you to supply a *role_password*. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-p | --password *password*

Specifies the password for the *user_name*. If you do not specify a password, the system prompts you for one. Passwords specified on the command line can be seen by any user on the system, hence this option is considered insecure.

-r | --rolename *role_name*

Specifies a role name for authentication. If you do not specify this option, no role is assumed.

-u | --username *user_name*

Specifies the user name for authentication. If you do not specify this option, the user identity running the console process is assumed.

--

This option is required and must always follow the preceding options. If you do not enter the preceding options, you must still enter the -- option.

subcommand_args Descriptions and other argument options that contain white spaces must be enclosed in double quotes.

-h

Displays the command's usage statement.

-n *zonename*

Specifies the zone name for the entry. This name is used when the zone is configured. See [zonecfg\(1M\)](#), under the -z zonename option, for the constraints on zone names. The specified zone name must be one of the configured zones on the system. The following command returns a list of configured zones:

```
/usr/sbin/zoneadm list -c
```

-l *label*

Specifies the label for the zone. This field is used to label the zone when the zone is booted. Each zone must have a unique label.

-x *polycymatch=0|1*

Specifies the policy match level for non-transport traffic. Only values of 0 (match the label) or 1 (be within the label range of the zone) are accepted.

ICMP packets that are received on the global zone IP address are accepted based on the label range of the global zone's security template if the global zone's *polycymatch* field is set to 1. When this field is set to 0 for a zone, the zone will not respond to an ICMP echo request from a host with a different label.

This subcommand argument is optional. If not specified, it will have a default value of 0.

-x *mlpzone*=""*port/protocol*

Specifies the multilevel port configuration entry for zone-specific IP addresses. Multiple *port/protocol* combinations are separated by a semi-colon. The empty string can be specified to remove all existing MLP zone values. This subcommand argument is optional.

An MLP is used to provide multilevel service in the global zone as well as in non-global zones. As an example of how a non-global zone can use an MLP, consider setting up two labeled zones, *internal* and *public*. The *internal* zone can access company networks; the *public* zone can access public internet but not the company's internal networks. For safe browsing, when a user in the *internal* zone wants to browse the Internet, the *internal* zone browser forwards the URL to the *public* zone, and the web content is then displayed in a *public* zone web browser. That way, if the download in *public* zone compromises the web browser, it cannot affect the company's internal network. To set this up, TCP port 8080 in the *public* zone is an MLP (8080/tcp), and the security template for the *public* zone has a label range from *PUBLIC* to *INTERNAL*.

-x *mlpshared*=""*port/protocol*

Specifies the multilevel port configuration entry for shared IP addresses. Multiple *port/protocol* combinations are separated by a semi-colon. The empty string can be specified to remove all existing MLP shared values. This subcommand argument is optional.

A shared IP address can reduce the total number of IP addresses that are needed on the system, especially when configuring a large number of zones. Unlike the case of the zone-specific IP address, when MLPs are declared on shared IP addresses, only the global zone can receive the incoming network traffic that is destined for the MLP.

- One of the following sets of arguments must be specified for subcommand *add*:

```
-n zonename -l label [-x polycymatch=policy-match-level \
-x mlpzone=port/protocol;... | \
-x mlpshared=port/protocol;... ]
-h
```

- One of the following sets of arguments must be specified for subcommand *modify*:

```
-n zonename [-l label] [-x polycymatch=policy-match-level \
-x mlpzone=port/protocol;... | \
-x mlpshared=port/protocol;... ]
-h
```

- One of the following arguments must be specified for subcommand *delete*:

-n *zonename* |
-h

- The following argument can be specified for subcommand `list`:

-n *zonename* |
-h

示例

示例 1 Adding a New Entry to the Zone Configuration Database

The admin role creates a new zone entry, `public`, with a label of `public`, a policy match level of 1, and a shared MLP port and protocol of 666 and TCP. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smtntzonecfg add -- -n public -l public \  
-x policymatch=1 -x mlpshared=666/tcp
```

示例 2 Modifying an Entry in the Zone Configuration Database

The admin role changes the `public` entry in the `tnzonecfg` database to `needtoknow`. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smtntzonecfg modify -- -n public -l needtoknow
```

示例 3 Listing the Zone Configuration Database

The admin role lists the entries in the `tnzonecfg` database. The administrator is prompted for the admin password.

```
$ /usr/sadm/bin/smtntzonecfg list --
```

退出状态

The following exit values are returned:

- 0 Successful completion.
- 1 Invalid command syntax. A usage message displays.
- 2 An error occurred while executing the command. An error message displays.

文件

The following files are used by the `smtntzonecfg` command:

`/etc/security/tsol/tnzonecfg` Trusted zone configuration database.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmgts
Interface Stability	Committed

另请参见

[attributes\(5\)](#)

附注

The functionality described on this manual page is available only if the system is configured with Trusted Extensions.

引用名	smtp-notify – 用于软件事件的电子邮件通知守护进程
用法概要	/usr/lib/fm/notify/smtp-notify
描述	smtp-notify 是用于订阅软件事件和故障管理生命周期事件并根据存储在 SMF 服务配置系统信息库中的一组通知首选项生成电子邮件通知的守护进程。 电子邮件通知首选项是使用 svccfg(1M) 设置的。 smtp-notify 由服务管理工具 smf(5) 管理，其服务 FMRI 为： <pre>svc:/system/fm/smtp-notify:default</pre> 可以使用 svcadm(1M) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 svcs(1) 命令来查询服务的状态。
属性	可以设置以下服务属性： <pre>config/rootdir</pre> 这是缺省值为 / 的一个 <code>astring</code> 属性。设置了此属性时，将为 smtp-notify 评估的所有路径名使用指定的根目录。 缺省情况下，电子邮件事件通知的正文将基于随 Oracle Solaris 操作系统提供的一组本地化消息模板。smtp-notify 将查找相应模板（根据事件类）并使用事件有效负荷中的元素填充模板。 通过指定 <code>msg_template</code> 通知首选项可以替换用户提供的消息正文模板。有关针对每个服务或每个事件类设置通知首选项的更多信息，请参阅 svccfg(1M)。 消息正文模板可以包含下面列出的扩展宏。在发送消息之前，将使用 smtp-notify 扩展这些宏。

Macro	Description
-----	-----
%%	expands to a literal % character
%<HOSTNAME>	expands to the hostname on which the event occurred
%<URL>	expands to the URL of the knowledge article associated with this event
%<CLASS>	expands to the event class
%<UUID>	expands to the UUID of the event
%<CODE>	expands to the knowledge article message ID
%<SEVERITY>	expands to the severity of the event

对于 SMF 服务状态转换事件，可以在消息模板中指定以下附加宏：

%<FMRI>	expands to FMRI of the affected service
%<FROM-STATE>	expands to the previous state of the service
%<TO-STATE>	expands to the new state of the service

为方便进行电子邮件过滤，smtp-notify 将为所有事件类的事件有效负荷的以下组件创建 X 头：

Header name	Value
-----	-----
X-FMEV-HOSTNAME	the name of the host on which the event occurred
X-FMEV-CLASS	the event class
X-FMEV-CODE	the Knowledge article message ID
X-FMEV-SEVERITY	the severity of the event
X-FMEV-UUID	the UUID of the event

SMF 服务状态转换事件也将包含下面的附加 X 头：

Header name	Value
-----	-----
X-FMEV-FMRI	the FMRI of the affected SMF service
X-FMEV-FROM-STATE	the previous state of the service
X-FMEV-TO-STATE	the new state of the service

示例

示例1 配置通知首选项

以下命令为 SMF 服务状态转换事件配置通知首选项。

```
# svccfg -s svc:/system/svc/global:default setnotify -g \
from-online,to-maintenance mailto:admin@somehost.com
```

示例2 启用电子邮件通知

以下命令为故障管理问题诊断事件启用电子邮件通知。

```
# svccfg setnotify problem-diagnosed mailto:admin@somehost.com
```

示例3 列出通知设置

以下命令列出故障管理问题诊断事件的所有通知设置。

```
# svccfg listnotify problem-diagnosed
```

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/fault-management/smtp-notify
接口稳定性	Unstable（不稳定）

另请参见

[svcs\(1\)](#)、[svcadm\(1M\)](#)、[svccfg\(1M\)](#)、[syslogd\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

引用名

sndradm – 控制 Sun StorageTek Availability Suite 远程镜像操作

用法概要

```
sndradm -I a master shadow bitmap
sndradm -I d master shadow bitmap
sndradm -h usage message
sndradm -v version information
sndradm [options] -e [sndr_set]
sndradm [options] -E [sndr_set]
sndradm [options] -d [sndr_set]
sndradm [options] -D block [sndr_set]
sndradm [options] -D noblock [sndr_set]
sndradm [options] -l [sndr_set]
sndradm [options] -m [sndr_set]
sndradm [options] -m -r [sndr_set]
sndradm [options] -u [sndr_set]
sndradm [options] -u -r [sndr_set]
sndradm [options] -w [sndr_set]
sndradm [options] -H [sndr_set]
sndradm [options] -p [sndr_set]
sndradm [options] -P [sndr_set]
sndradm [options] -q a volume [sndr_set]
sndradm [options] -q d [sndr_set]
sndradm [options] -q r volume [sndr_set]
sndradm [options] -i [sndr_set]
sndradm [options] -a value [sndr_set]
sndradm [options] -A value [sndr_set]
sndradm [options] -F value [sndr_set]
sndradm [options] -W value [sndr_set]
sndradm [options] -R
sndradm [options] -R b p bitmap [sndr_set]
sndradm [options] -R b s bitmap [sndr_set]
sndradm [options] -R C tag [sndr_set]
```



```

sndradm [options] -R g io_groupname [sndr_set]
sndradm [options] -R m sync [sndr_set]
sndradm [options] -R m async [sndr_set]
sndradm [options] -R -f volset-file
sndradm [options] -R r [sndr_set] *

```

描述

`/usr/sbin/sndradm` 命令是用于 Sun StorageTek Availability Suite 远程镜像软件的管理命令行接口。使用远程镜像，可以在物理分隔的不同 Sun 服务器之间实时复制磁盘。远程镜像在概念上与 RAID 1 的本地磁盘镜像方案相似，但它在较远距离执行复制操作。

如果未在命令行中指定远程镜像集 (*sndr_set*)，`sndradm` 基于配置的所有远程镜像集运行。

`sndradm` 命令在 Availability Suite 日志文件 `/var/adm/ds.log`（请参见 [ds.log\(4\)](#)）中为除输出 (`-p`、`-P` 和 `-i`)、帮助 (`-h`) 和版本 (`-v`) 以外的所有操作生成一个项。

选项

`sndradm` 实用程序支持下列选项：

`-f volset-file`

指定包含一个或多个远程镜像集的 *sndr_set* 信息的文件，该文件的格式与下面记录的完全指定命令行 *sndr_set* 的格式相同。

`-g io_groupname`

将操作仅限制为 *io_groupname* 包含的远程镜像集。

给定集合的 *io_groupname* 必须在主要主机和辅助主机上保持一致。

`-C tag`

在群集节点中，将操作仅限制为 *tag* 指定的群集资源组或磁盘组名称所包含的远程镜像集。此选项在非群集系统上无效。

`-n`

使用 `sndradm` 启动远程镜像操作后不会提示用户。对于输出、帮助和版本选项以外的所有选项，缺省行为是提示进行响应。例如，在启动从主卷到辅助卷的完全同步之后，远程镜像提示："Overwrite secondary with primary? (Y/N) [N]"。

sndr_set

指定远程镜像集。对于已启用的集合，此选项可以是 *shost:sdev* 格式的 *set_name*。可以提供一个格式与配置文件相同的完全指定的远程镜像集：

```
phost pdev pbitmap shost sdev sbitmap ip {sync | async} \
[g io_groupname] [C tag]
```

下面介绍了这些参数：

phost

指定主卷所在的服务器。

pdev

指定要复制的主卷分区。仅指定全路径名（例如，`/dev/rdisk/c0t1d0s2`）。

pbitmap

指定存储主分区的位图（记分板日志）的卷分区。仅指定全路径名（例如，`/dev/rdisk/c0t1d0s3`）。

shost

指定辅助卷所在的服务器。

sdev

指定辅助卷分区。仅指定全路径名（例如，`/dev/rdisk/c0t1d0s4`）。

sbitmap

指定存储辅助分区的位图（记分板日志）的卷分区。仅指定全路径名（例如，`/dev/rdisk/c0t1d0s5`）。

ip

指定网络传输协议。

sync | async

指定远程镜像运行模式。*sync* 是远程镜像模式，在此模式下，I/O 操作不会确认为已完成，直到已更新远程卷为止。*async* 是远程镜像模式，在此模式下，主要主机 I/O 操作确认为已完成，然后再更新远程卷。

io_groupname

指定远程镜像集所属的远程镜像一致性组的名称。在异步模式下，必须在远程镜像一致性组中的所有复制卷中保存写入顺序。这可以确保该组包含的辅助卷包含相应主卷的有效时间点副本。

向一致性组添加现有集合或在组中添加某个集合时，必须使用相同组名在主要主机和辅助主机上配置此集合。

tag (标记)

对于群集内的操作，如果卷名未暗含本地数据的磁盘组名称或资源标记以及位图卷，此选项会指定这些内容（例如，`/dev/rdisk/md/dg/vol` 和 `/dev/vx/rdisk/dg/vol` 均指示磁盘组名称 *dg*）。用户负责确保为远程镜像指定的群集标记与相应群集资源组标记匹配，并确保在重新配置群集资源组时所有 Availability Suite 服务都是最新的。在非群集系统上指定群集资源标记是非法的。

参数

有效 *sndradm* 命令必须指定下面列出的某个参数。

-I a master shadow bitmap

向 Availability Suite 配置文件添加具有指定主卷、影子卷和位图卷的 *ndr_ii* 项。请参见 [sndrsyncd\(1M\)](#)。如果相应 Point-in-Time Copy 集合不存在，则会在相关卷上执行下一 *sync* 命令时启用此集合。如果不再需要此集合，使用 *iiadm -d* 可禁用此 Point-in-Time Copy 集合。请参见 [iiadm\(1M\)](#)

-I *d master shadow bitmap*

从 Availability Suite 配置文件中删除具有指定主卷、影子卷和位图卷的 `ndr_ii` 项。使用 `dscfg` 命令列出 `ndr_ii` 配置项。

-a *value*

指定此集合的自动同步变量的值（打开或关闭）。针对某个集合请求 `autosync` 之后，从请求同步操作之时起，此功能处于活动状态，直到手动将此集合置于日志记录模式为止。手动将此集合置于日志记录模式后，`autosync` 功能将处于非活动状态，并在下次执行 `sync` 请求之前一直保持非活动状态。要检查 `autosync` 是否处于活动状态，请使用 `sndradm -P`。要检查是否已针对某个集合请求 `autosync`，请在 `dscfg -l` 输出中查找此集合的“`auto=on;`”标记。请参见 [sndrsyncd\(1M\)](#)。

-A *value*

指定在异步模式下复制集时为处理异步队列可创建的最大线程数。缺省值为 2。

-W *value*

指定可在异步模式下复制的集合中排队的最大写入数。缺省值为 4096。例如，将此值设置为 1 可确保辅助卷包含的写入操作数比主卷包含的写入操作数最多小 1。

调整最大写入数仅对使用基于内存的异步 I/O 队列有效。当使用基于磁盘的 I/O 队列时，将忽略此值。

-D {*block | noblock*}

切换基于磁盘的队列的 `block|noblock` 属性。缺省设置为 `block`。如果 I/O 填充速率远远大于清除速率以致无法填充此队列，则会阻止传入 I/O 直到队列中存在用于此 I/O 的充足空间为止。这是为了保存写入顺序，而不管是在一个卷还是在同一一致性组中的多个卷内。如果设置了 `noblock`，传入 I/O 将填充此队列，并且不会阻止 I/O。此集合将置于日志记录模式，并忽略磁盘队列内容。后继更新同步会将最新数据同步到辅助站点。

-F *value*

指定可在异步模式下复制的集合中排队的内核内存中的 512 字节 FBA 的最大数目。缺省值为 16384。

调整最大 FBA 数仅对使用基于内存的异步 I/O 队列的集合有效。使用基于磁盘的 I/O 队列时，将忽略此值。

-h

输出 `sndradm` 用法汇总。

-v

输出远程镜像版本号。

-e

为集合启用远程镜像，并启用记分板日志记录。设置记分板以便指示需要完全同步。集合详细信息保存在当前配置中。请参见 [dscfg\(1M\)](#)。针对存储卷驱动程序启用逻辑卷和位图卷（请参见 [sv\(7D\)](#)）。

-M

为集合启用远程镜像，并启用记分板日志记录。清除记分板以便指示已确保完全同步主卷和辅助卷。集合详细信息保存在当前配置中。请参见 [dscfg\(1M\)](#)。针对存储卷驱动程序启用逻辑卷和位图卷（请参见 [sv\(7D\)](#)）。

-d

为集合禁用远程镜像，并暂停任何当前的同步操作。`sndradm -d` 还会废弃用于跟踪主卷和辅助卷之间的临时差异的任何活动记分板。

-l

在主卷和辅助卷之间停止远程镜像复制和复制操作，并在这些卷上启动单独的远程镜像记分卡日志记录。复制一致性组中的所有集合时，这意味着辅助卷包含相应主卷的有效时间点副本。此种情况下，只要远程镜像集进入日志记录模式，`rdc` 内核模块会使此组中的所有其他集合自动进入日志记录模式。这确保辅助卷仍包含有效的时间点副本。要在使用 `-l` 参数后恢复远程镜像，请使用 `-m` 参数执行完全重新同步，或者使用 `-u` 参数执行更新重新同步（根据记分板）。

此选项不会在当前正在同步的任何卷的辅助卷上运行。

-w

等待同步复制完成或中止，或者立即返回（如果在辅助系统上调用）。

-H

报告指定卷集使用的网络链路的运行状况。此链路的运行状况报告为活动或非活动。活动表示此网络链路当前正积极用于复制或重新同步数据，因此运行状况良好。非活动表示网络链路当前未积极用于复制或重新同步数据，这可能表示此链路存在问题。

-p

显示配置的远程镜像卷或集合的列表。

-P

显示配置的远程镜像卷或集合的列表以及额外详细信息。（请参见下面的状态说明。）

-q a *volume*

向集合或组添加磁盘队列。集合或组处于日志记录模式时，此操作有效。

-q d

从集合或组中删除磁盘队列。集合或组处于日志记录模式时，此操作有效。

-q r *volume*

替换组或集合的磁盘队列。按照在队列禁用操作中的方式从集合或组中删除队列，并且按照在队列添加操作中的方式添加新磁盘队列。集合或组处于日志记录模式时，此操作有效。

-i

采用与 `volset-file` 相同的格式显示配置的远程镜像卷或集合的列表。

- R
尝试重置远程镜像集的错误条件，例如，故障位图。
- R **b p bitmap**
重新配置远程镜像集的主要位图。应在主要服务器和辅助服务器上输入此命令。一次只能为一个集合重新配置主要位图。
- R **b s bitmap**
重新配置远程镜像集的辅助位图。应在主要服务器和辅助服务器上输入此命令。一次只能为一个集合重新配置辅助位图。
- R **C tag**
重新配置远程镜像集的逻辑卷的群集标记或磁盘组名称（如果路径名未指示这些内容）。这不会影响远程卷。无法在非群集系统上使用此参数。
- R **m {sync | async}**
重新配置远程镜像集的复制模式。一致性组包含的集合必须全部处于同步或异步状态。无法在一个组中混合多种模式。
- R **g group**
重新配置远程镜像集的一致性组。应使用相同组名在主要服务器和辅助服务器上输入此命令。

要从一致性组中删除某个集合，请在重新配置一致性组时指定空字符串 ("")。

只能从主要服务器上发出下列参数：

- m
启动从主卷到辅助卷的完整卷复制，同时对新更新启用从主卷到辅助卷的远程镜像复制。当主卷和辅助卷可能互不相同，并且不存在以增量方式重新同步卷的日志记录信息时，请使用此参数。请参见“退出状态”。
- r
颠倒同步方向，以便从辅助卷同步主卷。将此参数与 -m 或 -u 参数配合使用。-m -r 启动从辅助（源）卷到主（目标）卷的完整卷复制，但同时启用从主（源）卷到辅助（目标）卷的新更新的远程镜像复制，从而确保卷集保持同步。当主卷和辅助卷内容可能互不相同，并且辅助卷具有所需内容，但不存在以增量方式重新同步卷（使用 -u）的日志记录信息时，请使用 -m -r。-u -r 使用在暂停复制时维护的远程镜像记分板日志从辅助（源）卷重新同步主（目标）卷。然后，恢复从主卷到辅助卷的新更新的远程镜像复制，以使卷集保持同步。在恢复/刷新操作期间停止卷集的工作负荷。此操作确保在恢复新更新复制之前匹配主卷和辅助卷。
- u
更新远程镜像卷集。此参数重新同步远程镜像卷集。仅更新远程镜像记分板中记录为已发生更改的块。启用对主卷的远程镜像复制，并使用远程镜像记分板日志启动重新同步过程，以使相应辅助卷与主卷匹配。

sndradm -P 返回的状态 下面列出了 sndradm -P 可返回的状态。

volume failed

逻辑数据卷的 I/O 操作已失败

bitmap failed

逻辑位图卷的 I/O 操作已失败

disk queue failed

磁盘队列卷的 I/O 操作已失败

need sync

已中断到此卷的同步。需要完成此同步（或通过 Point-in-Time Copy 恢复）。不能更改数据流方向，直到已完成一个或其他数据流为止。

need reverse sync

已中断到此卷的反向同步。需要完成此同步（或通过 Point-in-Time Copy 恢复）。不能更改数据流方向，直到已完成一个或其他数据流为止。

logging

仅在位图中记录传入写入。不会将数据复制到远程站点。need sync、need reverse sync 和 queuing 均为日志记录的子状态，以便在位图中记录（而不是复制）写入。队列模式（如下所述）将写入记录到位图，并对请求进行排队，以便在稍后由异步刷新程序复制。

reverse syncing

正在进行从辅助卷到主卷的复制。

syncing

正在进行从主卷到辅助卷的复制。

queuing

使用磁盘队列进行常规异步复制时，i/o 将放置到磁盘队列中，以便由异步刷新程序线程复制。如果发生临时链路故障，集合将切换到队列模式。不会废弃队列，因为它将与基于内存的队列一起运行。数据将记录到位图中，并放置在队列中。启用链路并发出 sndradm -u 时（通过为集合启用 autosync 来自动执行此操作），将重新启动刷新程序。这会在整个临时链路中断过程中保留写入顺序。如果不需要使用写入顺序，并且仅需要最新数据，则可将集合手动置于日志记录模式 (sndradm -l)，并发出更新同步 (sndradm -u)。此操作废弃队列中的数据，并使用位图进行快速重新同步。如果在链路恢复并发出更新同步之前填充队列，则会废弃此队列，并将集合置于日志记录模式以免应用程序挂起。

示例

示例1 启用远程镜像集

下面的命令在主机 example 上启用远程镜像异步集合，其中 example 是主要主机，example-remote 是辅助主机。

```
example% sndradm -e example /dev/rdisk/c1t0d0s1 /dev/rdisk/c1t1d0s3 \
example-remote /dev/rdisk/c2t3d0s5 /dev/rdisk/c2t4d0s5 ip async
```

示例2 向异步集合添加磁盘队列

下面的命令向异步集合添加磁盘队列卷。

```
example% sndradm -q a /dev/rdisk/c1t2d0s3 \
example-remote:/dev/rdisk/c2t3d0s5
```

示例3 从异步集合中删除磁盘队列

下面的命令从附加有磁盘队列卷的集合中删除磁盘队列卷。

```
example% sndradm -q d example-remote:/dev/rdisk/c2t3d0s5
```

示例4 禁用远程镜像集

下面的命令禁用在主机 example 上启用的远程镜像集。

```
example% sndradm -d example-remote:/dev/rdisk/c2t3d0s5
```

退出状态

0 命令已成功完成。

>0 出现错误。

在脚本中执行 -m 或 -u 选项时，这些选项后面的某个选项的退出状态始终返回成功，而不管远程镜像集的当前状态如何。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	storage/avs/avs-remote-mirror
接口稳定性	Committed（已确定）

另请参见 [dscfg\(1M\)](#)、[sndrd\(1M\)](#)、[sndrsyncd\(1M\)](#)、[ds.log\(4\)](#)、[rdc.cf\(4\)](#)、[attributes\(5\)](#)、[sv\(7D\)](#)

引用名	sndrd – Remote Mirror 守护进程
用法概要	<code>/usr/lib/sndrd [-c max_connections] [-l listen_backlog]</code>
描述	sndrd 守护进程处理客户机远程镜像请求。只有超级用户或具有同等特权的用户可以运行此守护进程。该守护进程是以运行级 2 自动调用的。sndrd 重新启动 TCP 传输层。 要更改 sndrd 的启动参数，管理员必须以 root 身份或等效身份更改 <code>/etc/default/sndr</code> 文件，而不是编辑 <code>/lib/svc/method/svc-rdcsyncd</code> 文件。请参见 sndr(4)。
选项	sndrd 守护进程支持以下选项： <code>-c max_connections</code> 设置面向连接的传输上的服务器允许的最大连接数。缺省情况下，连接数是 16。 <code>-l listen_backlog</code> 为面向连接的传输上的 RDC TCP 设置连接队列长度。缺省值为 10 个条目。
退出状态	0 守护进程启动成功。 >0 守护进程启动失败。 错误信息以 LOG_ERR 级别报告到 syslog 中。
文件	<code>/lib/svc/method/svc-rdcsyncd</code> 用于启动 sndrd 的 shell 脚本。 <code>/lib/svc/method/svc-rdc</code> 用于停止 sndrd 的 shell 脚本。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	storage/avs/avs-remote-mirror
接口稳定性	Committed（已确定）

另请参见 [svcadm\(1M\)](#)、[syslogd\(1M\)](#)、[ds.log\(4\)](#)、[attributes\(5\)](#)

附注

不要手动停止 sndrd 守护进程。如果需要手动停止 sndrd，请执行以下步骤。这将同时停止 sndrd 和 sndrsyncd 守护进程。

```
# svcadm disable svc:/system/nws_rdc
# svcadm disable svc:/system/nws_rdcsyncd
```

不要手动启动或重新启动 sndrd 守护进程。如果需要手动启动 sndrd，请执行以下步骤。这将同时启动 sndrd 和 sndrsyncd 守护进程。

```
# svcadm enable svc:/system/nws_rdc
```



```
# svcadm enable svc:/system/nws_rdcsyncd
```

有关其他信息，请参见 [svcadm\(1M\)](#)。

引用名	sndrsyncd – Availability Suite Remote Mirror 更新重新同步守护进程						
用法概要	/usr/lib/sndrsyncd						
描述	sndrsyncd 守护进程在网络或计算机发生故障后自动执行更新重新同步，并且在需要时调用 Point-in-Time Copy 副本以保护在重新同步过程中更新的数据卷。 当 Remote Mirror 使用的网络链路出现故障时，内核将通知该守护进程，该守护进程将使用 -u 选项调用 sndradm(1M) 命令来重新同步启用了自动同步功能且正在使用该网络链路的所有 Remote Mirror 集。有关如何为 Remote Mirror 集配置自动同步功能的详细信息，请参见 sndradm(1M)。 当有任何 Remote Mirror 重新同步启动或完成时，也会通知该守护进程。如果在 Availability Suite 配置文件中配置了守护进程，该守护进程将创建 Availability Suite 快照。在辅助服务器上，守护进程检查文件系统当前是否已挂载到辅助卷上，如果文件系统当前已挂载，则会通知内核不允许启动同步。						
退出状态	0 守护进程启动成功。 >0 守护进程启动失败。						
文件	/lib/svc/method/svc-rdcsyncd 用于启动 sndrsyncd 的 shell 脚本。 /lib/svc/method/svc-rdc 用于停止 sndrsyncd 的 shell 脚本。						
属性	有关以下属性的说明，请参见 attributes(5) ：						
<table><tr><th>属性类型</th><th>属性值</th></tr><tr><td>可用性</td><td>storage/avs/avs-remote-mirror</td></tr><tr><td>接口稳定性</td><td>Committed（已确定）</td></tr></table>		属性类型	属性值	可用性	storage/avs/avs-remote-mirror	接口稳定性	Committed（已确定）
属性类型	属性值						
可用性	storage/avs/avs-remote-mirror						
接口稳定性	Committed（已确定）						
另请参见	iiadm(1M) 、 sndradm(1M) 、 sndrd(1M) 、 svcadm(1M) 、 ds.log(4) 、 rdc.cf(4) 、 attributes(5)						
附注	不要手动停止 sndrsyncd 守护进程。如果需要手动停止 sndrsyncd，请执行以下步骤。这将同时停止 sndrd 和 sndrsyncd 守护进程。 <pre># svcadm disable svc:/system/nws_rdc # svcadm disable svc:/system/nws_rdcsyncd</pre> 不要手动启动或重新启动 sndrsyncd 守护进程。如果需要手动启动 sndrsyncd，请执行以下步骤。这将同时启动 sndrd 和 sndrsyncd 守护进程。 <pre># svcadm enable svc:/system/nws_rdc # svcadm enable svc:/system/nws_rdcsyncd</pre> 有关其他信息，请参见 svcadm(1M)。						

引用名	snmpdx – Sun Solstice Enterprise Master Agent														
用法概要	<pre>/usr/lib/snmp/snmpdx [-hy] [-a filename] [-c config-dir] [-d debug-level] [-i filename] [-m GROUP -m SPLIT] [-o filename] [-p port] [-r filename]</pre>														
描述	The Master Agent, snmpdx, is the main component of Solstice Enterprise Agent (SEA) technology. It runs as a daemon process and listens to User Datagram Protocol (UDP) port 161 for SNMP requests. The Master Agent also opens another port to receive SNMP trap notifications from various subagents. These traps are forwarded to various managers, as determined by the configuration file. Upon invocation, snmpdx reads its various configuration files and takes appropriate actions by activating subagents, determining the subtree Object Identifier (OID) for various subagents, populating its own Management Information Bases (MIBs), and so forth. The Master Agent invokes subagents, registers subagents, sends requests to subagents, receives responses from subagents, and traps notifications from subagents. The Master Agent is invoked by the service management facility smf(5) at boot time if <code>svc:/application/management/snmpdx</code> is enabled (see NOTES) and contents of the resource configuration file <code>/etc/snmp/conf/snmpdx.rsrc</code> are non-trivial. 注 – The SMA (Systems Management Agent) is the default SNMP agent in the Solaris operating system. snmpdx is Obsolete and may not be supported in a future release of Solaris.														
选项	The following options are supported: <table><tr><td>-a <i>filename</i></td><td>Specify the full path of the access control file used by the Master Agent. The default access control file is <code>/etc/snmp/conf/snmpdx.ac1</code>.</td></tr><tr><td>-c <i>config-dir</i></td><td>Specify the full path of the directory containing the Master Agent configuration files. The default directory is <code>/etc/snmp/conf</code>.</td></tr><tr><td>-d <i>debug-level</i></td><td>Debug. Levels from 0 to 4 are supported, giving various levels of debug information. The default is 0 which means no debug information is given.</td></tr><tr><td>-h</td><td>Help. Print the command line usage.</td></tr><tr><td>-i <i>filename</i></td><td>Specify the full path of the enterprise-name OID map. This file contains the PID used by the Master Agent for recovery after a crash. It contains tuples of the UNIX process ID, port number, resource name, and agent name. The default file is <code>/var/snmp/snmpdx.st</code>.</td></tr><tr><td>-m GROUP -m SPLIT</td><td>Specify the mode to use for forwarding of SNMP requests.</td></tr></table> <table><tr><td>GROUP</td><td>Multiple variables can be included in each request from the Master Agent to the subagents. This results in, at most, one send-request per agent.</td></tr></table>	-a <i>filename</i>	Specify the full path of the access control file used by the Master Agent. The default access control file is <code>/etc/snmp/conf/snmpdx.ac1</code> .	-c <i>config-dir</i>	Specify the full path of the directory containing the Master Agent configuration files. The default directory is <code>/etc/snmp/conf</code> .	-d <i>debug-level</i>	Debug. Levels from 0 to 4 are supported, giving various levels of debug information. The default is 0 which means no debug information is given.	-h	Help. Print the command line usage.	-i <i>filename</i>	Specify the full path of the enterprise-name OID map. This file contains the PID used by the Master Agent for recovery after a crash. It contains tuples of the UNIX process ID, port number, resource name, and agent name. The default file is <code>/var/snmp/snmpdx.st</code> .	-m GROUP -m SPLIT	Specify the mode to use for forwarding of SNMP requests.	GROUP	Multiple variables can be included in each request from the Master Agent to the subagents. This results in, at most, one send-request per agent.
-a <i>filename</i>	Specify the full path of the access control file used by the Master Agent. The default access control file is <code>/etc/snmp/conf/snmpdx.ac1</code> .														
-c <i>config-dir</i>	Specify the full path of the directory containing the Master Agent configuration files. The default directory is <code>/etc/snmp/conf</code> .														
-d <i>debug-level</i>	Debug. Levels from 0 to 4 are supported, giving various levels of debug information. The default is 0 which means no debug information is given.														
-h	Help. Print the command line usage.														
-i <i>filename</i>	Specify the full path of the enterprise-name OID map. This file contains the PID used by the Master Agent for recovery after a crash. It contains tuples of the UNIX process ID, port number, resource name, and agent name. The default file is <code>/var/snmp/snmpdx.st</code> .														
-m GROUP -m SPLIT	Specify the mode to use for forwarding of SNMP requests.														
GROUP	Multiple variables can be included in each request from the Master Agent to the subagents. This results in, at most, one send-request per agent.														

SPLIT Each variable in the incoming request results in one send-request to each subagent.

The default is GROUP.

- o*filename* Specify the full path of the file containing the tuple (enterprise-name, OID). For example, (Sun Microsystems, 1.3.1.6.1.4.32). The Master Agent uses this file as a base for look-up in the trap-filtering and forwarding process. The default file is /etc/snmp/conf/enterprises.oid.
- p*port* Specify the port number. The default port number is 161.
- r*filename* Specify the full path of the resource file to be used by the Master Agent. This file stores information about the subagents that the Master Agent invokes and manages. The default resource file is /etc/snmp/conf/snmpdx.rsrc.
- y Set a recovery indicator to invoke the recovery module. The recovery process discovers which subagents in the previous session are still active; those subagents not active are re-spawned by the Master Agent.

文件	/etc/snmp/conf/enterprises.oid	Enterprise-name OID map
	/etc/snmp/conf/snmpdx.acl	Access control file
	/etc/snmp/conf/snmpdx.rsrc	Resource configuration file
	/var/snmp/snmpdx.st	Master Agent status file
	/var/snmp/mib/snmpdx.mib	Master Agent MIB file

退出状态 The following error values are returned:

- 0 Successful completion.
- non-zero An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/management/snmp/sea
Interface Stability	Obsolete

另请参见 [attributes\(5\)](#), [smf\(5\)](#)

附注

The snmpdx service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/application/management/snmpdx
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	snmp-notify – 用于软件事件的 SNMP 通知守护进程
用法概要	/usr/lib/fm/notify/snmp-notify
描述	snmp-notify 是用于订阅软件事件及 FMA 诊断和修复事件并根据存储在 SMF 服务配置系统信息库中的一组通知首选项生成 SNMP 陷阱通知的守护进程。 SNMP 通知首选项是使用 svccfg(1M) 设置的。 snmp-notify 由服务管理工具 smf(5) 管理，其服务 FMRI 为： <pre>svc:/system/fm/snmp-notify:default</pre> 可以使用 svcadm(1M) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 svcs(1) 命令来查询服务的状态。
属性	可以设置以下服务属性： <pre>config/rootdir</pre> 这是缺省值为 / 的一个 <code>astring</code> 属性。设置了此属性时，将为 snmp-notify 评估的所有路径名使用指定的根目录。
示例	示例1 配置通知首选项 以下命令为 SMF 服务状态转换事件配置通知首选项。 <pre># svccfg setnotify -g from-online,to-maintenance snmp:active</pre> 示例2 启用 SNMP 通知 以下命令为故障管理诊断事件启用通知。 <pre># svccfg setnotify problem-diagnosed snmp:active</pre> 示例3 列出通知设置 以下命令列出故障管理诊断事件的所有通知设置。 <pre># svccfg listnotify fma-diagnosis</pre> 有关配置 SNMP 通知首选项的更多信息，请参阅 svccfg(1M)。
属性	有关以下属性的说明，请参见 attributes(5)：

属性类型	属性值
可用性	system/fault-management/snmp-notify
接口稳定性	Unstable（不稳定）

另请参见

[svcs\(1\)](#)、[svcadm\(1M\)](#)、[svccfg\(1M\)](#)、[syslogd\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

引用名	snmpXwbemd – SNMP Adapter Subagent for WBEM
用法概要	<code>/usr/sadm/lib/wbem/snmpXwbemd [-d] [-h] [-p <i>port</i>]</code>
描述	The <code>snmpXwbemd</code> daemon is a subagent in the Web-Based Enterprise Management (WBEM) services package. This daemon maps the Simple Network Management Protocol (SNMP) requests forwarded by the Solstice Enterprise Agents (SEA) Master Agent snmpdx(1M) into one or more equivalent WBEM Common Information Model (CIM) properties and instances. Further, it remaps the response from the CIM Object Manager into a SNMP response, which it passes back to <code>snmpdx(1M)</code>. A mapping file contains the corresponding Object Identifier (OID), class name, property name, and Abstract Syntax Notation 1 (ASN.1) type for each object. You can also create your own mapping file.
选项	The following options are supported: -d Displays all debug information. -h Displays help by printing the correct command line usage. -p Specifies the port number to use.
操作数	The following operand is supported: <i>port</i> Specifies the port number you want to use.
示例	示例 1 An Example of a 050SUNWwbcou.map File This mapping file that Sun Microsystems provides contains definitions of objects, in this format: <pre># #pragma ident "@(#)050SUNWwbcou.map 1.0 01/04/03 SMI" # # Copyright (c) 2001 by Sun Microsystems, Inc. # All rights reserved. # # *** Description of contents *** # # First non-commented non-blank line contains required Version label. # Remaining non-commented non-blank lines are considered map entries # used as described below: # # Column 1 - SNMP OID - Uniquely describes an SNMP variable # Column 2 - CIM Class Name - CIM class associated with this variable # Column 3 - CIM Property Name - CIM property that maps to SNMP OID variable # Column 4 - ASN.1 type - SNMP datatype that dictates how data is mapped # # to/from SNMP requests. Supported types are: SnmpString, SnmpOid,</pre>

示例 1 An Example of a 050SUNWwbcou.map File (续)

```
#           SnmpTimeticks, SnmpCounter, SnmpInt, SnmpGauge, SnmpIpAddress,
#           SnmpOpaque)
# Column 5 and greater are ignored
#
Version 1.0

1.3.6.1.2.1.1.1.0 Solaris_ComputerSystem Description SnmpString
1.3.6.1.2.1.1.3.0 Solaris_OperatingSystem LastBootUpTime SnmpTimeticks
1.3.6.1.2.1.1.4.0 Solaris_ComputerSystem PrimaryOwnerContact SnmpString
1.3.6.1.2.1.1.5.0 Solaris_ComputerSystem Name SnmpString

1.3.6.1.2.1.25.1.5.0 Solaris_OperatingSystem NumberOfUsers SnmpGauge
1.3.6.1.2.1.25.1.6.0 Solaris_OperatingSystem NumberOfProcesses SnmpGauge
1.3.6.1.2.1.25.1.7.0 Solaris_OperatingSystem MaxNumberOfProcesses SnmpGauge
1.3.6.1.2.1.25.1.2.0 Solaris_OperatingSystem LocalDateTime SnmpString
```

Each definition of an object in this file contains an OID, its corresponding CIM class name, its corresponding CIM property name, and its corresponding ASN.1 type. Each of these elements is separated by a space character.

文件

/var/sadm/wbem/snmp/map/050SUNWwbcou.map The SNMP Adapter Subagent for WBEM MIB-2 mapping file that Sun Microsystems provides contains SNMP Management Information Base (MIB) definitions for the CIM instrumentation that SNMP manages.

退出状态

The following exit values are returned:

- 0 Successful completion.
- 1 An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes.

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWwbco
CSI	Enabled
Interface Stability	Committed
MT-Level	Safe

另请参见 [snmpdx\(1M\), attributes\(5\)](#)

引用名 snoop – 捕获并检查网络包

用法概要

```
snoop [-aqrCDINPSvV] [-t [r | a | d]] [-c maxcount]
      [-d device] [-i filename] [-n filename] [-o filename]
      [-p first [, last]] [-s snaplen] [-x offset [, length]]
      [expression]
```

描述

snoop 从数据链路或 IP 接口捕获包并显示其内容。如果未指定数据链路或 IP 接口，snoop 将选取要使用的数据链路，并优先使用检测到 IP 通信的数据链路。snoop 使用 [pfmod\(7M\)](#) 和 [bufmod\(7M\)](#) STREAMS 模块从网络有效捕获包。捕获的包可在接收或保存到文件（该文件符合 RFC 1761）时显示，以便在以后检查。

snoop 可在单行汇总表或详细多行表单中显示这些包。在汇总表中，除某些 VLAN 包以外，仅显示与最高级别的协议相关的数据。如果包具有一个 VLAN 头，并且其 VLAN ID 为非零，snoop 将显示此包带有 VLAN 标记。例如，NFS 包将仅显示 NFS 信息。除刚才介绍的条件下的 VLAN 信息以外，系统将隐藏底层 RPC、UDP、IP 和以太网帧信息，但是，如果选择某个详细选项，则可以显示这些信息。

如果缺少名称服务（例如，LDAP 或 NIS），snoop 会将主机名显示为数字 IP 地址。

snoop 需要交互式接口。

选项

-C

列出根据内核包过滤器或 snoop 自己的过滤器的过滤器表达式生成的代码。

-D

在汇总行中显示捕获时删除的包的数目。

-N

根据捕获文件创建 IP 地址-名称文件。这必须与指定捕获文件的 **-i** 选项一起设置。地址-名称文件的名称与捕获文件相同，并附加了 **.names**。此文件在捕获站点记录 IP 地址到主机名的映射，并提高了捕获文件的可移植性。如果要在其他位置分析捕获文件，则生成一个 **.names** 文件。使用此标记时，不会显示包。

-I interface

使用 *interface* 指定的 IP 接口（例如，**lo0**）捕获网络中的 IP 包。[ipadm\(1M\)](#) 命令可用于列出可用 IP 接口。**-I** 选项和 **-d** 选项相互排斥。

-P

在非混杂模式下捕获包。仅显示发往主机的广播、多播或包。

-S

在汇总行中显示整个链路层帧的大小（以字节为单位）。

-V

详细汇总模式。从详细程度而言，此模式介于汇总模式和详细模式之间。此模式为包中的每个协议层显示一个汇总行，而不是仅为包中最高级别的协议显示汇总行。例如，对于 NFS 包，它将为 ETHER、IP、UDP、RPC 和 NFS 层各显示一行。通过 **grep** 可以轻松传输详细汇总模式输出以提取相关包。例如，要仅查看 RPC 汇总行，请输入以下内容：**example# snoop -i rpc.cap -V | grep RPC**

- a
在 `/dev/audio` 上侦听包（警告：可能有杂音）。
- c *maxcount*
在捕获 *maxcount* 包后退出。否则继续捕获，直到没有剩余磁盘空间或使用 Ctrl-C 中断为止。
- d *datalink*
使用 *datalink* 指定的 DLPI 数据链路（例如，`bge0` 或 `net0`）从网络中捕获链路层包。`dladm(1M)` `show-link` 子命令可用于列出可用数据链路。-d 选项和 -I 选项相互排斥。
- i *filename*
显示 *filename* 中以前捕获的包。如果没有此选项，snoop 将从网络接口读取包。如果存在 *filename.names* 文件，则会将其自动装载到 snoop IP 地址-名称映射表（请参见 -N 标志）。
- n *filename*
使用 *filename* 作为 IP 地址-名称映射表。此文件的格式必须与 `/etc/hosts` 文件的格式相同（后跟有主机名的 IP 地址）。
- o *filename*
在捕获包时将捕获到的包保存到 *filename* 中。（此 *filename* 称为“捕获文件”。）捕获文件的格式符合 RFC 1761。捕获包期间，将显示保存在此文件中的包的计数。如果您仅希望计算包的数目，而不保存到文件，请指定文件 `/dev/null`。
- p *first* [, *last*]
从捕获文件中选择一个或多个要显示的包。此文件中的 *first* 包即为第 1 个包。
- q
将网络包捕获到文件中时，不会显示包计数。这可以提高捕获包的性能。
- r
不会将 IP 地址解析为符号名称。这会防止 snoop 在捕获和显示包时生成网络通信。但是，如果使用 -n 选项，并在映射文件中找到某个地址，则将使用其相应名称。
- s *snaplen*
在 *snaplen* 字节后截断各个包。通常捕获整个包。仅当需要某些包头信息时，此选项才会有用。包截断操作是在内核中完成的，使您可以更好地利用流包缓冲区。这意味着，在高通信期间，因缓冲区溢出而删除包的可能性较小。此选项还会在将大型跟踪捕获到捕获文件时节省磁盘空间。要仅捕获 IP 头（无选项），请使用 *snaplen* 34。对于 UDP，请使用 42，对于 TCP，请使用 54。您可以使用包含 80 个字节的 *snaplen* 捕获 RPC 头。可以 120 个字节为单位捕获 NFS 头。
- t [*r*|*a*|*d*]
时间戳显示。时间戳精确至 4 微秒以内。缺省设置是以 *d*（增量）格式显示时间（自收到上一个包以来的时间）。选项 *a*（绝对）提供墙上时钟时间。选项 *r*（相对）提供相对于显示的第一个包的时间。此选项可以与 -p 选项配合使用，以便显示相对于任何选定包的时间。

-v

详细模式。详细输出包头。此显示针对每个包使用多个行，并且只能用于选定包。

-xoffset [, length]

以十六进制和 ASCII 格式显示包数据。offset 和 length 值选择要显示的包的部分。要显示整个包，请使用 offset 0。如果未提供 length 值，则会显示包的其余部分。

操作数

表达式

从网络或捕获文件中选择包。仅选择表达式为 true 的包。如果未提供表达式，则假定它为 true。

如果给定过滤器表达式，snoop 将为内核包过滤器或其自己的内部过滤器生成代码。如果使用网络接口捕获包，则将生成适用于内核包过滤器的代码。此过滤器作为流模块（即缓冲区模块的上游）实现。缓冲区模块会一直累积包，直到其已满为止，并将这些包传递到 snoop。在内核中不需要的包到达包缓冲区或 snoop 之前，内核包过滤器会拒绝这些包，因此此过滤器非常有效。内核包过滤器具有某些实现限制；它可以构建无法处理的过滤器表达式。此种情况下，snoop 会尝试拆分过滤器，并在内核中执行尽可能多的过滤操作。对于 snoop，其余过滤操作都由包过滤器完成。-c 标志可用于查看为内核包过滤器或 snoop 的包过滤器生成的代码。如果使用 -i 选项从捕获文件读取包，则仅使用 snoop 的包过滤器。

过滤器 *expression* 包含一个或多个布尔基元序列，这些基元可使用布尔运算符（AND、OR 和 NOT）组合在一起。应用布尔运算符的常规优先级规则。使用括号可以控制这些运算符的计算顺序。由于 shell 已知括号和其他过滤器表达式字符，因此通常需要将过滤器表达式包含在引号中。有关如何设置更有效的过滤器的信息，请参阅[示例 2](#)。

这些基元包括：

host hostname

如果源或目标地址是 hostname 的源或目标地址，则为 True。hostname 参数可以为文本地址。如果此名称不与其他表达式基元的名称冲突，则可以忽略关键字 host。例如，pinky 选择传送到主机 pinky 或从该主机接收的包，而 pinky and dinky 选择在主机 pinky AND dinky 之间交换的包。

使用的地址类型取决于 host 基元前面的基元。可能的限定符包括 inet、inet6、ether，或没有限定符。下面讨论了这三个基元。不提供任何基元等效于“inet host hostname or inet6 host hostname”。换言之，snoop 尝试根据与主机名关联的所有 IP 地址进行过滤。

inet 或 inet6

修改后面的 host 基元的限定符。如果为 inet，snoop 尝试根据名称查找返回的所有 IPv4 地址进行过滤。如果为 inet6，snoop 尝试根据名称查找返回的所有 IPv6 地址进行过滤。

ipaddr、atalkaddr 或 etheraddr

可以识别文本地址、IP 点分地址、AppleTalk 点分地址和以太网冒号地址。例如，

- “172.16.40.13”与采用该 IP 的所有包匹配
- “2::9255:a00:20ff:fe73:6e35”与采用此 IPv6 地址作为源或目标的所有包匹配；
- “65281.13”与采用此 AppleTalk 地址的所有包匹配；
- “8:0:20:f:b1:51”与采用此以太网地址作为源或目标的所有包匹配。

以字母开头的以太网地址解释为主机名。为了避免此种情况，请在指定地址时在此地址前面添加一个零。例如，如果以太网地址为 aa:0:45:23:52:44，通过添加一个前导零将此地址指定为 0aa:0:45:23:52:44。

from 或 src

修改以下的 host、net、ipaddr、atalkaddr、etheraddr、port 或 rpc 基元以便仅与源地址、端口或 RPC 回复匹配的限定符。

to 或 dst

修改以下的 host、net、ipaddr、atalkaddr、etheraddr、port 或 rpc 基元以便仅与目标地址、端口或 RPC 调用匹配的限定符。

ether

修改以下的 host 基元以便将名称解析为以太网地址的限定符。通常执行 IP 地址匹配操作。IPoIB (IP over InfiniBand) 等介质上不支持此选项。

ethertype *number*

如果以太网类型字段的值为 *number*，则为 True。如果 *number* 不为 0x8100 (VLAN)，并且包带有 VLAN 标记，表达式将与封装的以太网类型相匹配。

ip、ip6、arp、rarp、pppoed、pppoes

如果包是适当的 ethertype，则为 True。

vlan

如果包 ethertype 为 VLAN，并且 VLAN ID 不为零，则为 True。

vlan-id *id*

对于 ethertype 为 VLAN 且 ID 为 *id* 的包，则为 True。

pppoe

如果包 ethertype 为 pppoed 或 pppoes，则为 True。

broadcast

如果包为广播包，则为 True。对于以太网，等效于 ether[2:4] = 0xffffffff。IPoIB (IP over InfiniBand) 等介质上不支持此选项。

multicast

如果包为多播包，则为 True。在以太网上等效于“ether[0] & 1 = 1”。IPoIB (IP over InfiniBand) 等介质上不支持此选项。

bootp, dhcp

如果包是未分段的 IPv4 UDP 包，并且其源端口为 BOOTPS (67) 且目标端口为 BOOTPC (68)，或者源端口为 BOOTPC (68) 且目标为 BOOTPS (67)，则为 True。

dhcp6

如果包是未分段的 IPv6 UDP 包，并且其源端口为 DHCPV6-SERVER (547) 且目标端口为 DHCPV6-CLIENT (546)，或者源端口为 DHCPV6-CLIENT (546) 且目标为 DHCPV6-SERVER (547)，则为 True。

apple

如果包为 Apple Ethertalk 包，则为 True。等效于“ethertype 0x809b or ethertype 0x80f3”。

decnet

如果包为 DECNET 包，则为 True。

greater *length*

如果包长于 *length*，则为 True。

less *length*

如果包短于 *length*，则为 True。

udp, tcp, icmp, icmp6, ah, esp

如果 IP 或 IPv6 协议为适当的类型，则为 True。

net *net*

如果 IP 源或目标地址的网络号为 *net*，则为 True。*from* 或 *to* 限定符可用于选择网络号仅存在于源或目标地址中的包。

port *port*

如果源或目标端口为 *port*，则为 True。*port* 可以是来自 */etc/services* 的端口号或名称。*tcp* 或 *udp* 基元只能用于选择 TCP 或 UDP 端口。*from* 或 *to* 限定符可用于选择 *port* 仅作为源或目标的包。

rpc *prog* [, *vers* [, *proc*]]

如果包是 RPC 调用或 *prog* 标识的协议的回复包，则为 True。*prog* 可以是来自 */etc/rpc* 的 RPC 协议的名称，也可以是程序编号。*vers* 和 *proc* 可用于进一步限定程序 *version* 和 *procedure* 编号，例如，*rpc nfs, 2, 0* 选择 NFS 空过程的所有调用和回复。*to* 或 *from* 限定符可用于仅选择调用或回复包。

zone *zoneid*

如果 *zoneid* 与 *ipnet* 设备上接收到的包的源或目标 *zoneid* 匹配，则为 True。

ldap

如果包是端口 389 上的 LDAP 包，则为 True。

gateway *host*

如果包已使用 *host* 作为网关，也即，以太网源或目标地址用于 *host*（而非 IP 地址），则为 True。等效于“ether host *host* and not host *host*”。

nofrag

如果包未分段或是 IP 分段序列中的第一个包，则为 True。等效于 *ip[6:2] & 0x1fff = 0*。

expr relop expr

如果保持关系，其中 *relop* 为 >、<、>=、<=、=、!= 之一，并且 *expr* 是由数字、包字段选择器、*length* 基元和算术运算符 +、-、*、&、|、^ 和 % 组成的算术表达式，则为 True。在算术运算符之间应用关系运算符和常规优先级规则之前，系统会先计算 *expr* 中的算术运算符，例如，在执行加法运算之前先执行乘法运算。括号可用于控制计算顺序。要使用包中某个字段的值，请使用下面的语法：

```
base[expr [: size ] ]
```

其中 *expr* 计算包中离 *base* 偏移量的偏移量的值，前一个偏移量可以为 *ether*、*ip*、*ip6*、*udp*、*tcp* 或 *icmp*。size 值指定字段大小。如果未指定，则假定为 1。其他合法值为 2 和 4。例如，

```
ether[0] & 1 = 1
```

等效于 multicast

```
ether[2:4] = 0xffffffff
```

等效于 broadcast。

```
ip[ip[0] & 0xf * 4 : 2] = 2049
```

等效于 udp[0:2] = 2049

```
ip[0] & 0xf > 5
```

使用选项选择 IP 包。

```
ip[6:2] & 0x1fff = 0
```

消除 IP 分段。

```
udp and ip[6:2]&0x1fff = 0 and udp[6:2] != 0
```

使用 UDP 校验和查找所有包。

length 基元可用于获取包的长度。例如，“*length* > 60”等效于“greater 60”，“*ether*[*length* - 1]”获取包中最后一个字节的值。

与

在两个布尔值之间执行逻辑 AND 运算。通过并列的两个布尔表达式来暗指 AND 运算，例如，“dinky pinky”与“dinky AND pinky”相同。

or 或，

在两个布尔值之间执行逻辑 OR 运算。可改用逗号，例如，“dinky,pinky”与“dinky OR pinky”相同。

not 或 !

对下列布尔值执行逻辑 NOT 运算。在 AND 或 OR 之前计算此运算符。

slp

如果包为 SLP 包，则为 True。

sctp

如果包为 SCTP 包，则为 True。

ospf

如果包为 OSPF 包，则为 True。

示例

示例1 使用 snoop 命令

捕获所有包，并在接收这些包时显示它们：

```
example# snoop
```

使用主机 funky 作为源或目标来捕获包，并在接收这些包时显示它们：

```
example# snoop funky
```

捕获 funky 和 pinky 之间的包，并将其保存到文件。然后，使用相对于捕获的第一个包的时间（以秒为单位）检查包：

```
example# snoop -o cap funky pinky
```

```
example# snoop -i cap -t r | more
```

要查看其他捕获文件中的选定包：

```
example# snoop -i pkts -p 99,108
```

```

99  0.0027  boutique -> sunroof      NFS C GETATTR FH=8E6
100 0.0046  sunroof -> boutique      NFS R GETATTR OK
101 0.0080  boutique -> sunroof NFS C RENAME FH=8E6C MTra00192 to .nfs08
102 0.0102  marmot -> viper          NFS C LOOKUP FH=561E screen.r.13.i386
103 0.0072  viper -> marmot          NFS R LOOKUP No such file or directory
104 0.0085  bugbomb -> sunroof      RLOGIN C PORT=1023 h
105 0.0005  kandinsky -> sparky      RSTAT C Get Statistics
106 0.0004  beebledbrox -> sunroof  NFS C GETATTR FH=0307
107 0.0021  sparky -> kandinsky      RSTAT R
108 0.0073  office -> jeremiah      NFS C READ FH=2584 at 40960 for 8192
```

要更详细地查看包 101：

```
example# snoop -i pkts -v -p101
```

```
ETHER: ----- Ether Header -----
```

```
ETHER:
```

```
ETHER: Packet 101 arrived at 16:09:53.59
```

```
ETHER: Packet size = 210 bytes
```

```
ETHER: Destination = 8:0:20:1:3d:94, Sun
```

```
ETHER: Source      = 8:0:69:1:5f:e, Silicon Graphics
```

```
ETHER: Ethertype = 0800 (IP)
```

```
ETHER:
```

```
IP: ----- IP Header -----
```

```
IP:
```

```
IP: Version = 4, header length = 20 bytes
```

```
IP: Type of service = 00
```

示例1 使用 snoop 命令 (续)

```

IP:      ..0. .... = routine
IP:      ...0 .... = normal delay
IP:      .... 0... = normal throughput
IP:      .... .0.. = normal reliability
IP: Total length = 196 bytes
IP: Identification 19846
IP: Flags = 0X
IP: .0.. .... = may fragment
IP: ..0. .... = more fragments
IP: Fragment offset = 0 bytes
IP: Time to live = 255 seconds/hops
IP: Protocol = 17 (UDP)
IP: Header checksum = 18DC
IP: Source address = 172.16.40.222, boutique
IP: Destination address = 172.16.40.200, sunroof
IP:
UDP: ----- UDP Header -----
UDP:
UDP: Source port = 1023
UDP: Destination port = 2049 (Sun RPC)
UDP: Length = 176
UDP: Checksum = 0
UDP:
RPC: ----- SUN RPC Header -----
RPC:
RPC: Transaction id = 665905
RPC: Type = 0 (Call)
RPC: RPC version = 2
RPC: Program = 100003 (NFS), version = 2, procedure = 1
RPC: Credentials: Flavor = 1 (Unix), len = 32 bytes
RPC:   Time = 06-Mar-90 07:26:58
RPC:   Hostname = boutique
RPC:   Uid = 0, Gid = 1
RPC:   Groups = 1
RPC: Verifier   : Flavor = 0 (None), len = 0 bytes
RPC:
NFS: ----- SUN NFS -----
NFS:
NFS: Proc = 11 (Rename)
NFS: File handle = 000016430000000100080000305A1C47
NFS:           597A0000000800002046314AFC450000
NFS: File name = MTra00192
NFS: File handle = 000016430000000100080000305A1C47
NFS:           597A0000000800002046314AFC450000
NFS: File name = .nfs08
NFS:

```

示例1 使用 snoop 命令 (续)

要仅查看 sunroof 和 boutique 之间的 NFS 包：

```
example# snoop -i pkts rpc nfs and sunroof and boutique
1  0.0000  boutique -> sunroof    NFS C GETATTR FH=8E6C
2  0.0046  sunroof -> boutique    NFS R GETATTR OK
3  0.0080  boutique -> sunroof    NFS C RENAME FH=8E6C MTra00192 to .nfs08
```

要将这些包保存到新捕获文件：

```
example# snoop -i pkts -o pkts.nfs rpc nfs sunroof boutique
```

要查看封装的包，包中将存在一个封装指示器：

```
example# snoop ip-in-ip
sunroof -> boutique ICMP Echo request    (1 encaps)
```

如果在封装包中使用 -V：

```
example# snoop -V ip-in-ip
sunroof -> boutique  ETHER Type=0800 (IP), size = 118 bytes
sunroof -> boutique  IP D=172.16.40.222 S=172.16.40.200 LEN=104, ID=27497
sunroof -> boutique  IP D=10.1.1.2 S=10.1.1.1 LEN=84, ID=27497
sunroof -> boutique  ICMP Echo request
```

示例2 设置更有效的过滤器

要设置更有效的过滤器，应在表达式末尾使用下列过滤器，以便可以在内核中设置表达式的第一个部分：greater、less、port、rpc、nofrag 和 relop。使用无法在内核中设置的这些基元时，如果存在 OR，则很难拆分过滤。而应改用括号以便强制应执行 OR 运算的基元。

要在 port 80 上封装 funky 和 pinky 之间的 tcp 或 udp 类型的包：

```
example# snoop funky and pinky and port 80 and tcp or udp
```

由于内核过滤器无法处理基元 port，并且表达式中还存在 OR，一种更有效的过滤方法是将 OR 移至表达式末尾，并使用括号在 tcp 和 udp 之间强制 OR：

```
example# snoop funky and pinky and (tcp or udp) and port 80
```

退出状态

- 0 成功完成。
- 1 出现错误。

文件

- /dev/audio 指向系统的主要音频设备的符号链接。
- /dev/null 空文件。
- /etc/hosts 主机名数据库。
- /etc/rpc RPC 程序编号数据库。

/etc/services Internet 服务和别名。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/network/network-clients
接口稳定性	Obsolete（已过时）

另请参见 [dladm\(1M\)](#)、[ipadm\(1M\)](#)、[netstat\(1M\)](#)、[hosts\(4\)](#)、[rpc\(4\)](#)、[services\(4\)](#)、[attributes\(5\)](#)、[audio\(7I\)](#)

Callaghan, B. 和 Gilligan, R. 合著的《RFC 1761, Snoop Version 2 Packet Capture File Format》。Network Working Group 出版。1995 年 2 月。

警告 实时包解释的处理开销要高得多。因此，包删除计数可能较高。为了进行更可靠的捕获，请使用 -o 选项将原始包输出到文件，并脱机分析这些包。

未过滤的包捕获会对主机施加大量处理负载，特别是当实时解释捕获的包时。如果使用详细选项，此处理负载将进一步增加。由于大量使用 snoop 会拒绝其他进程的计算资源，因此不应在生产服务器上使用它。仅限在专用计算机上大量使用 snoop。

snoop 不会重装 IP 分段。解释较高级别的协议将在第一个 IP 分段末尾暂停。

使用 snoop 的负面影响是可能会生成额外的包。例如，它可以使用网络名称服务 (NIS) 将 IP 地址转换为主机名以便显示。通过捕获到文件以便在以后显示，可以延迟地址-名称映射，直到完成捕获会话后为止。捕获到挂载了 NFS 的文件中还会生成额外的包。

将 snaplen（-s 选项）设置为较小值会删除解释较高级别的协议所需的头信息。确切截止值取决于使用的网络和协议。对于在 10 Mb/s 以太网上使用 UDP 的 NFS 版本 2 通信，请不要将 snaplen 设置为小于 150 字节的值。对于在 100 Mb/s 以太网上使用 TCP 的 NFS 版本 3 通信，snaplen 应为 250 字节或更大值。

snoop 需要来自 RPC 请求的信息，以便完全解释 RPC 回复。如果捕获文件或包范围中的 RPC 回复前面没有请求，则仅显示 RPC 回复头。

引用名	soconfig – configure transport providers for use by sockets
用法概要	<pre> /usr/sbin/soconfig -d <i>dir</i> /usr/sbin/soconfig -f <i>file</i> /usr/sbin/soconfig <i>family type protocol</i> [<i>module</i> <i>path</i>] /usr/sbin/soconfig -l [-np] </pre>
描述	The soconfig utility configures the transport provider driver for use with sockets. It specifies how the family, type, and protocol parameters in the socket(3SOCKET) call are mapped to the name of a transport provider such as <code>/dev/tcp</code>. This utility can be used to add an additional mapping or remove a previous mapping. The init(1M) utility uses soconfig with the sock2path.d(4) file during the booting sequence.
选项	The following options are supported: -d <i>dir</i> Set up the soconfig configuration for each driver according to the information stored in the files in <i>dir</i>. -f <i>file</i> Set up the soconfig configuration for each driver according to the information stored in <i>file</i>. A soconfig file consists of lines of at least the first three fields listed below, separated by spaces: <pre><i>family type protocol</i> [<i>module</i> <i>path</i>]</pre> These fields are described in the OPERANDS section below. An example of <i>file</i> can be found in the EXAMPLES section below. -l [-np] Dump the in-kernel socket configuration table. This option has the following suboptions: -n Do not convert socket family, type, and protocol to string representation. -p Display machine-parseable output.
操作数	The following operands are supported: <i>family</i> The protocol family as listed in the <code>/usr/include/sys/socket.h</code> file, expressed as an integer. <i>type</i> The socket type as listed in the <code>/usr/include/sys/socket.h</code> file, expressed as an integer.

protocol
The protocol number as specified in the family-specific include file, expressed as an integer. For example, for AF_INET this number is specified in /usr/include/netinet/in.h. An unspecified protocol number is denoted with the value zero.

module | path
The module name or path name of a device that corresponds to the transport provider, such as tcp or /dev/tcp. Modules must reside in kernel/socketmod. A device name must begin with /dev. If this parameter is specified, the configuration will be added for the specified family, type, and protocol. If this parameter is not specified, the configuration will be removed.

示例

示例1 Using soconfig
The following example sets up a module for family AF_INET and type SOCK_STREAM:
example# **soconfig 2 2 0 tcp**
The following example sets up /dev/tcp for family AF_INET and type SOCK_STREAM:
example# **soconfig 2 2 0 /dev/tcp**
The following is a sample file used with the -f option. Comment lines begin with a hash mark (#):

#	Family	Type	Protocol	Module Path
	2	2	0	tcp
	2	2	6	tcp
	2	1	0	udp
	2	1	17	udp
	1	2	0	/dev/ticotsord
	1	1	0	/dev/ticlts
	2	4	0	icmp

文件

/etc/sock2path.d
Directory containing files with mappings from sockets to transport providers.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[init\(1M\)](#), [sock2path.d\(4\)](#), [attributes\(5\)](#)

《[Network Interface Guide](#)》

引用名

soladdapp – add an application to the Solstice application registry

用法概要

```
/usr/snadm/bin/soladdapp [-r registry] -n name -i icon -e executable
[args]
```

描述

soladdapp adds an application to the Solstice application registry. After it is added, the application is displayed in the Solstice Launcher main window (see [solstice\(1M\)](#)).

选项

-r registry

Define the full path name of the Solstice registry file.

-n name

Define the name of the tool to be registered.

-i icon

Define the full path name of the tool icon.

-e executable

Define the full path name of the tool.

args

Specify any arguments to use with the tool.

When executed without options, soladdapp uses /opt/SUNWadm/etc/.solstice_registry (the default registry path).

返回值

0

on success

1

on failure

2

if the registry is locked

3

if the entry is a duplicate.

示例

示例 1

A sample display of the soladdapp command.

The following adds an application called Disk Manager to the Solstice application registry for display in the Solstice Launcher main window.

```
# soladdapp -r /opt/SUNWadm/etc/.solstice_registry -n "Disk Manager"
-i /opt/SUNWdsk/etc/diskmgr.xpm -e /opt/SUNWdsk/bin/diskmgr
```

文件

/opt/SUNWadm/etc/.solstice_registry

The default registry path.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWsadml

另请参见 [soldelapp\(1M\)](#), [solstice\(1M\)](#), [attributes\(5\)](#)

附注 Globally registered applications are used by local and remote users sharing the software in a particular /opt directory. They can be added only using soladdapp.

引用名	<code>soldelapp</code> – remove an application from the Solstice application registry
用法概要	<code>/usr/snadm/bin/soldelapp [-r <i>registry</i>] -n <i>name</i></code>
描述	<code>soldelapp</code> removes an application from the Solstice application registry. After removal, the application is no longer displayed in the Solstice Launcher main window (see solstice(1M)).
选项	<code>-r <i>registry</i></code> Define the full path name of the Solstice registry file. <code>-n <i>name</i></code> Define the name of the tool to be removed. When executed without options, <code>soldelapp</code> uses <code>/opt/SUNWadm/etc/.solstice_registry</code> (the default registry path).
返回值	0 on success 1 on failure 2 if the registry is locked 3 if <i>name</i> is not found in the registry 4 if the named registry or default registry is not found
示例	示例 1 A sample display of the <code>soldelapp</code> command. The following removes an application called Disk Manager from the Solstice application registry and the Solstice Launcher main window. <pre># soldelapp -r /opt/SUNWadm/etc/.solstice_registry -n "Disk Manager"</pre>
文件	<code>/opt/SUNWadm/etc/.solstice_registry</code> The default registry file.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWsadml

另请参见 [soladdapp\(1M\)](#), [solstice\(1M\)](#), [attributes\(5\)](#)

附注 Globally registered applications are used by local and remote users sharing the software in a particular `/opt` directory. They can be removed only using `soldelapp`.

引用名

solstice – access system administration tools with a graphical user interface

用法概要

/bin/solstice

描述

solstice used on a system presents the Solstice Launcher, a graphical user interface that provides access to the Solstice AdminSuite product family of system administration tools. The tools that appear in the launcher depend on what Solstice products you installed on your system.

Help is available by using the Help button.

用法

The Solstice Launcher allows you to do the following tasks:

Launch applications	Use the Solstice Launcher to launch system administration tools.
Register applications	Use the Solstice Launcher to add and register applications locally with the launcher.
Remove applications	Use the Solstice Launcher to remove locally registered applications.
Customize application properties	Use the Solstice Launcher to show, hide, or remove applications in the launcher, reorder the icons, change the launcher window width, modify applications properties, and add applications.

文件

/ \$HOME/.solstice_registry

Local registry information.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	SUNWsadml

另请参见 [soladdapp\(1M\)](#), [soldelapp\(1M\)](#), [attributes\(5\)](#)

附注

The Solstice Launcher adds or removes local applications that are private to the user (not local to the system) only. The properties of globally registered applications that are used by local and remote users sharing the software from a particular /opt directory cannot be modified from the Solstice Launcher. To register global applications for use by local and remote users, use the [soladdapp\(1M\)](#) command. To remove globally registered applications, use the [soldelapp\(1M\)](#) command.

引用名	spptun – PPP tunneling driver utility	
用法概要	<pre>spptun plumb spptun plumb [-s sap] protocol device spptun unplumb interface spptun query</pre>	
描述	The spptun utility is used to configure and query the Solaris PPP tunneling device driver, /dev/spptun. Currently, only PPP over Ethernet (PPPoE) is supported, so the plumb and unplumb arguments are used to specify Ethernet interfaces that are to be used for PPPoE, and the query option lists the plumbed interfaces. The use of spptun to add interfaces is similar to the use of ifconfig(1M) to add interfaces to IP. The plumbing is done once for each interface, preferably at system start-up time, and is not normally manipulated on a running system. If multiple instances of PPP are run over a single interface, they share the plumbing to that interface. Plumbing for each session is not required (and not possible for PPPoE). The proper way to plumb interfaces for PPPoE is to list the interfaces, one per line, in the /etc/ppp/pppoe.if file. If alternate Ethertype (SAP) values are necessary, then include the PPPoE Session and Discovery Stage values as hexadecimal on the same line. The line format is: <pre>interface [session [discovery]]</pre> The defaults are the Ethertypes specified in RFC 2516, and most users should not need to set these values. See the examples for one possible use.	
用法	<pre>spptun plumb</pre> <pre>spptun plumb [-s sap] protocol device</pre>	When specified with no additional arguments, the plumb argument lists the protocols that are supported by the utility. These are the strings that are used as the <i>protocol</i> argument below. This plumbs a new interface into the driver. The <i>protocol</i> parameter is pppoe for the PPP-carrying “Session Stage” connection or pppoed for the PPPoE “Discovery Stage” connection. Both connections must be present for each Ethernet interface that is to be used for PPPoE. The <i>device</i> parameter is the path name of the Ethernet interface to use (use ifconfig(1M) to list available devices). If the path begins with /dev/, then this portion may be omitted. The -s sap option can be specified to use an alternate Ethertype (SAP) for the selected protocol. The sap value must be given in

示例

	hexadecimal. Some access servers use Ethertypes for PPPoE different from those in RFC 2516. The defaults are 8864 for pppoe and 8863 for pppoeed.
<code>sppptun unplumb <i>interface</i></code>	This removes an existing interface from the driver and terminates any PPP sessions that were using the interface. The <i>interface</i> parameter is the name of the interface as reported when the interface was plumbed.
<code>sppptun query</code>	Displays the canonical names of all interfaces plumbed into the /dev/sppptun device driver.

示例 1 Setting up to Use PPPoE on hme0
Plumb the hme0 interface.

```
# sppptun plumb pppoe hme0
hme0:pppoed
# sppptun plumb pppoe hme0
hme0:pppoe
```

Remove the hme0 interface.

```
# sppptun unplumb hme0:pppoed
# sppptun unplumb hme0:pppoe
```

示例 2 Script to Remove All Plumbed Interfaces

```
#!/bin/sh
for intf in `sppptun query`
do
    sppptun unplumb $intf
done
```

示例 3 Interoperating with 3COM HomeConnect Dual Link ADSL

```
# dladm show-link
LINK      CLASS  MTU   STATE  OVER
nge0      phys   1500  up     --
# echo nge0 3c13 3c12 > /etc/ppp/pppoe.if
```

退出状态

The following exit values are returned:

- 0 Successful completion.
- 1 One or more errors occurred.

文件 /etc/ppp/pppoe.if list of Ethernet interfaces to be plumbed at boot time

 /usr/sbin/spptun executable command

 /dev/spptun Solaris PPP tunneling device driver

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/ppp/tunnel

另请参见 [pppd\(1M\)](#), [pppoec\(1M\)](#), [pppoed\(1M\)](#), [spptun\(7M\)](#)

RFC 2516, Method for Transmitting PPP Over Ethernet (PPPoE), Mamakos et al, February 1999

引用名	spray – spray packets								
用法概要	<pre>/usr/sbin/spray [-c count] [-d delay] [-l length] [-t nettype] host</pre>								
描述	spray sends a one-way stream of packets to <i>host</i> using RPC, and reports how many were received, as well as the transfer rate. The <i>host</i> argument can be either a name or an Internet address. spray is not useful as a networking benchmark, as it uses unreliable connectionless transports, UDP for example. spray can report a large number of packets dropped when the drops were caused by spray sending packets faster than they can be buffered locally, that is, before the packets get to the network medium.								
选项	<table><tr><td>-c <i>count</i></td><td>Specify how many packets to send. The default value of <i>count</i> is the number of packets required to make the total stream size 100000 bytes.</td></tr><tr><td>-d <i>delay</i></td><td>Specify how many microseconds to pause between sending each packet. The default is 0.</td></tr><tr><td>-l <i>length</i></td><td>The <i>length</i> parameter is the numbers of bytes in the Ethernet packet that holds the RPC call message. Since the data is encoded using XDR, and XDR only deals with 32 bit quantities, not all values of <i>length</i> are possible, and spray rounds up to the nearest possible value. When <i>length</i> is greater than 1514, then the RPC call can no longer be encapsulated in one Ethernet packet, so the <i>length</i> field no longer has a simple correspondence to Ethernet packet size. The default value of <i>length</i> is 86 bytes, the size of the RPC and UDP headers.</td></tr><tr><td>-t <i>nettype</i></td><td>Specify class of transports. Defaults to netpath. See rpc(3NSL) for a description of supported classes.</td></tr></table>	-c <i>count</i>	Specify how many packets to send. The default value of <i>count</i> is the number of packets required to make the total stream size 100000 bytes.	-d <i>delay</i>	Specify how many microseconds to pause between sending each packet. The default is 0.	-l <i>length</i>	The <i>length</i> parameter is the numbers of bytes in the Ethernet packet that holds the RPC call message. Since the data is encoded using XDR, and XDR only deals with 32 bit quantities, not all values of <i>length</i> are possible, and spray rounds up to the nearest possible value. When <i>length</i> is greater than 1514, then the RPC call can no longer be encapsulated in one Ethernet packet, so the <i>length</i> field no longer has a simple correspondence to Ethernet packet size. The default value of <i>length</i> is 86 bytes, the size of the RPC and UDP headers.	-t <i>nettype</i>	Specify class of transports. Defaults to netpath. See rpc(3NSL) for a description of supported classes.
-c <i>count</i>	Specify how many packets to send. The default value of <i>count</i> is the number of packets required to make the total stream size 100000 bytes.								
-d <i>delay</i>	Specify how many microseconds to pause between sending each packet. The default is 0.								
-l <i>length</i>	The <i>length</i> parameter is the numbers of bytes in the Ethernet packet that holds the RPC call message. Since the data is encoded using XDR, and XDR only deals with 32 bit quantities, not all values of <i>length</i> are possible, and spray rounds up to the nearest possible value. When <i>length</i> is greater than 1514, then the RPC call can no longer be encapsulated in one Ethernet packet, so the <i>length</i> field no longer has a simple correspondence to Ethernet packet size. The default value of <i>length</i> is 86 bytes, the size of the RPC and UDP headers.								
-t <i>nettype</i>	Specify class of transports. Defaults to netpath. See rpc(3NSL) for a description of supported classes.								

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/network-clients

另请参见 [rpc\(3NSL\)](#), [attributes\(5\)](#)

引用名	srptadm – 管理 SRP 目标
用法概要	<pre>srptadm modify-target [-e, --enable] [-d, --disable] [-r, --reset] <i>hca_guid</i> srptadm list-target [<i>hca_guid</i>] srptadm modify-defaults [-e, --enable] [-d, --disable] srptadm list-defaults</pre>
描述	srptadm 命令管理 stmfadm(1M) 和 libstmf(3LIB) 中描述的 SCSI 目标模式框架内的 SCSI RDMA 协议 (SCSI RDMA Protocol, SRP) 目标端口。这允许 SRP 启动器使用 SCSI 协议访问 SCSI 目标模式框架 (SCSI Target Mode Framework, STMF) 逻辑单元。 srptadm 是作为一组子命令以及用于各个子命令的选项和操作数实现的。下文中的相应部分分别描述了这些子命令。
子命令	list-defaults 列出关于缺省属性的信息。此子命令没有选项。 list-target <i>hca_guid</i> 如果指定了 <i>hca_guid</i>，则列出目标 HCA 的属性。否则，将列出所有 HCA 的属性。 modify-defaults [-e, --enable] [-d, --disable] 修改缺省参数。 -e, --enable 为没有使用 modify-target 显式禁用的所有 HCA 启用 SRP 目标创建。 -d, --disable 为没有使用 modify-target 显式启用的所有 HCA 禁用 SRP 目标创建。 modify-target [-e, --enable] [-d, --disable] [-r, --reset] <i>hca_guid</i> 为指定的 HCA 设置 SRP 目标属性。 -e, --enable 在此 HCA 上启用 SRP 目标创建。 -d, --disable 在此 HCA 上禁用 SRP 目标创建。 -r, --reset 清除特定于 HCA 的信息并重置为缺省值。如果指定了此选项，则与此 HCA 关联的 SRP 目标（如果有）将不进行任何修改，直至重新启动 SRP 目标 SMF 服务。
操作数	list-target 和 modify-target 子命令具有以下操作数。 <i>hca_guid</i> 此系统上可以为提供 SRP 目标服务的 InfiniBand 主机通道适配器 (Host Channel Adapter, HCA) 的 GUID。GUID 必须为以下格式之一：

3BA000100CD18	基础十六进制格式。
0003BA000100CD18	带有前导零的基础十六进制格式。
hca:3BA000100CD18	cfgadm(1M) 中的格式。
eui.0003BA000100CD18	EUI 格式。

用法 如果在 SRP 服务处于 online 状态时更改缺省状态，则现有目标的状态不会更改，直至服务重新启动。

如果 SRP 目标服务处于 online 状态，则更改目标状态会立即生效。设置为 disabled 的目标将脱机并被删除；设置为 enabled 的目标将立即创建。

示例 示例1 列出缺省属性
以下命令列出缺省的 SRP 目标服务属性。

```
# srptadm list-defaults
```

示例2 更改缺省行为
以下命令将 SRP 目标服务的缺省行为更改为服务处于启用状态时不创建 SRP 目标。

```
# srptadm modify-defaults -d
```

示例3 列出特定 HCA 的属性
以下命令列出特定 HCA 的 SRP 目标属性。

```
# cfgadm | grep hca
hca:3BA000100CD18      IB-HCA      connected   configured  ok
hca:3BA000100D030      IB-HCA      connected   configured  ok

# srptadm list-target hca:3BA000100CD18
Target HCA 3BA000100CD18:
  Enabled      : true
  SRP Target Name  : eui.0003BA000100CD18
  Operational Status : online
```

示例4 为特定 HCA 禁用服务
以下命令为特定 HCA 禁用 SRP 目标服务。

```
# srptadm modify-target -d 3BA000100CD18
```

```
可使用 list-target 命令查看更改：

# srptadm list-target hca:3BA000100CD18
Target HCA 3BA000100CD18:
  Enabled      : false
  SRP Target Name  : eui.0003BA000100CD18
```


示例 4 为特定 HCA 禁用服务 (续)

Operational Status : -

示例 5 重新启用目标服务

以下命令序列重新启用 SRP 目标服务并显示 SRP 和 STMF 目标属性。

```
# srptadm modify-target -e 3BA000100CD18
# srptadm list-target hca:3BA000100CD18
Target HCA 3BA000100CD18:
    Enabled           : true
    SRP Target Name   : eui.0003BA000100CD18
    Operational Status : online

# stmfadm list-target -v eui.0003BA000100CD18
Target: eui.0003BA000100CD18
    Operational Status: Online
    Provider Name      : srpt
    Alias              : -
    Protocol           : SRP
    Sessions           : 0
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/storage/scsi-rdma/scsi-rdma-target
接口稳定性	Committed (已确定)

另请参见 [cfgadm\(1M\)](#)、[stmfadm\(1M\)](#)、[libsprt\(3LIB\)](#)、[libstmf\(3LIB\)](#)、[attributes\(5\)](#)、[srpt\(7D\)](#)

引用名	sshd – secure shell daemon
用法概要	<pre>sshd [-deiqtD46] [-b bits] [-f config_file] [-g login_grace_time] [-h host_key_file] [-h PKCS#11 URI] [-p port] [-V client_protocol_id]</pre>
描述	The sshd (Secure Shell daemon) is the daemon program for ssh(1). Together these programs replace rlogin and rsh, and provide secure encrypted communications between two untrusted hosts over an insecure network. The programs are intended to be as easy to install and use as possible. sshd is the daemon that listens for connections from clients. It forks a new daemon for each incoming connection. The forked daemons handle key exchange, encryption, authentication, command execution, and data exchange. This implementation of sshd supports SSH protocol version 2 only.
SSH Protocol Version 1	This protocol version is no longer supported in the daemon program. It is supported only in the client. See ssh(1) manual page for more information.
SSH Protocol Version 2	Each host has a host-specific DSA/RSA key used to identify the host. Forward security is provided through a Diffie-Hellman key agreement. This key agreement results in a shared session key. The rest of the session is encrypted using a symmetric cipher, currently 128/192/256-bit AES CBC or CTR, 128/256-bit ARCFOUR, Blowfish, or 3DES.. The client selects the encryption algorithm to use from those offered by the server. Additionally, session integrity is provided through a cryptographic message authentication code (hmac-sha1 or hmac-md5). Protocol version 2 provides a public key based user authentication method (PubKeyAuthentication) GSS-API based user authentication, conventional password authentication, and a generic prompt/reply protocol for password-based authentication.
Command Execution and Data Forwarding	If the client successfully authenticates itself, a dialog for preparing the session is entered. At this time the client can request things like allocating a pseudo-tty, forwarding X11 connections, forwarding TCP/IP connections, or forwarding the authentication agent connection over the secure channel. Finally, the client either requests a shell or execution of a command. The sides then enter session mode. In this mode, either side may send data at any time, and such data is forwarded to/from the shell or command on the server side, and the user terminal on the client side. When the user program terminates and all forwarded X11 and other connections have been closed, the server sends command exit status to the client, and both sides exit. sshd can be configured using command-line options or the configuration file /etc/ssh/ssh_config, described in ssh_config(4). Command-line options override values specified in the configuration file.

sshd rereads its configuration file when it receives a hangup signal, SIGHUP, by executing itself with the name it was started as, that is, `/usr/lib/ssh/sshd`.

Host Access Control

The sshd daemon uses TCP Wrappers to restrict access to hosts. It uses the service name of sshd for `hosts_access()`. For more information on TCP Wrappers see `tcpd(1M)` and `hosts_access(3)` man pages, which are part of the `SUNWsman` package (they are not SunOS man pages). TCP wrappers binaries, including `libwrap`, are in `security/tcp-wrapper`, a required package for `service/network/ssh`, the package containing sshd.

选项

The options for sshd are as follows:

`-b bits`

Specifies the number of bits in the server key (the default is 768).

`-d`

Debug mode. The server sends verbose debug output to the system log, and does not put itself in the background. The server also will not fork and will only process one connection. This option is only intended for debugging for the server. Multiple `-d` options increase the debugging level. Maximum is 3.

`-e`

When this option is specified, sshd will send the output to standard error instead of to the system log.

`-f configuration_file`

Specifies the name of the configuration file. The default is `/etc/ssh/sshd_config`. sshd refuses to start if there is no configuration file.

`-g login_grace_time`

Gives the grace time for clients to authenticate themselves (the default is 300 seconds). If the client fails to authenticate the user within this number of seconds, the server disconnects and exits. A value of zero indicates no limit.

`-h host_key_file`

Specifies a file from which a host key is read. This option must be given if sshd is not run as root (as the normal host key files are normally not readable by anyone but root). The default is `/etc/ssh/ssh_host_key` for protocol version 1, and `/etc/ssh/ssh_host_rsa_key` and `/etc/ssh/ssh_host_dsa_key` for protocol version 2. It is possible to have multiple host key files for the different protocol versions and host key algorithms.

`-h PKCS#11 URI`

Instead of working with a host key file, work with a certificate and a private key stored in the PKCS#11 token. See also `Using X.509 Certificates` section below.

`-i`

Specifies that sshd is being run from `inetd`. sshd is normally not run from `inetd` because it needs to generate the server key before it can respond to the client, and this may take tens of

seconds. Clients would have to wait too long if the key was regenerated every time. However, with small key sizes (for example, 512) using `sshd` from `inetd` may be reasonable.

-o *option*

Can be used to specify options in the format used in the configuration file. This is useful for specifying options for which there are no separate command-line flags.

-p *port*

Specifies the port on which the server listens for connections (the default is 22).

-q

Quiet mode. Nothing is sent to the system log. Normally the beginning, authentication, and termination of each connection is logged.

-t

Test mode. Check only the validity of the configuration file and the sanity of the keys. This is useful for updating `sshd` reliably as configuration options might change.

-D

When this option is specified `sshd` does not detach and does not become a daemon. This allows easy monitoring of `sshd`.

-4

Forces `sshd` to use IPv4 addresses only.

-6

Forces `sshd` to use IPv6 addresses only.

扩展描述

authorized_keys File Format

The `$HOME/.ssh/authorized_keys` file lists the public keys that are permitted for public key authentication (`PubkeyAuthentication`) in protocol version 2. The `AuthorizedKeysFile` configuration option can be used to specify an alternative file.

Each line of the file contains one key (empty lines and lines starting with a hash mark `#` are ignored as comments).

For each RSA key for protocol version 1, the file consists of the following space-separated fields:

options bits exponent modulus comment

For the public key for protocol version 2, the file consists of the following space-separated fields:

options key-type base64-encoding-key comment

For protocol version 2, *key-type* is one of `ssh-rsa` or `ssh-dsa`.

The options field is optional; its presence is determined by whether the line starts with a number. (The option field never starts with a number.) The bits, exponent, and modulus fields give the RSA key; the comment field is a convenient place for you to identify the key.

Lines in this file are usually several hundred bytes long (because of the size of the key modulus). You will find it very inconvenient to type them in; instead, copy the public key file and edit it.

Permissions of this file must be set so that it is not world or group writable. See the `StrictModes` option of `sshd_config(4)`.

The options (if present) consist of comma-separated option specifications. No spaces are permitted, except within double quotes. The following option specifications are supported:

`from="pattern-list"`

Specifies that, in addition to public key authentication, the canonical name of the remote host must be present in the comma-separated list of patterns ('*' and '?' serve as wildcards). The list can also contain negated patterns by prefixing the patterns with '!'. If the canonical host name matches a negated pattern, the key is not accepted.

The purpose of this option is to give you the option of increasing security: public key authentication by itself does not trust the network or name servers or anything but the key. However, if someone manages to steal the key, possession of the key would permit the intruder to log in from anywhere in the world. This option makes using a stolen key more difficult, because name servers and routers would have to be compromised, in addition to just the key.

`command="command"`

Specifies that the *command* is executed whenever this key is used for authentication. The command supplied by the user (if any) is ignored. The command is run on a pty if the client requests a pty; otherwise it is run without a tty. If an 8-bit clean channel is required, one must not request a pty or should specify `no-pty`. You can include a quote in the command by escaping it with a backslash. This option might be useful to restrict certain public keys from performing a specific operation. An example is a key that permits remote backups but nothing else. Note that the client can specify TCP/IP and/or X11 forwarding unless they are explicitly prohibited from doing so. Also note that this option applies to shell, command, or subsystem execution.

`environment="NAME=value"`

Specifies that the string *NAME=value* is to be added to the environment when logging in using this key. Environment variables set this way override other default environment values. Multiple options of this type are permitted. Environment processing is disabled by default and is controlled via the `PermitUserEnvironment` option.

`no-port-forwarding`

Forbids TCP/IP forwarding when this key is used for authentication. Any port forward requests by the client will return an error. This might be used, for example, in connection with the `command` option.

no-X11-forwarding

Forbids X11 forwarding when this key is used for authentication. Any X11 forward requests by the client will return an error.

no-agent-forwarding

Forbids authentication agent forwarding when this key is used for authentication.

no-pty

Prevents `tty` allocation (a request to allocate a `pty` will fail).

permitopen="*host:port*"

Limit local `ssh -L` port forwarding such that it can connect only to the specified host and port. IPv6 addresses can be specified with an alternative syntax: *host/port*. You can invoke multiple `permitopen` options, with each instance separated by a comma. No pattern matching is performed on the specified hostnames. They must be literal domains or addresses.

ssh_known_hosts File Format

The `/etc/ssh/ssh_known_hosts` and `$HOME/.ssh/known_hosts` files contain host public keys for all known hosts. The global file should be prepared by the administrator (optional), and the per-user file is maintained automatically: whenever the user connects from an unknown host its key is added to the per-user file.

For the RSA key for protocol version 1, these files consist of the following space-separated fields:

hostnames bits exponent modulus comment

For the public key for protocol version 2, these files consist of the following space-separated fields:

hostnames key-type base64-encoding-key comment

For protocol version 2, *key-type* is one of `ssh-rsa` or `ssh-dsa`.

Hostnames is a comma-separated list of patterns (`*` and `?` act as wildcards); each pattern in turn is matched against the canonical host name (when authenticating a client) or against the user-supplied name (when authenticating a server). A pattern can also be preceded by `!` to indicate negation: if the host name matches a negated pattern, it is not accepted (by that line) even if it matched another pattern on the line.

Alternately, hostnames can be stored in a hashed form, which hides host names and addresses should the file's contents be disclosed. Hashed hostnames start with a vertical bar (`|`) character. Only one hashed hostname can appear on a single line and none of the above negation or wildcard operators may be applied.

Bits, exponent, and modulus are taken directly from the RSA host key; they can be obtained, for example, from `/etc/ssh/ssh_host_rsa_key.pub`. The optional comment field continues to the end of the line, and is not used.

Lines starting with a hash mark (`#`) and empty lines are ignored as comments.

When performing host authentication, authentication is accepted if any matching line has the proper key. It is thus permissible (but not recommended) to have several lines or different host keys for the same names. This will inevitably happen when short forms of host names from different domains are put in the file. It is possible that the files contain conflicting information; authentication is accepted if valid information can be found from either file.

The lines in these files are typically hundreds of characters long. You should definitely not type in the host keys by hand. Rather, generate them by a script or by taking `/etc/ssh/ssh_host_rsa_key.pub` and adding the host names at the front.

环境变量

sshd sets the following environment variables for commands executed by ssh users:

DISPLAY

Indicates the location of the X11 server. It is automatically set by sshd to point to a value of the form *hostname:n*, where *hostname* indicates the host where the shell runs, and *n* is an integer greater than or equal to 1. ssh uses this special value to forward X11 connections over the secure channel. Unless you have important reasons to do otherwise, you should not set DISPLAY explicitly, as that will render the X11 connection insecure and will require you to manually copy any required authorization cookies.

HOME

Set to the path of the user's home directory.

LANG, LC_ALL, LC_COLLATE, LC_CTYPE, LC_MESSAGES, LC_MONETARY, LC_NUMERIC, LC_TIME

A locale setting. The locale defaults to that of sshd (usually the system-wide default locale), or is negotiated between the client and server during initial key exchange (as per RFC 4253).

Following initial key exchange, each of the variables can be overridden in the following sequence:

1. If a locale setting is set in a client's environment and that client supports “Environment Variable Passing” (see RFC 4254), then the setting will be passed over to the server side.
2. If the public key authentication method was used to authenticate the server and the `PermitUserEnvironment` variable in `sshd_config(4)` is set to `yes` on the server side, then the setting can be changed through the use of the `environment` option in the client's `AuthorizedKeysFile` file.
3. The setting can be change in the client's `~/.ssh/environment` file on the server.

See `PermitUserEnvironment` in `sshd_config(4)` as to when the `AuthorizedKeysFile` and `~/.ssh/environment` files are processed and used for setting the user environment.

LOGNAME

Synonym for `USER`. Set for compatibility with systems that use this variable.

MAIL

Set to point to the user's mailbox.

SSH_AUTH_SOCKET

Indicates the path of a unix-domain socket used to communicate with the agent.

SSH_CONNECTION

Identifies the client and server ends of the connection. The variable contains four space-separated values: client IP address, client port number, server IP address and server port number.

SSH_CLIENT

Identifies the client end of the connection. The variable contains three space-separated values: client IP address, client port number, and server port number.

SSH_TTY

Set to the name of the `tty` (path to the device) associated with the current shell or command. If the current session has no `tty`, this variable is not set.

TZ

Indicates the present timezone, if `TIMEZONE` is set in `/etc/default/login` or if `TZ` was set when the daemon was started.

HZ

If set in `/etc/default/login`, the daemon sets it to the same value.

SHELL

The user's shell, if `ALTSHELL=YES` in `/etc/default/login`.

PATH

Set to the value of `PATH` or `SUPATH` (see [login\(1\)](#)) in `/etc/default/login`, or, if not set, to `/usr/bin:/bin`.

USER

Set to the name of the user logging in.

Additionally, `sshd` reads `$HOME/.ssh/environment` and adds lines of the format `VARNAME=value` to the environment.

示例

In the following examples, certain lines might wrap due to line length limits in your display. You should nevertheless consider the wrapped line as a single line.

示例 1 `authorized_key` File Entries

The following are examples of `authorized_key` file entries for protocol 1:

```
1024 33 12121...312314325 ylo@foo.bar
```

```
from="*.niksula.hut.fi,!pc.niksula.hut.fi" 1024 35 23...2334 ylo@niksula
```

```
command="dump /home",no-pty,no-port-forwarding 1024 33 23...2323 backup.hut.fi
```


示例 2 authorized_key File Entries for Protocol 2

The following are examples of authorized_key file entries for protocol 2:

```
ssh-rsa AAAAB3NzaC1y....EU88ovYKg4GfcLWGCFTuw8= ylo@foo.bar
from="*.niksula.hut.fi" ssh-rsa AAAAB3NzaC...uw8= ylo@niksula
command="dump /home",no-pty,no-port-forwarding ssh-rsa AA..8= backup.hut.fi
```

示例 3 ssh_known_hosts File Entries for Protocol 1

The following are examples of ssh_known_hosts file entries for protocol 1:

```
closenet,closenet.hut.fi,...,130.233.208.41 1024 37 159...93 closenet.hut.fi
```

示例 4 ssh_known_hosts File Entries for Protocol 2

The following are examples of ssh_known_hosts file entries for protocol 2:

```
closenet,closenet.hut.fi,...,130.233.208.41 ssh-rsa AA..8= closenet.hut.fi
```

示例 5 Using X.509 Public Key Authentication

The following example of user authentication should help you understand how the X.509 public key authentication works. Steps to use X.509 certificates/keys for host authentication are very similar. Note that we use the same token (softtoken) for both the TA certificate and the user certificate, which would not normally be the case.

Create a self-signed trusted anchor certificate and a user certificate signed by the trusted anchor private key. Configure the SSH daemon for certificate validation, and run the SSH client with the user certificate generated in the previous step as its user identity.

Replace user name “*PUT-YOUR-USERNAME-HERE*” with an existing (your) user name on the server side.

```
# Generate a trusted anchor certificate/key pair in the
# default token (softtoken).
pktool gencert keystore=pkcs11 label=authority \
    subject="CN=authority" serial=0x01

# Export the trusted anchor certificate (and, after that,
# put it to /etc/ssh/cert directory).
pktool export keystore=pkcs11 outfile=ta.cert objtype=cert
    label=authority outformat=pem

# Generate a user certificate signed by the trusted anchor
# private key and then import the certificate into the
# token.
pktool genscr keystore=pkcs11 label=user \
    outcsr=user.csr subject="CN=<PUT-YOUR-USERNAME-HERE>"
pktool signcsr keystore=pkcs11 signkey=authority \
    csr=user.csr serial=0x02 issuer="CN=authority" \
```

示例 5 Using X.509 Public Key Authentication (续)

```

    outcert=user.cert
pktool import keystore=pkcs11 infile=user.cert label=user

# Create a new policy file.
kmfcfg create dbfile=/etc/ssh/policy.xml policy=ssh \
    ta-name=search mapper-name=cn

# Start the SSH daemon in a debug mode to test the
# configuration. Use a different port so that you do not
# clash with the already running daemon.
/usr/lib/ssh/sshd -p 2222 -ddd -o \
    -o TrustedAnchorKeystore=/etc/ssh/cert \
    -o KMFPolicyDatabase=/etc/ssh/policy.xml \
    -o KMFPolicyName=ssh

# Run the client in a debug mode to test the configuration
# (will ask for a PIN for the token, use the 'pinfile'
# attribute if you do not want to set PIN for each
# invocation).
ssh -p 2222 -vvv -o \
    IdentityFile=\
    "pkcs11:object=user;token=Sun Software PKCS#11 softtoken" \
    <PUT-YOUR-USERNAME-HERE>@localhost

```

Before you use the softtoken you should set its PIN by means of the `setpin` subcommand. See [pktool\(1\)](#) for more information.

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

文件

/etc/default/login

Contains defaults for several `sshd_config` parameters, environment variables, and other environmental factors.

The following parameters affect environment variables (see [login\(1\)](#) and descriptions of these variables, above):

- TIMEZONE
- HZ
- ALTSHELL
- PATH
- SUPATH

The following `/etc/default/login` parameters supply default values for corresponding [sshd_config\(4\)](#) parameters:

- CONSOLE (see PermitRootLogin in [sshd_config\(4\)](#))
- PASSREQ (see PermitEmptyPasswords in [sshd_config\(4\)](#))
- TIMEOUT (see LoginGraceTime in [sshd_config\(4\)](#))

The following `/etc/default/login` parameters:

- UMASK
- ULIMIT

...set the [umask\(2\)](#) and file size limit of, respectively, the shells and commands spawned by `sshd`.

Finally, two `/etc/default/login` parameters affect the maximum allowed login attempts per-connection using interactive user authentication methods (for example, `keyboard-interactive` but not `publickey`), as per [login\(1\)](#):

- RETRIES
- SYSLOG_FAILED_LOGINS

`/etc/ssh/sshd_config`

Contains configuration data for `sshd`. This file should be writable by root only, but it is recommended (though not necessary) that it be world-readable.

`/etc/ssh/ssh_host_key`

`/etc/ssh/ssh_host_dsa_key`

`/etc/ssh/ssh_host_rsa_key`

Contains the private part of the host key. This file should only be owned by root, readable only by root, and not accessible to others. `sshd` does not start if this file is group/world-accessible.

`/etc/ssh/ssh_host_key.pub`

`/etc/ssh/ssh_host_dsa_key.pub`

`/etc/ssh/ssh_host_rsa_key.pub`

Contains the public part of the host key. This file should be world-readable but writable only by root. Its contents should match the private part. This file is not used for encryption; it is provided only for the convenience of the user so its contents can be copied to known hosts files. These two files are created using [ssh-keygen\(1\)](#).

`/var/run/sshd.pid`

Contains the process ID of the `sshd` listening for connections. If there are several daemons running concurrently for different ports, this contains the pid of the one started last. The content of this file is not sensitive; it can be world-readable. You can use the `PidFile` keyword in `sshd_config` to specify a file other than `/var/run/sshd.pid`. See [sshd_config\(4\)](#).

`/etc/ssh/ssh_known_hosts` and `$HOME/.ssh/known_hosts`

These files are consulted when using `rhosts` with public key host authentication to check the public key of the host. The key must be listed in one of these files to be accepted. The client uses the same files to verify that the remote host is the one it intended to connect.

These files should be writable only by root or the owner. `/etc/ssh/ssh_known_hosts` should be world-readable, and `$HOME/.ssh/known_hosts` can but need not be world-readable.

`/etc/nologin`

If this file exists, `sshd` refuses to let anyone except root log in. The contents of the file are displayed to anyone trying to log in, and non-root connections are refused. The file should be world-readable.

`$HOME/.ssh/authorized_keys`

Lists the public keys (RSA or DSA) that can be used to log into the user's account. This file must be readable by root. This might, on some machines, imply that it is world-readable if the user's home directory resides on an NFS volume. It is recommended that it not be accessible by others. The format of this file is described above. Users will place the contents of their `identity.pub`, `id_dsa.pub` and/or `id_rsa.pub` files into this file, as described in [ssh-keygen\(1\)](#).

`$HOME/.rhosts`

This file contains host-username pairs, separated by a space, one per line. The given user on the corresponding host is permitted to log in without password. The same file is used by `rlogin` and `rshd`. The file must be writable only by the user; it is recommended that it not be accessible by others. It is also possible to use `netgroups` in the file. Either host or user name may be of the form `+@groupname` to specify all hosts or all users in the group.

`$HOME/.shosts`

For `ssh`, this file is exactly the same as for `.rhosts`. However, this file is not used by `rlogin` and `rshd`, so using this permits access using SSH only.

`/etc/hosts.equiv`

This file is used during `.rhosts` authentication. In its simplest form, this file contains host names, one per line. Users on these hosts are permitted to log in without a password, provided they have the same user name on both machines. The host name can also be followed by a user name; such users are permitted to log in as any user on this machine (except root). Additionally, the syntax `+@group` can be used to specify `netgroups`. Negated entries start with a hyphen (-).

If the client host/user is successfully matched in this file, login is automatically permitted, provided the client and server user names are the same. Additionally, successful RSA host authentication is normally required. This file must be writable only by root; it is recommended that it be world-readable.

Warning: It is almost never a good idea to use user names in `hosts.equiv`. Beware that it really means that the named user(s) can log in as anybody, which includes `bin`, `daemon`, `adm`, and other accounts that own critical binaries and directories. For practical purposes, using a user name grants the user root access. Probably the only valid use for user names is in negative entries. This warning also applies to `rsh/rlogin`.

`/etc/ssh/moduli`

A private file.

`/etc/ssh/shosts.equiv`

This file is processed exactly as `/etc/hosts.equiv`. However, this file might be useful in environments that want to run both `rsh/rlogin` and `ssh`.

`$HOME/.ssh/environment`

This file is read into the environment at login (if it exists). It can contain only empty lines, comment lines (that start with `#`), and assignment lines of the form *name=value*. The file should be writable only by the user; it need not be readable by anyone else. Environment processing is disabled by default and is controlled by means of the `PermitUserEnvironment` option.

`$HOME/.ssh/rc`

If this file exists, it is run with `/bin/sh` after reading the environment files but before starting the user's shell or command. If X11 spoofing is in use, this will receive the proto cookie pair in standard input (and `DISPLAY` in environment). This must call `xauth` in that case.

The primary purpose of `$HOME/.ssh/rc` is to run any initialization routines that might be needed before the user's home directory becomes accessible; AFS is a particular example of such an environment. If this file exists, it is run with `/bin/sh` after reading the environment files, but before starting the user's shell or command. It must not produce any output on `stdout`; `stderr` must be used instead. If X11 forwarding is in use, it will receive the proto cookie pair in its standard input and `DISPLAY` in its environment. The script must call `xauth` because `sshd` will not run `xauth` automatically to add X11 cookies.

This file will probably contain some initialization code followed by something similar to:

```
if read proto cookie && [ -n "$DISPLAY" ]
then
  if [ 'echo $DISPLAY | cut -c1-10' = 'localhost:' ]
  then
    # X11UseLocalhost=yes
    echo add unix:'echo $DISPLAY |
    cut -c11-' $proto $cookie
  else
    # X11UseLocalhost=no
    echo add $DISPLAY $proto $cookie
  fi | xauth -q -
fi
```

If this file does not exist, `/etc/ssh/sshrc` is run, and if that does not exist, `xauth` is used to store the cookie. `$HOME/.ssh/rc` should be writable only by the user, and need not be readable by anyone else.

`/etc/ssh/sshrd`

Similar to `$HOME/.ssh/rc`. This can be used to specify machine-specific login-time initializations globally. This file should be writable only by root, and should be world-readable.

安全

sshd supports the use of several user authentication mechanisms: a public key system where keys are associated with users (through users' `authorized_keys` files), a public key system where keys are associated with hosts (see the `HostbasedAuthentication` configuration parameter), a GSS-API based method (see the `GssAuthentication` and `GssKeyEx` configuration parameters) and three initial authentication methods: `none`, `password`, and a generic prompt/reply protocol, `keyboard-interactive`.

sshd negotiates the use of the GSS-API with clients only if it has a GSS-API acceptor credential for the “host” service. This means that, for GSS-API based authentication, the server must have a Kerberos V keytab entry (see below) or the equivalent for any other GSS-API mechanism that might be installed.

In order for Kerberos authentication to work, a host/`<FQDN>` Kerberos principal must exist for each Fully Qualified Domain Name associated with the `in.sshd` server. Each of these host/`<FQDN>` principals must have a keytab entry in the `/etc/krb5/krb5.keytab` file on the `in.sshd` server. An example principal might be:

`host/bigmachine.eng.example.com`

See [kadmin\(1M\)](#) or [gkadmin\(1M\)](#) for instructions on adding a principal to a `krb5.keytab` file. See 《Oracle Solaris 管理：安全服务》 for a discussion of Kerberos authentication.

GSS-API authorization is covered in [gss_auth_rules\(5\)](#).

sshd uses [pam\(3PAM\)](#) for the three initial authentication methods as well as for account management, session management, and password management for all authentication methods.

Specifically, sshd calls `pam_authenticate()` for the “none,” “password” and “keyboard-interactive” SSHv2 `userauth` types. Other SSHv2 authentication methods do not call `pam_authenticate()`. `pam_acct_mgmt()` is called for each authentication method that succeeds.

`pam_setcred()` and `pam_open_session()` are called when authentication succeeds and `pam_close_session()` is called when connections are closed.

`pam_open_session()` and `pam_close_session()` are also called when SSHv2 channels with ptys are opened and closed.

Each SSHv2 `userauth` type has its own PAM service name:

SSHv2 Userauth	PAM Service Name
none	sshd-none
password	sshd-password
keyboard-interactive	sshd-kbdint
pubkey	sshd-pubkey
hostbased	sshd-hostbased
gssapi-with-mic	sshd-gssapi
gssapi-keyex	sshd-gssapi

If `pam_acct_mgmt()` returns `PAM_NEW_AUTHTOK_REQD` (indicating that the user's authentication tokens have expired), then `sshd` forces the use of “keyboard-interactive” userauth, if version 2 of the protocol is in use. The “keyboard-interactive” userauth will call `pam_chauthtok()` if `pam_acct_mgmt()` once again returns `PAM_NEW_AUTHTOK_REQD`. By this means, administrators are able to control what authentication methods are allowed for SSHv2 on a per-user basis.

Setting up Host-based Authentication

To establish host-based authentication, you must perform the following steps:

- Configure the client.
- Configure the server.
- Publish known hosts.
- Make appropriate entries in `/etc/ssh/shosts.equiv` and `~/.shosts`.

These steps are expanded in the following paragraphs.

- On a client machine, in the system-wide client configuration file, `/etc/ssh/ssh_config`, you must have the entry:


```
HostbasedAuthentication yes
```

See [ssh_config\(4\)](#) and [ssh-keysign\(1M\)](#).
- On the server, in the system-wide server configuration file, `/etc/ssh/sshd_config`, you must have the entry:


```
HostbasedAuthentication yes
```

If per-user `.shost` files are to be allowed (see last step), in the same file, you must have:

```
IgnoreRhosts no
```

See [sshd_config\(4\)](#) for a description of these keywords.
- To publish known hosts, you must have entries for the clients from which users will be allowed host-based authentication. Make these entries in either or both of the system-wide file (`/etc/ssh/ssh_known_hosts`) or the per-user file (`~/.ssh/known_hosts`).

- Note that sshd uses .shosts, not .rhosts. If you want the functionality provided by .rhosts, but do not want to use rlogin or rsh because of their security shortcomings, you can use .shosts in conjunction with sshd. To use this feature, make appropriate entries in /etc/ssh/shosts.equiv and ~/.shosts, in the format specified in [rhosts\(4\)](#).

For the vast majority of network environments, .shosts is preferred over .rhosts.

Using X.509 Certificates
as Host Keys and User
Identities

SSH commands can work with X.509v3 certificates when used as host keys and/or user identities. When used this way, PKCS#11 URIs can be used instead of plain filenames to locate keys and certificates. Host keys, user identities, and their corresponding certificates used for X.509 public key authentication can be stored only in PKCS#11 tokens. The only PKCS#11 URI attributes used by SSH are object, token, and pinfile. If no hardware keystore is available, the PKCS#11 softtoken keystore can be used.

For certificate validation (a server that authenticates a user, or a client that authenticates a server), the KMF policy needs a certificate mapper. Mappers are not set in the default policy file. See TrustedAnchorKeystore, KMFPolicyDatabase, and KMFPolicyName options in [sshd_config\(4\)](#), and see [libkmf\(3LIB\)](#) for more information.

Self-signed certificates can only be used as host keys, never as user identities. If a self-signed certificate is used or if a host certificate cannot be validated on the client side due to a missing relevant trusted anchor certificate, the SSH client offers to store the key in the known_host database, as is currently done with plain public keys. Only the public key is extracted from the certificate and stored in the database, thus maintaining the current format. The reason for accepting such a certificate on the client side is that the initial host Key Exchange is not restartable. Given this, if a server uses certificates as its only host key(s), a client that is unable to validate the certificate would otherwise be denied from logging in.

See the EXAMPLES section for guidance on how to use the X.509 authentication.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/ssh
Interface Stability	Committed

The interface stability of /etc/ssh/moduli is Private.

另请参见

[kmfcfg\(1\)](#), [login\(1\)](#), [pktool\(1\)](#), [scp\(1\)](#), [ssh\(1\)](#), [ssh-add\(1\)](#), [ssh-agent\(1\)](#), [ssh-keygen\(1\)](#), [svcs\(1\)](#), [gkadmin\(1M\)](#), [kadmin\(1M\)](#), [sftp-server\(1M\)](#), [ssh-keysign\(1M\)](#), [svcadm\(1M\)](#), [libkmf\(3LIB\)](#), [pam\(3PAM\)](#), [rhosts\(4\)](#), [ssh_config\(4\)](#), [sshd_config\(4\)](#), [attributes\(5\)](#), [gss_auth_rules\(5\)](#), [kerberos\(5\)](#), [pam_roles\(5\)](#), [pkcs11_softtoken\(5\)](#), [smf\(5\)](#)

《Oracle Solaris 管理：安全服务》

附注

The sshd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/ssh:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

sshd always sets PAM_RHOST and sets PAM_AUSER in the case of host-based userauth. This behavior allows for remote logins to roles using host-based authentication. See [pam_roles\(5\)](#).

引用名 ssh-keysign – ssh helper program for host-based authentication

用法概要 ssh-keysign

描述 ssh-keysign is used by [ssh\(1\)](#) to access the local host keys and generate the digital signature required during host-based authentication with SSH protocol version 2. This signature is of data that includes, among other items, the name of the client host and the name of the client user.

ssh-keysign is disabled by default and can be enabled only in the global client configuration file /etc/ssh/ssh_config by setting HostbasedAuthentication to yes.

ssh-keysign is not intended to be invoked by the user, but from ssh. See [ssh\(1\)](#) and [sshd\(1M\)](#) for more information about host-based authentication.

文件

/etc/ssh/ssh_config	Controls whether ssh-keysign is enabled.
/etc/ssh/ssh_host_dsa_key	These files contain the private parts of the host keys used to generate the digital signature. They should be owned by root, readable only by root, and not accessible to others. Because they are readable only by root, ssh-keysign must be set -uid root if host-based authentication is used.
/etc/ssh/ssh_host_rsa_key	

安全 ssh-keysign will not sign host-based authentication data under the following conditions:

- If the HostbasedAuthentication client configuration parameter is not set to yes in /etc/ssh/ssh_config. This setting cannot be overridden in users' ~/.ssh/ssh_config files.
- If the client hostname and username in /etc/ssh/ssh_config do not match the canonical hostname of the client where ssh-keysign is invoked and the name of the user invoking ssh-keysign.

In spite of ssh-keysign's restrictions on the contents of the host-based authentication data, there remains the ability of users to use it as an avenue for obtaining the client's private host keys. For this reason host-based authentication is turned off by default.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	network/ssh
Interface Stability	Committed

另请参见 [ssh\(1\)](#), [sshd\(1M\)](#), [ssh_config\(4\)](#), [attributes\(5\)](#)

Authors Markus Friedl, markus@openbsd.org

History ssh-keysign first appeared in OpenBSD 3.2.

引用名 statd – network status monitor

用法概要 /usr/lib/nfs/statd

描述 statd is an intermediate version of the status monitor. It interacts with [lockd\(1M\)](#) to provide the crash and recovery functions for the locking services on NFS. statd keeps track of the clients with processes which hold locks on a server. When the server reboots after a crash, statd sends a message to the statd on each client indicating that the server has rebooted. The client statd processes then inform the lockd on the client that the server has rebooted. The client lockd then attempts to reclaim the lock(s) from the server.

statd on the client host also informs the statd on the server(s) holding locks for the client when the client has rebooted. In this case, the statd on the server informs its lockd that all locks held by the rebooting client should be released, allowing other processes to lock those files.

lockd is started by [automountd\(1M\)](#), [mount_nfs\(1M\)](#), and [share\(1M\)](#) if NFS automounts are needed.

文件	/var/statmon/sm	lists hosts and network addresses to be contacted after a reboot
	/var/statmon/sm.bak	lists hosts and network addresses that could not be contacted after last reboot
	/var/statmon/state	includes a number which changes during a reboot
	/usr/include/rpcsvc/sm_inter.x	contains the rpcgen source code for the interface services provided by the statd daemon.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/file-system/nfs

另请参见 [svcs\(1\)](#), [automountd\(1M\)](#), [lockd\(1M\)](#), [mount_nfs\(1M\)](#), [share\(1M\)](#), [svcadm\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

《Oracle Solaris 管理：IP 服务》

附注 The crash of a server is only detected upon its recovery.

The statd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/network/nfs/status

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

If it is disabled, it will be enabled by [mount_nfs\(1M\)](#), [share_nfs\(1M\)](#), and [automountd\(1M\)](#) unless its `application/auto_enable` property is set to `false`.

引用名

用法概要

stclient – Service Tag Administration Program

```
stclient -x | -E [-r root_dir]
```

```
stclient -a [-i instance_URN] -p product_name -e product_version
-t product_URN [-F parent_URN] -P product_parent
[-I product_defined_instance_id] -m product_vendor -A platform_arch
-z container -S source [-r root_dir]
```

```
stclient -u -i instance_URN -F parent_URN -I product_defined_instance_id
[-r root_dir]
```

```
stclient -d -i instance_URN [-r root_dir]
```

```
stclient -g -i instance_URN [-r root_dir]
```

```
stclient -f -t product_URN [-r root_dir]
```

```
stclient -h
```

```
stclient -v
```

描述

The `stclient` command displays, finds, adds, updates and deletes records in the Service Tag registry. The registry is in the XML file `/var/sadm/servicetag/registry/servicetag.xml`, and contains the inventory of the product instances installed in the machine. Each record has a unique product instance identifier which is generated when the service tag is added in the registry. This product instance identifier is used as a key when finding, updating or deleting the service tag records. The `extract` option prints out the registry contents in XML format in `stdout`.

The `stclient` command also runs in interactive mode. This mode is invoked by running `stclient` without any parameters. A menu of all the available options are displayed, and the user is prompted to enter different parameters depending on the option chosen.

Any user can extract or get the contents of the registry, but only users with the appropriate privileges, the “svctag” user, or the creator of the service tag is authorized to update or delete a service tag record.

选项

The following options are supported:

-p <i>product_name</i>	Sets the product name of the service tag to be added. For example, “stclient -p “Solaris 10 Operating System”” sets the product name to the “Solaris 10 Operating System”.

<code>-e <i>product_version</i></code>	Sets the product version of the service tag to be added. For example, “ <code>stclient -e 5.10</code> ” sets the product version to “5.10”.

-t <i>product_URN</i>	Sets the Sun product unique identifier of the service tag to be added. For example, “stclient -t urn:uuid:5005588c-36f3-11d6-9cec-fc96f718e113” sets

	the Sun product unique identifier to “urn:uuid:5005588c-36f3-11d6-9cec-fc96f718e113”.
<i>-i instance_URN</i>	Sets the product instance unique identifier of the service tag. For example, “stclient -i 3a444ab2-1dd2-11b2-a69d-830020782a6b” sets the product instance unique identifier to “3a444ab2-1dd2-11b2-a69d-830020782a6b”. This field may also be set when the service tag is added but is not typically used. The generation of the instance ID is normally left to the stclient.
<i>-F product_parent_URN</i>	Sets the Sun product unique identifier of parent product. For example, “stclient -F urn:uuid:5005588c-36f3-11d6-9cec-fc96f718e113” sets the Sun product unique identifier of parent product to “urn:uuid:5005588c-36f3-11d6-9cec-fc96f718e113”.
<i>-P product_parent</i>	Sets the parent product of the service tag. For example, “stclient -P JES” sets the product parent to “JES”.
<i>-I product_defined_instance_id</i>	Sets the product defined instance identifier. For example, “stclient -I ser:abc-klm-000-7190” sets the product defined instance identifier to “ser:abc-klm-000-7190”.
<i>-m product_vendor</i>	Sets the vendor of the product. For example, “stclient -m Sun” sets the product vendor to “Sun”.
<i>-A platform_arch</i>	Sets the platform architecture that the product is running on. For example, “stclient -A sparc” sets the platform architecture to “sparc”.
<i>-z container</i>	Sets the container that the product is running on. For example, “stclient -z “global zone”” sets the container to “global zone”.
<i>-S source</i>	Sets the source of this service tag, naming the entity that created it.
<i>-r root_dir</i>	Sets the root directory where the command searches for the registry. The command searches /var/sadm/servicetag/registry by default. When this option is used, the command looks for the Registry in the specified root directory, (for example, “-r /home/user1” searches /home/user1/var/sadm/servicetag). This is typically used for testing.

操作数

Function Letters	The function portion of the key is specified by one of the following letters:
x	Extract. Extracts and prints the contents of the Service Tag registry in XML format. An alternate root directory may be specified for testing purposes.
a	Add. Adds a service tag in the registry. A unique product instance identifier may be supplied and is automatically generated if not provided. This field is returned by the command. The required fields for add are product name, product version, product URN, product parent, vendor, platform architecture, container and source.
u	Update. Updates a service tag in the registry using the product instance URN as the key. The parent URN and product defined instance id fields can be updated.
d	Delete. Deletes a service tag in the registry using the product instance URN as the key.
E	Extract Environment. Enumerates to standard output the environmental or “agent” information associated with the registered Service Tags on this system.
g	Get. Gets and prints a service tag from the registry using the product instance URN as the key.
f	Find. Finds and prints all the instance URNs of a given product URN.
h	Help. Displays the command options.
v	Version. Displays the version number of the command.

退出状态 The following exit values are returned:

- 0 Successful completion.
- > 0 An error occurred.

示例

示例 1 Adding a Service Tag in the Registry
To add a service tag, enter the following:

```
# stclient -a -p "Java Enterprise Web Server 6.0" -e 6.0 \  
-t urn:uuid:f2b8b59c-6eba-11d7-986f-9f5d47ec72fe \  
-P Java Enterprise Server -m Sun -A sparc -z global -S patch
```

The screen displays the following:

```
Java Enterprise Web Server 6.0 6.0 added  
Product instance URN=urn:st:7fc61914-1dd2-11b2-b992-830020782a6b
```

示例 2 Updating a Service Tag in the Registry
To update a service tag, enter the following:

示例 2 Updating a Service Tag in the Registry (续)

```
# stclient -u -i 7fc61914-1dd2-11b2-b992-830020782a6b \
-I urn:st:product.defined.id
```

The screen displays the following:

```
Service tag updated
```

示例 3 Extracting a Service Tag Registry

To extract a service tag, enter the following:

```
# stclient -x
```

The screen displays output similar to the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<registry urn="urn:uuid:1234ab-00e1-11b3-98737646873" version="1.0">
  <service_tag>
    .
    .
    .
  </service_tag>
</registry>
```

示例 4 Finding all product instances

To find all product instances, enter the following:

```
# stclient -f -t urn:uuid:f2b8b59c-6eba-11d7-986f-9f5d47ec72fe \
fc61914-1dd2-11b2-b992-830020782a6b
```

示例 5 Listing the Environmental Information

To list the environmental information associated with the registered Service Tags on this system, enter the following:

```
# stclient -E
```

The screen displays output similar to the following:

```
<?xml version="1.0" encoding="UTF-8"?>
<agent>
  <agent_urn>urn:st:af15ee62-0bb3-ef2d-fa96-85a11996cc71</agent_urn>
  .
  .
  .
  </system_info>
</agent>
```

环境变量 See [environ\(5\)](#) for descriptions of the following environment variables that affect the execution of `stclient`: `LANG`, `LC_ALL`, `LC_CTYPE`, `LC_MESSAGES`, and `NLSPATH`.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/management/service-tag
Interface Stability	Committed

另请参见 [in.stdiscover\(1M\)](#), [in.stlisten\(1M\)](#)

引用名	stmfadm – SCSI 目标模式框架命令行接口
用法概要	<pre> stmfadm add-hg-member -g, --group-name group-name group-member... stmfadm add-tg-member -g, --group-name group-name group-member... stmfadm add-view [-n, --lun logical-unit-number -t, --target-group group-name -h, --host-group group-name] lu-name stmfadm create-hg group-name stmfadm create-lu [-p, --lu-prop logical-unit-property=val -s, --size size] lu-file stmfadm create-tg group-name stmfadm delete-hg group-name stmfadm delete-lu lu-name stmfadm delete-tg group-name stmfadm import-lu lu-file stmfadm list-hg [-v] [host-group-name...] stmfadm list-tg [-v] [target-group-name...] stmfadm list-lu [-v] [lu-name...] stmfadm list-target [-v] [target-name...] stmfadm list-view -l, --lu-name lu-name [entry-name...] stmfadm list-state stmfadm modify-lu [-p, --lu-prop logical-unit-property=val -s, --size size, -f, --file] lu-name lu-file stmfadm online-lu lu-name stmfadm offline-lu lu-name stmfadm online-lu target-name stmfadm offline-lu target-name stmfadm remove-hg-member -g, --group-name group-name group-member... stmfadm remove-tg-member -g, --group-name group-name group-member... stmfadm remove-view -l, --lu-name lu-name entry-name </pre>

描述 stmfadm 命令用于在 SCSI 目标模式框架 (SCSI Target Mode Framework, STMF) 内配置逻辑单元。该框架与本手册页使用以下术语：

启动器

负责向 SCSI 目标和逻辑单元发出 SCSI I/O 命令的设备。

target

负责为逻辑单元接收 SCSI I/O 命令的设备。

逻辑单元

目标内负责执行 SCSI I/O 命令的设备。

逻辑单元号

目标内的逻辑单元的标识符。

启动器组

启动器组是指一个或多个可组合用于同一**视图**（请参见下文）的启动器的集合。一个启动器不能充当多个启动器组的成员。

目标组

目标组是指一个或多个在创建**视图**（请参见下文）时可采用相同方式进行处理 SCSI 目标端口的集合。一个特定 SCSI 启动器可以看到的逻辑单元集是由视图的组合集确定的。

每个逻辑单元都具有一组视图条目，每个视图条目指定一个目标组、主机组和一个 LUN。在通过该目标组进行连接时，来自该主机组的启动器能够识别并连接到使用指定 LUN 的该逻辑单元。您可以使用视图来限制特定启动器可以看到的逻辑单元集，并分配将使用的 LUN 集。

视图

视图定义了启动器组、目标组和逻辑单元号与指定逻辑单元的关联关系。添加到逻辑单元的任何视图条目不能与该逻辑单元中现有的视图条目相冲突。尝试复制任何指定启动器、目标与逻辑单元号的关联关系时，系统会将视图条目视为冲突条目。例如，逻辑单元 LU_0 包含与之关联的以下视图条目：

```
Logical Unit: LU_0
    View Entry: 0
        initiator group: HostA
        target group: All targets
        logical unit number: 32
```

如果您尝试了以下操作：

```
# stmf add-view -n 31 -h HostA LU_0
```

...此操作将返回一个错误并显示一条消息，指出视图条目与一个或多个现有视图条目相冲突。之所以会产生此冲突，是因为现有视图条目 0 已经为启动器组 HostA 将 LU_0 映射到了逻辑单元号 32。

子命令

stmfadm 命令支持下列子命令。

```
add-view [-n, --lun logical-unit-number -t, --target-group group-name -h, --host-group group-name] lu-name
```

将一个逻辑单元视图条目添加到逻辑单元 *lu-name*，其中 *lu-name* 是逻辑单元的 STMF 名称，如 list-lu 子命令所示。add-view 子命令为用户提供了针对逻辑单元的访问控制实施机制，并为包含启动器和目标的给定集合的逻辑单元提供了逻辑单

元号的分配方法。必须至少有一个应用的视图，然后才能将逻辑单元提供给启动器使用。每个视图条目都将分配一个条目名称，在 `remove-view` 和 `list-view` 子命令中可使用该名称来引用条目。

`add-view` 支持以下选项：

- n, --lun *logical-unit-number*
logical-unit-number 是为视图条目分配给逻辑单元的一个整数，其范围是 0-16383。如果未指定此选项，则将由 STMF 框架来分配逻辑单元号。
- t, --target-group *group-name*
group-name 是之前使用 STMF `create-tg` 子命令创建的目标组的名称。如果未指定此选项，将可通过所有目标使用此逻辑单元。
- h, --host-group *group-name*
group-name 是之前使用 STMF `create-hg` 子命令创建的主机组的名称。如果未指定此选项，此逻辑单元将可供登录到 STMF 框架的所有启动器使用。

`add-hg-member -g group-name group member...`

向主机组添加主机组成员。*group-name* 必须是使用 `create-hg` 子命令创建的一个现有组。*group member* 可以指定为 *name_type.name_value*，其中 *name_type* 可以为下列项之一：

```
wwn
iqn
eui
```

...*name_value* 是启动器名称的值。例如，要向主机组 HostA 添加光纤通道启动器端口全局名称 200000e08b909221，应使用的命令为：

```
# stmfadm add-hg-member -g HostA wwn.200000e08b909221
```

要向 HostA 添加名为 iqn.1986-03.com.sun:01.46f7e262 的 iSCSI 启动器节点成员，应使用的命令为：

```
# stmfadm add-hg-member -g HostA iqn.1986-03.com.sun:01.46f7e262
```

此外，还可以使用成员的 SCSI 名称字符串标识符来指定成员。要使用 SCSI 名称字符串标识符添加上述两个启动器，应使用的命令为：

```
# stmfadm add-hg-member -g HostA eui.200000e08b909221
# stmfadm add-hg-member -g HostA iqn.1986-03.com.sun:01.46f7e262
```

一个主机组的成员不能充当多个主机组的成员。

`add-tg-member -g group-name group member...`

向目标组添加目标组成员。*group-name* 必须是使用 `create-tg` 子命令创建的一个现有组。*group member* 可以指定为 *name_type.name_value*，其中 *name_type* 可以为下列项之一：

```
wwn
iqn
eui
```

`...name_value` 是目标名称的值。例如，要向目标组 `TG0` 添加光纤通道目标端口全局名称 `501000e092376af7`，应使用的命令为：

```
# stmfadm add-tg-member -g TG0 wwn.501000e092376af7
```

要向 `TG0` 添加名为 `iqn.1986-03.com.sun:target.01.01110` 的 iSCSI 目标成员，应使用的命令为：

```
# stmfadm add-tg-member -g TG0 iqn.1986-03.com.sun:target.01.01110
```

此外，还可以使用成员的 SCSI 名称字符串标识符来指定成员。要使用 SCSI 名称字符串标识符添加上述两个目标，应使用的命令为：

```
# stmfadm add-tg-member -g TG0 eui.501000e092376af7
```

```
# stmfadm add-tg-member -g TG0 iqn.1986-03.com.sun:target.01.01110
```

一个目标组的成员不能充当多个目标组的成员。

`create-hg group-name`

创建名为 `group-name` 的启动器组。`group-name` 是最大长度为 255 个 Unicode 字符的字符串。组名在 STMF 系统内必须是唯一的。

`create-lu [-p, --lu-prop logical-unit-property=val --s, --size size] lu-file`

创建一个可以向 STMF 注册的逻辑单元。对于 `-p` 选项，`logical-unit-property` 可以是以下属性之一：

`alias`

最多为 255 个字符，表示设备的用户自定义名称。缺省值为后备存储的名称。

`blk`

指定设备的块大小。缺省值为 512，但是不要将该值设置为高于 4096 KB。

`guid`

表示有效的 NAA 注册扩展标识符，为三十二位十六进制 ASCII 字符。缺省值由 STMF 设置为一个生成的值。

`host-id`

表示分配的主机 ID 的八位十六进制 ASCII 字符。这将用来为逻辑单元生成全局唯一标识符 (Globally Unique Identifier, GUID)。缺省值是由 `hostid(1)` 返回的值。

`meta`

元数据文件名。该属性在指定后将用于保存逻辑单元的 SCSI 元数据。无缺省值。

`mgmt-url`

最多为 1024 个字符，表示管理网络地址 URL。通过在一对引号 (") 内包含以空格分隔的 URL，可将多个 URL 作为一个参数进行传递。

`oui`

组织唯一标识符。表示分配的 IEEE OUI 公司 ID 的六位十六进制 ASCII 字符。这将用来生成设备标识符 (GUID)。缺省值为 `00144F`。

pid

符合 SCSI SPC-3 规范的十六字符产品标识。此值将反映在为设备返回的标准 INQUIRY 数据中。缺省值为 COMSTAR。

serial

序列号, 指定 SCSI 重要产品数据序列号（第 80h 页）。这是一个最大长度为 252 字节的字符值。没有缺省值。

vid

符合 SCSI SPC-3 规范的八字符供应商标识。此值将反映在为设备返回的标准 INQUIRY 数据中。缺省值为 SUN。

wcd

禁用回写高速缓存。指定 **true** 或 **false** 以确定回写高速缓存禁用行为。缺省值是由 *lu-file* 参数指定的后备存储设备的回写高速缓存设置。

wp

写保护位。指定 **true** 或 **false** 以确定设备是否报告为有写保护。缺省值是 **/**。

对于 **-s** 选项, *size* 是一个整数, 后跟下列字母之一, 表示大小的单位:

k 千字节
m 兆字节
g 千兆字节
t 兆兆字节
p 千兆兆字节
e 艾字节

lu-file 是用作逻辑单元的后备存储的文件。如果未指定 **-s** 选项, 将使用指定的 *lu-file* 的大小作为逻辑单元的大小。向 STMF 注册的逻辑单元需要空间来存储元数据。当将 *zvol* 指定为后备存储设备时, 缺省情况下将使用 *zvol* 的一个特殊属性来包含元数据。对于所有其他设备, 缺省行为是使用设备的前 64k 空间。另一种替代方法是在 **create-lu** 命令中使用 *meta* 属性指定其他文件来包含元数据。建议使用可为逻辑单元元数据提供足够空间的文件, 最好为 64k。

create-tg *group-name*

创建名为 *group-name* 的目标组。*group-name* 是最大长度为 255 个 Unicode 字符的字符串。组名在 STMF 系统内必须是唯一的。

delete-hg *group-name*

删除由 *group-name* 标识的主机组。

delete-lu *lu-name*

删除使用 **stmfadm create-lu** 创建的现有逻辑单元。这可以有效地从 STMF 框架中卸载逻辑单元。逻辑单元上的任何现有数据将保持不变。

`delete-tg group-name`

删除由 *group-name* 标识的目标组。

`import-lu lu-file`

将之前使用 `stmfadm create-lu` 创建且后来又使用 `stmfadm delete-lu` 从 STMF 中删除的逻辑单元导入并装入到 STMF 中。成功后，该逻辑单元将重新可供 STMF 使用。*lu-file* 是在 `stmfadm create-lu` 命令中使用的文件名。如果此逻辑单元使用一个单独的元数据文件，则在此处必须使用在 `stmfadm create-lu` 命令中使用的 *meta* 属性值中的文件名。

`list-hg [-v, --verbose] [host-group-name...]`

列出系统中 *host-group-name* 所指的主机组的相关信息。如果未指定 *host-group-name*，则会列出系统中的所有主机组。如果指定了 `--v` 或 `--verbose` 选项，则会显示主机组内的所有成员。

`list-lu [-v, --verbose] [lu-name...]`

列出系统中 *lu-name* 所指的逻辑单元的相关信息。如果未指定 *lu-name*，则会列出系统中的所有逻辑单元。如果指定了 `-v` 或 `--verbose` 选项，则会显示关于逻辑单元的更多信息。

`list-target [-v, --verbose] [target-name...]`

列出系统中 *target-name* 所指的目标端口的相关信息。如果未指定目标名称，则会列出系统中的所有目标端口。如果指定了 `-v` 或 `--verbose` 选项，则会显示关于目标的更多信息以及已登录启动器的 SCSI 会话信息。

`list-tg [-v, --verbose] [target-group-name...]`

列出系统中 *target-group-name* 所指的目标组的相关信息。如果未指定 *target-group-name*，则会列出系统中的所有目标组。如果指定了 `--v` 或 `--verbose` 选项，则会显示目标组内的所有成员。

`list-view --l, --lu-name lu-name [entry-name...]`

为 *lu-name* 所指的逻辑单元列出视图条目。如果未指定 *entry-name*，则会为指定的逻辑单元列出所有视图条目。

`modify-lu [-p, --lu-prop logical-unit-property=val --s, --size size, -f, --file]`

lu-name|lu-file

修改使用 `stmfadm create-lu` 命令创建的逻辑单元的属性。对于 `-p` 选项，*logical-unit-property* 可以是以下属性之一：

alias

最多为 255 个字符，表示设备的用户自定义名称。缺省值为后备存储的名称。

mgmt-url

最多为 1024 个字符，表示管理网络地址 URL。通过在一对引号 (") 内包含以空格分隔的 URL，可将多个 URL 作为一个参数进行传递。

wcd

禁用回写高速缓存。指定 `true` 或 `false` 以确定回写高速缓存禁用行为。缺省值是由 *lu-file* 参数指定的后备存储设备的回写高速缓存设置。

wp

写保护位。指定 **true** 或 **false** 以确定设备是否报告为有写保护。缺省值是 **/**。

对于 **-s** 选项, *size* 是一个整数, 后跟下列字母之一, 表示大小的单位:

k 千字节
m 兆字节
g 千兆字节
t 兆兆字节
p 千兆兆字节
e 艾字节

lu-name 是要修改的逻辑单元的 **guid**。如果指定了 **-f** 选项, 此操作数将解释为文件名。这提供了修改当前未导入到 STMF 中的逻辑单元的功能。

online-lu *lu-name*

使当前已向 STMF 注册的某个逻辑单元联机。

online-target *target-name*

使指定的目标联机。

offline-lu *lu-name*

使当前已向 STMF 注册的某个逻辑单元脱机。

offline-target *target-name*

使指定的目标脱机。

list-state

列出 STMF 的运行状态和配置状态。

remove-hg-member -g *group-name* *group member*

从主机组中删除主机组成员。*group-name* 必须是使用 **create-hg** 子命令创建的一个现有组。*group member* 可以指定为 *name_type.name_value*, 其中 *name_type* 可以为下列项之一:

wwn
iqn
eui

...name_value 是启动器名称的值。例如, 要从 HostA 主机组中删除光纤通道启动器端口全局名称 **200000e08b909221**, 应使用的命令为:

```
# stmfadm remove-hg-member -g HostA wwn.200000e08b909221
```

要从 HostA 删除名为 **iqn.1986-03.com.sun:01.46f7e262** 的 iSCSI 启动器节点成员, 应使用的命令为:

```
# stmfadm remove-hg-member -g HostA iqn.1986-03.com.sun:01.46f7e262
```

此外，还可以使用成员的 SCSI 名称字符串标识符来指定成员。要使用 SCSI 名称字符串标识符删除上述两个启动器，应使用的命令为：

```
# stmfadm remove-hg-member -g HostA eui.200000e08b909221
# stmfadm remove-hg-member -g HostA iqn.1986-03.com.sun:01.46f7e262
```

remove-tg-member -g *group-name* *group member*

从目标组中删除目标组成员。*group-name* 必须是使用 **create-tg** 子命令创建的一个现有组。*group member* 可以指定为 *name_type.name_value*，其中 *name_type* 可以为下列项之一：

```
wwn
iqn
eui
```

...*name_value* 是目标名称的值。例如，要从目标组 TG0 中删除光纤通道目标端口全局名称 501000e092376af7，应使用的命令为：

```
# stmfadm remove-tg-member -g TG0 wwn.501000e092376af7
```

要从 TG0 中删除名为 iqn.1986-03.com.sun:target.01.01110 的 ISCSI 目标成员，应使用的命令为：

```
# stmfadm remove-tg-member -g TG0 iqn.1986-03.com.sun:target.01.01110
```

此外，还可以使用成员的 SCSI 名称字符串标识符来指定成员。要使用 SCSI 名称字符串标识符删除上述两个目标，应使用的命令为：

```
# stmfadm remove-tg-member -g TG0 eui.501000e092376af7
# stmfadm remove-tg-member -g TG0 iqn.1986-03.com.sun:target.01.01110
```

remove-view --l, --lu-name *lu-name* *entry-name*

从逻辑单元中删除一个或多个逻辑单元视图条目。

示例

示例1 创建具有两个启动器端口的主机组

以下命令使用 **create-hg** 和 **add-hg-member** 子命令创建一个主机组并向该主机组添加两个启动器端口。

```
# stmfadm create-hg host-group-a
# stmfadm add-hg-member -g host-group-a wwn.210105b0000d92d0
```

示例2 向逻辑单元添加视图条目

以下命令使用了 **add-view** 子命令，因此可以从 host-group-a 访问某个逻辑单元。

```
# stmfadm add-view -h host-group-a 6000AE40C5000000000046FC4FEA001C
```

示例3 列出视图条目

以下命令使用 **list-view** 子命令为指定的逻辑单元列出所有视图条目。

```
# stmfadm list-view -l 6000AE40C5000000000046FC4FEA001C
View Entry: 0
            Host group      : host-group-a
```

示例 3 列出视图条目 (续)

```
Target group : All
LUN          : 0
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/storage/scsi-target-mode-framework
接口稳定性	Committed（已确定）

另请参见 [hostid\(1\)](#)、[sbdadm\(1M\)](#)、[attributes\(5\)](#)

引用名	stmsboot – administration program for the Solaris I/O multipathing feature
用法概要	<pre>/usr/sbin/stmsboot [[-D (fp mpt mpt_sas iscsi)] -d -e -u] -L -l controller_number]</pre>
描述	The Solaris I/O multipathing feature is a multipathing solution for storage devices that is part of the Solaris operating environment. This feature was formerly known as Sun StorEdge Traffic Manager (STMS) or MPxIO. The stmsboot program is an administrative command to manage enumeration of multipath-capable devices with Solaris I/O multipathing. Solaris I/O multipathing-enabled devices are enumerated under scsi_vhci(7D), providing multipathing capabilities. Solaris I/O multipathing-disabled devices are enumerated under the physical controller. In the /dev and /devices trees, Solaris I/O multipathing-enabled devices receive new names that indicate that they are under Solaris I/O multipathing control. This means a device will have a different name from its original name (after enabling) when it is under Solaris I/O multipathing control. The stmsboot command automatically updates /etc/vfstab and dump configuration to reflect the device names changes when enabling or disabling Solaris I/O multipathing. One reboot is required for changes to take effect.
选项	The following options are supported: -e [-E fp mpt mpt_sas iscsi] Enables Solaris I/O multipathing on all supported multipath-capable controller ports, including fp(7d), mpt(7D), mpt_sas(7D), and iscsi(7D) port drivers. Multipath-capable ports include fibre channel (fp(7d)) controller ports and SAS (mpt(7D) or mpt_sas(7D)) controller ports. Following this enabling, you are prompted to reboot. During the reboot, vfstab and the dump configuration will be updated to reflect the device name changes. Specifying -D mpt, -D mpt_sas, or -D fp limits the enabling operation to ports attached using the specified driver. -d [-D fp mpt mpt_sas iscsi] Disables Solaris I/O multipathing on all supported multipath-capable controller ports, including fp(7d), mpt(7D), mpt_sas(7D), and iscsi(7D) port drivers. Multipath-capable ports include fibre channel (fp(7d)) controller ports and SAS (mpt(7D) or mpt_sas(7D)) controller ports. Following this disabling, you are prompted to reboot. During the reboot, vfstab and the dump configuration will be updated to reflect the device name changes. Specifying -D mpt, -D mpt_sas, or -D fp limits the disabling operation to ports attached using the specified driver. -u [-D fp mpt mpt_sas iscsi] Updates vfstab and the dump configuration after you have manually modified the configuration to have Solaris I/O multipathing enabled or disabled on specific fp(7d), mpt(7D), mpt_sas(7D), and iscsi(7D) controller ports. This option prompts you to reboot. During the reboot, vfstab and the dump configuration will be updated to reflect the device name changes.

-L

Display the device name changes from non-Solaris I/O multipathing device names to Solaris I/O multipathing device names for multipath-enabled controller ports. If Solaris I/O multipathing is not enabled, then no mappings are displayed.

-l *controller_number*

Display the device name changes from non-Solaris I/O multipathing device names to Solaris I/O multipathing device names for the specified controller. If Solaris I/O multipathing is not enabled, then no mappings are displayed.

Note that [mpt_sas\(7D\)](#) has MPxIO turned on by default. This means that when using the -L or -l option with -D mpt_sas, stmsboot does not display any non-multipathed and multipathed device names.

用法

The primary function of stmsboot is to control the enabling and disabling of Solaris I/O multipathing on the host. The utility automatically updates [vfstab\(4\)](#) and [dumpadm\(1M\)](#) configuration to reflect device name changes. The system administrator is responsible for modifying application configuration (for example, backup software, DBMS, and so forth) to reflect updated device names.

The -L and -l options display the mapping between multipathed and non-multipathed device names. These options function only after changes to the Solaris I/O multipathing configuration have taken effect, that is, following the reboot after invoking stmsboot -e.

ZFS datasets, including ZFS root datasets, are correctly handled by stmsboot.

示例

示例1 Enabling Solaris I/O Multipathing

To enable Solaris I/O multipathing for all multipath-capable controllers, run:

```
# stmsboot -e
```

To enable Solaris I/O multipathing on multipath-capable [mpt\(7D\)](#) controller ports, enter:

```
# stmsboot -D mpt -e
```

To enable Solaris I/O multipathing on multipath-capable [mpt_sas\(7D\)](#) controller ports, enter:

```
# stmsboot -D mpt_sas -e
```

To enable Solaris I/O Multipathing on multipath-capable fibre channel controller ports, enter:

```
# stmsboot -D fp -e
```

To enable Solaris I/O Multipathing on multipath-capable iSCSI controller ports, enter:

```
# stmsboot -D iscsi -e
```

示例 2 Disabling Solaris I/O Multipathing

To disable Solaris I/O multipathing on all multipath-capable controllers, enter:

```
# stmsboot -d
```

To disable Solaris I/O multipathing on multipath-capable [mpt\(7D\)](#) controller ports, enter:

```
# stmsboot -D mpt -d
```

To disable Solaris I/O multipathing on multipath-capable [mpt_sas\(7D\)](#) controller ports, enter:

```
# stmsboot -D mpt_sas -d
```

To disable Solaris I/O multipathing on multipath-capable iSCSI controller ports, enter:

```
# stmsboot -D iscsi -d
```

To disable Solaris I/O multipathing on multipath-capable fibre channel controller ports, enter:

```
# stmsboot -D fp -d
```

示例 3 Enabling Solaris I/O Multipathing on Selected Ports

To enable Solaris I/O multipathing on specific fibre channel controller ports and disable the feature on others, manually edit the `/kernel/drv/fp.conf` file. (See [fp\(7d\)](#).) The following command will update [vfstab\(4\)](#) and [dumpadm\(1M\)](#) configurations to reflect the changed device names:

```
# stmsboot -u
```

A similar procedure involving the `/kernel/drv/mpt.conf` file should be followed for devices attached by means of the [mpt\(7D\)](#) driver. For devices attached by means of the [iscsi\(7D\)](#) driver, follow a similar procedure that uses the `/kernel/drv/iscsi.conf` file.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os, system/library
Interface Stability	Obsolete

另请参见

[dumpadm\(1M\)](#), [fsck\(1M\)](#), [mpathadm\(1M\)](#), [ufsdump\(1M\)](#), [zfs\(1M\)](#), [zpool\(1M\)](#), [ufsdump\(4\)](#), [vfstab\(4\)](#), [emlxs\(7D\)](#), [fcp\(7D\)](#), [fp\(7d\)](#), [iscsi\(7D\)](#), [mpt\(7D\)](#), [mpt_sas\(7D\)](#), [qlc\(7D\)](#), [scsi_vhci\(7D\)](#)

Solaris SAN Configuration and Multipathing Guide (see <http://docs.oracle.com>)

Consult a particular storage product's system administrator's guide and release notes for further information specific to that product.

附注	Solaris I/O multipathing is not supported on all devices. After enabling Solaris I/O multipathing, only supported devices are placed under Solaris I/O multipathing control. Non-supported devices remain unchanged. For Solaris releases prior to the current release, the <code>-e</code> and <code>-d</code> options replace <code>mpxio-disable</code> property entries with a global <code>mpxio-disable</code> entry in <code>fp.conf</code>.
Enabling Solaris I/O Multipathing on a Sun StorEdge Disk Array	The following applies to Sun StorEdge T3, 3910, 3960, 6120, and 6320 storage subsystems. To place your Sun StorEdge disk subsystem under Solaris I/O multipathing control, in addition to enabling Solaris I/O multipathing, the <code>mp_support</code> of the subsystem must be set to <code>mpxio</code> mode. The preferred sequence is to change the subsystem's <code>mp_support</code> to <code>mpxio</code> mode, then run <code>stmsboot -e</code>. If Solaris I/O multipathing is already enabled but the subsystem's <code>mp_support</code> is not in <code>mpxio</code> mode, then change the <code>mp_support</code> to <code>mpxio</code> mode and run <code>stmsboot -u</code>. Refer to the Sun StorEdge Administrator's Guide for your subsystem for more details.
Using ufsdump	The ufsdump(1M) command records details of filesystem dumps in <code>/etc/dumpdates</code> (see ufsdump(4)). Among other items, the entries contain device names. An effect of the “active” <code>stmsboot</code> options (<code>-e</code>, <code>-d</code>, and <code>-u</code>) is to change the device name of a storage device. Because <code>stmsboot</code> does not modify <code>dumpdates</code>, entries will refer to obsolete device names, that is, device names that were in effect before Solaris I/O multipathing configuration changes were performed. In this situation <code>ufsdump</code> will behave as if no previous dump of the filesystem had been performed. A level 0 dump will be performed.
Procedure to Use stmsboot in Conjunction with Sun Cluster	If possible, invoke <code>stmsboot -e</code> before installing Sun Cluster software. After executing <code>stmsboot</code>, install Sun Cluster software normally. If Sun Cluster software is installed before executing <code>stmsboot</code>, follow this procedure: On each machine in the cluster where Solaris I/O multipathing is required, execute: <pre># stmsboot -e</pre> ...and allow the system to reboot. When the system comes up, enter the following two commands: 1. <code># /usr/cluster/bin/scdidadm -C</code> 2. <code># /usr/cluster/bin/scdidadm -r</code> The preceding commands update <code>did</code> mappings with new device names while preserving <code>did</code> instance numbers for disks that are connected to multiple cluster nodes. <code>did</code> instance numbers of the local disks might not be preserved. For this reason, the <code>did</code> disk names for local disks might change. The remaining steps are required only if local <code>did</code> number changes require editing of <code>/etc/vfstab</code>. If such editing is not required, run <code>/usr/cluster/bin/scgdevs</code> from each node in the cluster to complete the procedure.

3. Update `/etc/vfstab` to reflect any new did disk names for your local disks.
4. Reboot the system.

To disable the Solaris multipathing feature, use `stmsboot -d` (instead of `stmsboot -e`), then follow the procedure above.

To view mappings between the old and new device names, run `stmsboot -L`. To view did device name mappings, run `/usr/cluster/bin/scdidadm -L`.

With active-passive storage arrays, it is possible that while your host is rebooting the array controller could failover the path that a particular target is using. In this scenario, [fsck\(1M\)](#) will fail to open the physical path listed in `/etc/vfstab`. The `svc:/system/filesystem/local:default` SMF service will transition to a maintenance state as a result. To rectify this, consult the documentation for your storage array to failback the path. The [mpathadm\(1M\)](#) can assist with determining the active and passive path(s).

Limitations

On x86 platforms, the current Solaris release does not support disabling Solaris I/O multipathing of boot devices attached by means of fibre channel. Solaris I/O multipathing is always enabled for supported fibre channel-attached boot devices. Disabling Solaris I/O multipathing in this situation must be performed on a per-port basis. See [fp\(7d\)](#).

Executing `devfsadm -C` removes obsolete device entries that `stmsboot` relies on. This will prevent correct operation of the `-d` option for boot devices (regardless of platform type) and the `-L` option.

引用名

strace – print STREAMS trace messages

用法概要

strace [*mid sid level*]. . .

描述

strace without arguments writes all STREAMS event trace messages from all drivers and modules to its standard output. These messages are obtained from the STREAMS log driver (see [log\(7D\)](#)). If arguments are provided, they must be in triplets of the form *mid*, *sid*, *level*, where *mid* is a STREAMS module ID number, *sid* is a sub-ID number, and *level* is a tracing priority level. Each triplet indicates that tracing messages are to be received from the given module/driver, sub-ID (usually indicating minor device), and priority level equal to, or less than the given level. The token `all` may be used for any member to indicate no restriction for that attribute.

The format of each trace message output is:

<*seq*> <*time*> <*ticks*> <*level*> <*flags*> <*mid*> <*sid*> <*text*>

<*seq*> trace sequence number

<*time*> time of message in *hh:mm:ss*

<*ticks*> time of message in machine ticks since boot

<*level*> tracing priority level

<*flags*> E : message is also in the error log
F : indicates a fatal error
N : mail was sent to the system administrator (hardcoded as root)

<*mid*> module ID number of source

<*sid*> sub-ID number of source

<*text*> formatted text of the trace message

Once initiated, strace will continue to execute until terminated by the user.

示例

示例 1 A sample output of the strace command:

The following example outputs all trace messages from the module or driver whose module ID is 41:

```
strace 41 all all
```

The following example outputs those trace messages from driver or module ID 41 with sub-IDs 0, 1, or 2:

```
strace 41 0 1 41 1 1 41 2 0
```

Messages from sub-IDs 0 and 1 must have a tracing level less than or equal to 1. Those from sub-ID 2 must have a tracing level of 0.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[attributes\(5\)](#), [log\(7D\)](#)
[《STREAMS Programming Guide》](#)

附注

- There is no restriction to the number of `strace` processes opening the STREAMS log driver at a time.
- The log-driver records the list of the triplets specified in the command invocation, and compares each potential trace message against this list to decide if it should be formatted and sent up to the `strace` process. Hence, long lists of triplets will have a greater impact on overall STREAMS performance. Running `strace` will have the most impact on the timing of the modules and drivers generating the trace messages that are sent to the `strace` process. If trace messages are generated faster than the `strace` process can handle them, some of the messages will be lost. This last case can be determined by examining the sequence numbers on the trace messages output.

引用名	strclean – STREAMS error logger cleanup program				
用法概要	strclean [-a <i>age</i>] [-d <i>logdir</i>]				
描述	strclean is used to clean up the STREAMS error logger directory on a regular basis (for example, by using cron. By default, all files with names matching <code>error.*</code> in <code>/var/adm/streams</code> that have not been modified in the last three days are removed.				
选项	The following options are supported: -a <i>age</i> The maximum age in days for a log file can be changed using the -a option. -d <i>logdir</i> A directory other than <code>/var/adm/streams</code> can be specified using the -d option.				
示例	示例 1 A sample of using the strclean command. This example has the same result as running strclean with no arguments: <pre>example% strclean -d /var/adm/streams -a 3</pre>				
文件	<code>/var/adm/streams/error.*</code>				
属性	See attributes(5) for descriptions of the following attributes:				
<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>system/core-os</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	system/core-os
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	system/core-os				
另请参见	cron(1M) , strerr(1M) , attributes(5) 《STREAMS Programming Guide》				
附注	strclean is typically run from cron on a daily or weekly basis.				

引用名 strerr – STREAMS error logger daemon

用法概要 strerr

描述

strerr receives error log messages from the STREAMS-based log driver (see [log\(7D\)](#)) and appends them to a log file. The resultant error log files reside in the directory `/var/adm/streams`, and are named `error.mm-dd`, where *mm* is the month and *dd* is the day of the messages contained in each log file.

The format of an error log message is:

<seq> <time> <ticks> <flags> <mid> <sid> <text>

<seq> error sequence number

<time> time of message in hh:mm:ss

<ticks> time of message in machine ticks since boot priority level

<flags> T : the message was also sent to a tracing process F : indicates a fatal error N : send mail to the system administrator (hardcoded as root)

<mid> module ID number of source

<sid> sub-ID number of source

<text> formatted text of the error message

Messages that appear in the error log are intended to report exceptional conditions that require the attention of the system administrator. Those messages which indicate the total failure of a STREAMS-based driver or module should have the F flag set. Those messages requiring the immediate attention of the administrator will have the N flag set, which causes the error logger to send the message to the system administrator using `mail`. The priority level usually has no meaning in the error log but will have meaning if the message is also sent to a tracer process.

Once initiated, `strerr` continues to execute until terminated by the user. It is commonly executed asynchronously.

文件 `/var/adm/streams/error.mm-dd` error log file.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [attributes\(5\)](#), [log\(7D\)](#)

《STREAMS Programming Guide》

附注

There is no restriction to the number of `strerr` processes opening the STREAMS-based log driver at a time.

If a module or driver is generating a large number of error messages, running the error logger will cause a degradation in STREAMS performance. If a large burst of messages are generated in a short time, the log driver may not be able to deliver some of the messages. This situation is indicated by gaps in the sequence numbering of the messages in the log files.

引用名	sttydefs – maintain line settings and hunt sequences for TTY ports
用法概要	<pre> /usr/sbin/sttydefs -a ttylabel [-b] [-f final-flags] [-i initial-flags] [-n nextlabel] /usr/sbin/sttydefs -l [ttylabel] /usr/sbin/sttydefs -r ttylabel </pre>
描述	sttydefs is an administrative command that maintains the line settings and hunt sequences for the system's TTY ports by making entries in, and deleting entries from the <code>/etc/ttydefs</code> file. sttydefs with a <code>-a</code> or <code>-r</code> option may be invoked only by the super-user. sttydefs with <code>-l</code> may be invoked by any user on the system.
选项	The following options are supported: <code>-a ttylabel</code> Add a record to the <code>tttydefs</code> file, using <i>ttylabel</i> as its label. The following describes the effect of the <code>-b</code>, <code>-n</code>, <code>-i</code>, or <code>-f</code> options when used in conjunction with the <code>-a</code> option: <code>-b</code> Enable autobaud. Autobaud allows the system to set the line speed of a given TTY port to the line speed of the device connected to the port without the user's intervention. <code>-f final-flags</code> Specify the value to be used in the <i>final-flags</i> field in <code>/etc/ttydefs</code>. <i>final-flags</i> must be in a format recognized by the <code>stty</code> command. <i>final-flags</i> are the termio(7I) settings used by <code>ttymon</code> after receiving a successful connection request and immediately before invoking the service on the port. If this option is not specified, sttydefs will set <i>final-flags</i> equal to the termio(7I) flags <code>9600</code> and <code>sane</code>. <code>-i initial-flags</code> Specify the value to be used in the <i>initial-flags</i> field in <code>/etc/ttydefs</code>. <i>initial-flags</i> must be in a format recognized by the <code>stty</code> command. These flags are used by <code>ttymon</code> when searching for the correct baud rate. They are set prior to writing the prompt. If this option is not specified, sttydefs will set <i>initial-flags</i> equal to the termio(7I) flag <code>9600</code>. <code>-n nextlabel</code> Specify the value to be used in the <i>nextlabel</i> field in <code>/etc/ttydefs</code>. If this option is not specified, sttydefs will set <i>nextlabel</i> equal to <i>ttylabel</i>. <code>-l[ttylabel]</code> If a <i>ttylabel</i> is specified, sttydefs displays the record from <code>/etc/ttydefs</code> whose TTY label matches the specified <i>ttylabel</i>. If no <i>ttylabel</i> is specified, sttydefs displays the entire contents of <code>/etc/ttydefs</code>. sttydefs verifies that each entry it displays is correct and that the entry's <i>nextlabel</i> field references an existing <code>-r ttylabel</code> Remove any record in the <code>tttydefs</code> file that has <i>ttylabel</i> as its label.

Output If successful, `sttydefs` will exit with a status of 0. `sttydefs -l` will generate the requested information and send it to standard output.

示例 示例1 A sample of `sttydefs` command.
The following command lists all the entries in the `ttydefs` file and prints an error message for each invalid entry that is detected.

```
example# sttydefs -l
```

The following shows a command that requests information for a single label and its output:

```
example# sttydefs -l 9600
-----
9600:9600 hupcl erase ^h:9600 sane ixany tab3 hupcl erase ^h::4800
-----
ttylabel:    9600
initial flags:  9600 hupcl erase ^h
final flags:   9600 sane ixany tab3 hupcl erase ^h
autobaud:     no
nextlabel:    4800
```

The following sequence of commands will add the labels 1200, 2400, 4800, and 9600 and put them in a circular list:

```
sttydefs -a 1200 -n 2400 -i 1200 -f "1200 sane"
sttydefs -a 2400 -n 4800 -i 2400 -f "2400 sane"
sttydefs -a 4800 -n 9600 -i 4800 -f "4800 sane"
sttydefs -a 9600 -n 1200 -i 9600 -f "9600 sane"
```

文件 /etc/ttydefs

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [attributes\(5\)](#), [termio\(7I\)](#)

引用名 su – become superuser or another user

用法概要 su [-] [*username* [*arg...*]]

描述 The su command allows one to become another user without logging off or to assume a role. The default user *name* is root (superuser).

To use su, the appropriate password must be supplied (unless the invoker is already root). If the password is correct, su creates a new shell process that has the real and effective user ID, group IDs, and supplementary group list set to those of the specified *username*. Additionally, the new shell's project ID is set to the default project ID of the specified user. See [getproject\(3PROJECT\)](#), [setproject\(3PROJECT\)](#). The new shell will be the shell specified in the shell field of *username*'s password file entry (see [passwd\(4\)](#)). If no shell is specified, /usr/bin/sh is used (see [sh\(1\)](#)). If superuser privilege is requested and the shell for the superuser cannot be invoked using [exec\(2\)](#), /sbin/sh is used as a fallback. To return to normal user ID privileges, type an EOF character (CTRL-D) to exit the new shell.

Any additional arguments given on the command line are passed to the new shell. When using programs such as sh, an *arg* of the form -c *string* executes *string* using the shell and an *arg* of -r gives the user a restricted shell.

To create a login environment, the command “su -” does the following:

- In addition to what is already propagated, the LC* and LANG environment variables from the specified user's environment are also propagated.
- Propagate TZ from the user's environment. If TZ is not found in the user's environment, su uses the TZ value from the TIMEZONE parameter found in /etc/default/login.
- Set MAIL to /var/mail/*new_user*.

If the first argument to su is a dash (-), the environment will be changed to what would be expected if the user actually logged in as the specified user. Otherwise, the environment is passed along, with the exception of \$PATH, which is controlled by PATH and SUPATH in /etc/default/su.

All attempts to become another user using su are logged in the log file /var/adm/sulog (see [sulog\(4\)](#)).

安全 su uses [pam\(3PAM\)](#) with the service name su for authentication, account management, and credential establishment.

示例 示例 1 Becoming User bin While Retaining Your Previously Exported Environment

To become user bin while retaining your previously exported environment, execute:

```
example% su bin
```


示例 2 Becoming User bin and Changing to bin's Login Environment

To become user bin but change the environment to what would be expected if bin had originally logged in, execute:

example% **su - bin**

示例 3 Executing command with user bin's Environment and Permissions

To execute command with the temporary environment and permissions of user bin, type:

example% **su - bin -c "command args"**

环境变量

Variables with LD_ prefix are removed for security reasons. Thus, su bin will not retain previously exported variables with LD_ prefix while becoming user bin.

If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see [environ\(5\)](#)) are not set in the environment, the operational behavior of su for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables are set in the environment, the “C” (U.S. style) locale determines how su behaves.

- LC_CTYPE Determines how su handles characters. When LC_CTYPE is set to a valid value, su can display and handle text and filenames containing valid characters for that locale. su can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. su can also handle EUC characters of 1, 2, or more column widths. In the “C” locale, only characters from ISO 8859-1 are valid.
- LC_MESSAGES Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the “C” locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

文件

- \$HOME/.profile
user's login commands for sh and ksh
- /etc/passwd
system's password file
- /etc/profile
system-wide sh and ksh login commands
- /var/adm/sulog
log file
- /etc/default/su
the default parameters in this file are:

- SULOG If defined, all attempts to su to another user are logged in the indicated file.
- CONSOLE If defined, all attempts to su to root are logged on the console.
- PATH Default path. (/usr/bin:)
- SUPATH Default path for a user invoking su to root. (/usr/sbin:/usr/bin)
- SYSLOG Determines whether the [syslog\(3C\)](#) LOG_AUTH facility should be used to log all su attempts. LOG_NOTICE messages are generated for su's to root, LOG_INFO messages are generated for su's to other users, and LOG_CRIT messages are generated for failed su attempts.

/etc/default/login
the default parameters in this file are:

- SLEEPTIME If present, sets the number of seconds to wait before login failure is printed to the screen and another login attempt is allowed. Default is 4 seconds. Minimum is 0 seconds. Maximum is 5 seconds.

Both su and [login\(1\)](#) are affected by the value of SLEEPTIME.
- TIMEZONE Sets the TZ environment variable of the shell.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[csh\(1\)](#), [env\(1\)](#), [ksh\(1\)](#), [login\(1\)](#), [roles\(1\)](#), [sh\(1\)](#), [syslogd\(1M\)](#), [exec\(2\)](#), [getprojent\(3PROJECT\)](#), [setproject\(3PROJECT\)](#), [pam\(3PAM\)](#), [pam_authenticate\(3PAM\)](#), [pam_acct_mgmt\(3PAM\)](#), [pam_setcred\(3PAM\)](#), [pam.conf\(4\)](#), [passwd\(4\)](#), [profile\(4\)](#), [slog\(4\)](#), [syslog\(3C\)](#), [attributes\(5\)](#), [environ\(5\)](#)

引用名	sulogin – access single-user mode				
用法概要	sulogin				
描述	The sulogin utility is automatically invoked by init when the system is first started. It prompts the user to type a user name and password to enter system maintenance mode (single-user mode) or to type EOF (typically CTRL-D) for normal startup (multi-user mode). The user should never directly invoke sulogin. The user must have the solaris.system.maintenance authorization. The sulogin utility can prompt the user to enter the root password on a variable number of serial console devices, in addition to the traditional console device. See consadm(1m) and msglog(7D) for a description of how to configure a serial device to display the single-user login prompt.				
文件	/etc/default/sulogin Default value can be set for the following flag: PASSREQ Determines if login requires a password. Default is PASSREQ=YES. /etc/default/login Default value can be set for the following flag: SLEEPTIME If present, sets the number of seconds to wait before login failure is printed to the screen and another login attempt is allowed. Default is 4 seconds. Minimum is 0 seconds. Maximum is 5 seconds. Both su(1M) and login(1) are affected by the value of SLEEPTIME.				
属性	See attributes(5) for descriptions of the following attributes:				
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>system/core-os</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	system/core-os
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	system/core-os				
另请参见	auths(1) , login(1) , consadm(1m) , init(1M) , su(1M) , attributes(5) , msglog(7D)				
附注	By default, the root user has all authorizations. Granting the solaris.system.maintenance authorization to the Console User Rights Profile may have an undesirable side effect of granting the currently logged in user maintenance mode access. The solaris.system.maintenance authorization should be directly granted to appropriate users rather than through the Console User Rights Profile.				

引用名	svadm – 是用来控制 Availability Suite 存储卷操作的命令行接口
用法概要	<pre>svadm -h svadm -v svadm [-C tag] svadm [-C tag] -i svadm [-C tag] -e {-f config_file volume} svadm [-C tag] -d {-f config_file volume} svadm [-C tag] -r {-f config_file volume}</pre>
描述	svadm 命令提供相应的工具来为指定的卷启用和禁用存储卷 (Storage Volume, SV) 驱动程序以及动态地重新配置系统，从而对 SV 驱动程序进行控制。
选项	如果没有为 svadm 命令指定参数，则该实用程序将显示当前处于 SV 控制下的卷的列表。svadm 支持以下选项： -C tag 在群集节点上，用于将操作限制到属于标记指定的群集资源组或磁盘组名称的那些卷。在非群集系统上，此选项是非法的。可以使用特殊标记 local 将操作限制到无法转移到群集中其他节点的那些卷。 -d 禁用在命令行上或配置文件中指定的 SV 设备。如果同时指定了此选项与 -c 标记，则卷应该在此群集磁盘组中。 -e 启用在命令行上或配置文件中指定的 SV 设备。卷的详细信息将保存在当前配置中。请参见 dscfg(1M)。如果同时指定了此选项与 -c 标记，则卷应该在此群集磁盘组中。 -f config_file 指定包含卷列表的配置文件。命令将读取该卷列表，然后执行操作。config_file 的格式是卷路径名的简单列表，一行一个路径名。将忽略空白行以及以注释字符 (#) 开头的行。 -h 显示 svadm 用法摘要。 -i 显示当前处于 SV 控制下的卷的扩展状态。 -r 当指定了 config_file 时，将重新配置正在运行的系统以使其与 config_file 中指定的配置相匹配。当指定了 -c 选项时，将比较每个卷的群集标记并将其更改为 cluster_tag。如果随此选项指定了一个卷，则只有重新配置与此卷关联的群集标记有效。要启用或禁用单个卷，应该使用 -e 或 -d 选项。

-v
显示 SV 版本号。

用法 启用 SV 卷后，对设备的常规系统调用访问（请参见 [Intro\(2\)](#)）将被重定向到 StoreEdge 体系结构软件。这样，标准应用程序可以使用 StorageTek 功能，如 Oracle StorageTek Point-in-Time Copy 和 Remote Mirror Software。

svadm 命令在执行启用 (-e) 和禁用 (-d) 操作时，会在 Availability Suite 日志文件 /var/adm/ds.log 中生成一个条目（请参见 [ds.log\(4\)](#)）。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	driver/storage/sv、driver/storage/sv
接口稳定性	Committed（已确定）

另请参见 [dscfg\(1M\)](#)、[ds.log\(4\)](#)、[attributes\(5\)](#)、[sv\(7D\)](#)

引用名	svcadm – 处理服务实例
用法概要	<pre>/usr/sbin/svcadm [-v] enable [-rst] {FMRI pattern}... /usr/sbin/svcadm [-v] disable [-st] {FMRI pattern}... /usr/sbin/svcadm [-v] restart {FMRI pattern}... /usr/sbin/svcadm [-v] refresh {FMRI pattern}... /usr/sbin/svcadm [-v] clear {FMRI pattern}... /usr/sbin/svcadm [-v] mark [-It] instance_state {FMRI pattern}... /usr/sbin/svcadm [-v] milestone [-d] milestone_FMRI</pre>
描述	svcadm 发出针对在服务管理工具中执行的服务的操作请求（请参见 smf(5) ）。针对某个服务的操作是由分配给它的服务重启程序代理执行的。缺省的服务重启程序为 <code>svc.startd</code> （请参见 svc.startd(1M) ）。
选项	所支持的选项如下： <code>-v</code> 将操作详细地输出到标准输出。
子命令	
常用操作	下面列出了在对服务实例的典型管理中使用的子命令。 对于接受一个或多个操作数的子命令，如果操作数指定了一个服务（而不是一个服务实例）且该服务仅有一个实例，则 <code>svcadm</code> 对该实例执行操作。如果某个缩写的 <i>FMRI</i>（故障管理资源标识符）或模式与多个服务相匹配，则显示一条警告消息并忽略该操作数。请参见 smf(5)。 如果服务具有多个实例，<code>svcadm</code> 将返回非零退出状态。 <code>enable [-rst] {FMRI pattern}...</code> 启用由操作数指定的服务实例。对于每个服务实例，所分配的重启程序会尝试使其进入 online（联机）状态。此操作需要相应的权限来修改服务实例的 "general"（通用）属性组（请参见 smf_security(5)）。 如果指定了 <code>-r</code> 选项，<code>svcadm</code> 将启用每个服务实例并递归启用其依赖项。 如果指定了 <code>-s</code> 选项，<code>svcadm</code> 将启用每个服务实例，然后等待每个服务实例进入 online（联机）或 degraded（降级）状态。如果 <code>svcadm</code> 确定服务不能在不经管理员干预的情况下达到这些状态，它会提早返回。

	如果指定了 -t 选项，svcadm 会临时启用每个服务实例。临时启用仅会持续到重新引导之前。此操作需要相应的权限来修改服务实例的 "restarter_actions" 属性组（请参见 smf_security(5)）。缺省情况下，enable 在重新引导后仍然有效。
disable [-st] {FMRI pattern}...	禁任由操作数指定的服务实例。对于每个服务实例，所分配的重启程序会尝试使其进入 disabled（禁用）状态。此操作需要相应的权限来修改服务实例的 "general"（通用）属性组（请参见 smf_security(5)）。 如果指定了 -s 选项，svcadm 将禁用每个服务实例，然后等待每个服务实例进入 disabled（禁用）状态。如果 svcadm 确定服务不能在不经管理员干预的情况下达到此状态，它会提早返回。 如果指定了 -t 选项，svcadm 会临时禁用每个服务实例。临时禁用仅会持续到重新引导之前。此操作需要相应的权限来修改服务实例的 "restarter_actions" 属性组（请参见 smf_security(5)）。缺省情况下，disable 在重新引导后仍然有效。
restart {FMRI pattern}...	请求重新启动由操作数指定的服务实例。此操作需要相应的权限来修改服务实例的 "restarter_actions" 属性组（请参见 smf_security(5)）。大多数重启程序将重新启动服务操作实施为一个完整的服务 "stop"（停止）操作，后跟一个 "start"（启动）操作。 此子命令只能重新启动处于 online（联机）或 degraded（降级）状态的服务，因为这些状态是在 smf(5) 中定义的。
refresh {FMRI pattern}...	针对操作数指定的每个服务实例，请求所分配的重启程序使用来自当前配置的值更新该服务的运行配置快照。其中某些值是立即生效的（例如，依赖项更改）。其他值在服务下次 restart（重新启动）之前不会生效。有关更多信息，请参见重启程序和服务文档。

如果服务是由 `svc.startd(1M)` 管理的，将调用 `refresh` 方法（如果存在）来请求服务重新读取其自己的配置。对于其他重启程序，请参见重启程序文档。

此操作需要相应的权限来修改服务实例的 "restarter_actions" 属性组（请参见 `smf_security(5)`）。

`clear {FMRI|pattern}...`

针对操作数指定的每个服务实例，如果它处于 `maintenance`（维护）状态，则向所分配的重启程序发出服务已修复的信号。如果实例处于 `degraded`（降级）状态，则请求所分配的重启程序将该服务更改为 `online`（联机）状态。此操作需要相应的权限来修改服务实例的 "restarter_actions" 属性组（请参见 `smf_security(5)`）。

例外操作

下列子命令用于服务开发与临时管理处理。

`mark [-It] instance_state {FMRI|pattern}...`

如果 `instance_state` 为 "maintenance"（维护），则针对操作数指定的每个服务，`svcadm` 请求所分配的重启程序将该服务置于 `maintenance`（维护）状态。有关为每个重启程序采取的操作的详细描述，请参见 `svc.startd(1M)` 和 `inetd(1M)`。

如果 `instance_state` 为 "degraded"（降级），则针对操作数指定的处于联机状态的服务，`svcadm` 请求分配给这些服务的重启程序将其置于 `degraded`（降级）状态。

如果指定了 `-I` 选项，该请求被标记为即时的。

`-t` 选项仅对维护请求有效。如果指定了此选项，该请求将被标记为临时的，它的效果仅会持续到下次重新引导之前。

`milestone [-d] milestone_FMRI`

如果 `milestone_FMRI` 为关键字 "none"，则将临时禁用除主重启程序 `svc:/system/svc/restarter:default` 以外的所有服务。

如果 *milestone_FMRI* 为关键字 "all", 则将废除对所有服务的临时启用和禁用请求。

如果 *milestone_FMRI* 为下列值之一：

```
svc:/milestone/single-user:default
svc:/milestone/multi-user:default
svc:/milestone/multi-user-server:default
```

则将废除对指定的服务及其依赖的（直接或间接）的所有服务的临时启用和禁用请求。将临时禁用所有其他服务。

使用 "milestone" 子命令更改系统的当前里程碑将不会更改系统的当前运行级。要更改系统的运行级，请直接调用 `/usr/sbin/init`。

此操作需要相应的权限来修改

`svc:/system/svc/restarter:default` 服务实例的 "options_ovr" 属性组（请参见 [smf_security\(5\)](#)）。

如上所述，`-d` 选项会即时将里程碑更改为所需的里程碑。此外，它将使指定的里程碑成为缺省的引导里程碑，缺省的引导里程碑在重新引导后仍然有效。缺省的里程碑是由主重启程序

`svc:/system/svc/restarter:default` 中的 `options/milestone` 属性定义的。如果此属性不存在，缺省值为 "all"。此操作需要相应的权限来修改

`svc:/system/svc/restarter:default` 服务实例的 "options"（选项）属性组（请参见 [smf_security\(5\)](#)）。

操作数

支持下列操作数：

FMRI 一个 *FMRI*，用于指定一个或多个实例。可以通过指定实例名称或指定服务名的后缀部分以缩写形式指定 *FMRI*。例如，对于以下 *FMRI*：

```
svc:/network/smtp:sendmail
```

下列各项都是有效的缩写：

```
sendmail
:sendmail
```

```
smtp
smtp:sendmail
network/smtp
```

而下列项则无效：

```
mail
network
network/smt
```

如果 *FMRI* 指定了一个服务，则命令将应用于该服务的所有实例。*FMRI* 的缩写形式不稳定，不应在脚本或其他永久性工具中使用。

pattern 服务实例的 *FMRI* 根据 [fnmatch\(5\)](#) 描述的 "globbing"（通配）规则与之匹配的模式。如果该模式不以 "svc:" 开头，则会在其前面加上 "svc:/"。

如果某个缩写的 *FMRI* 或模式与多个服务相匹配，则显示一条警告消息并忽略该操作数。

示例

示例1 重新启动服务实例

以下命令重新启动 NFS 服务器。缺省服务实例的完整 *FMRI* 为：`svc:/network/nfs/server:default`

不过，您可以将该完整的 *FMRI* 缩写为以下形式：

```
# svcadm restart nfs/server
```

示例2 禁用标准的 HTTP 服务器

以下命令禁用标准的 HTTP 服务器，其中使用了缩写的 *FMRI*：

```
$ svcadm disable http
```

示例3 启用实例及其依赖的实例

以下命令启用 `foo:bar` 实例及其依赖的所有实例：

```
$ svcadm enable -r foo:bar
```

示例4 以同步方式启用实例

以下命令启用 `foo:bar` 实例。在实例联机或 `svcadm` 断定该服务无法联机之前，此命令不会返回。

```
$ svcadm enable -s foo:bar
```

示例5 限制和恢复正在运行的服务

以下命令将正在运行的服务限制到单用户模式：

```
# svcadm milestone milestone/single-user
```

示例5 限制和恢复正在运行的服务 (续)

以下命令恢复正在运行的服务：

```
# svcadm milestone all
```

退出状态

将返回下列退出值：

- 0 成功完成。
- 1 发生了致命错误。会在标准错误输出中显示一条或多条错误消息。
- 2 指定的命令行选项无效。
- 3 **svcadm** 断定由于它等待的服务实例本身存在某个问题，不经管理员干预该服务实例无法达到所需的状态。
- 4 **svcadm** 断定由于它等待的服务实例的依赖项存在某个问题，不经管理员干预该服务实例无法达到所需的状态。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	请参见下文。

交互式输出为 **Uncommitted**（未确定）。调用和非交互式输出为 **Committed**（已确定）。

另请参见

[svccprop\(1\)](#)、[svcs\(1\)](#)、[inetd\(1M\)](#)、[init\(1M\)](#)、[svccfg\(1M\)](#)、[svc.startd\(1M\)](#)、[libscf\(3LIB\)](#)、

附注

svcadm 等待服务及其依赖项更改状态时花费的时间量由其方法超时隐式限制。例如，如果某个服务使用的缺省重启程序的 **start** 方法挂起，则超时时间到期后，该服务将转换为 **maintenance**（维护）状态。**svcadm** 将认为该服务在不经管理员干预的情况下无法联机。

尝试以同步方式启用依赖于（直接或间接）某个文件的服务时，如果调用者不具备搜索包含某个文件的目录所需的特权，该尝试可能会失败且退出状态指出不满足依赖项。将来的 Solaris 发行版中可能会取消这一限制。

引用名	svccfg – 导入、导出和修改服务配置
用法概要	<pre> /usr/sbin/svccfg [-v] [-s FMRI] /usr/sbin/svccfg [-v] [-s FMRI] subcommand [args]... /usr/sbin/svccfg [-v] [-s FMRI] -f command-file </pre>
描述	svccfg 命令处理服务配置系统信息库中的数据。svccfg 可以通过单个子命令或指定一个包含一系列子命令的命令文件以交互方式调用。 对系统信息库中现有服务所做的更改通常在下次刷新服务实例之后才会对该服务生效。有关更多详细信息，请参见 svcadm(1M) 手册页上的 <code>refresh</code> 子命令。
选项	所支持的选项如下： -f <i>command-file</i> 读取并执行 <i>command-file</i> 中的 svccfg 子命令。 -s <i>FMRI</i> 在执行任何子命令之前选择 <i>FMRI</i>（fault management resource identifier，故障管理资源标识符）所表示的实体。请参见 smf(5)。 -v 详细模式。
子命令	子命令分为几类，如以下各小节所述。 所有接受 <i>FMRI</i> 的子命令还接受缩写或 <code>glob</code> 模式。可以通过指定实例名称或服务名称的结尾部分使用实例或服务的缩写形式。例如，对于以下 <i>FMRI</i>： <pre>svc:/network/smtp:sendmail</pre> 下列各项都是有效的缩写： <pre> sendmail :sendmail smtp smtp:sendmail network/smtp </pre> 而下列项则无效： <pre> mail network network/smt </pre> <i>FMRI</i> 的缩写形式不稳定，不应在脚本或其他永久性工具中使用。如果模式与多个实例或服务匹配，则会输出一条错误消息并且不采取任何操作。

常见子命令

```
end
exit
quit
```

立即退出。

```
repository [-p prefix] repfile
```

将 *repfile* 用作系统信息库。缺省情况下，svccfg(1M) 使用系统信息库。

仅将系统信息库用于您当前正在运行的相同版本的 Solaris 中的文件（包括修补程序）。不要将该子命令用于系统信息库 `/etc/svc/repository.db`。

如果在部署时间之前使用 `svccfg repository` 预填充 SMF 系统信息库，则使用 `-p` 选项为使用 `import` 导入的清单的系统标准位置指定根前缀。系统信息库位于活动系统上后，该前缀将被 `/lib/svc/manifest` 和 `/var/svc/manifest` 替换。如果运行时 `-p` 目录的清单未显示在系统标准位置中，将删除与它们关联的服务。

```
set [-v|-V]
```

设置可选行为。如果未指定选项，设置会显示当前有效的选项。

```
-v
```

启用详细模式。

```
-V
```

禁用详细模式。

服务清单和配置文件子命令

```
apply [-n] [-v] file | directory
```

如果参数是服务配置文件或清单，则将配置应用到 SMF 系统信息库的管理层。将根据需要创建服务、实例、属性组和属性。

如果未指定属性或属性组的类型特性，将尝试从现有的类型设置或服务模板确定类型。如果无法确定类型，将显示一条警告并跳过服务，这样就不会在服务 and 实例中引入不一致的数据。忽略不存在的服务和实例。

要使用配置文件中的宽松元素定义，需要将以下定义添加到 `DOCTYPE` 条目中：

```
<!ENTITY % profile "INCLUDE">
<!ENTITY % manifest "IGNORE">
```

将刷新配置文件修改的服务和实例。如果指定了 `-n`，将处理配置文件并且不对 SMF 系统信息库应用更改。将在标准错误中报告已找到的任何语法错误并返回退出代码 1。有关服务配置文件的说明，请参见 [smf\(5\)](#)。该命令需要具有修改服务和实例中的属性的特权。有关修改属性所需的特权，请参见 [smf_security\(5\)](#)。

将针对清单和系统信息库中的模板数据验证清单或配置文件中的服务和实例，并发出所有违反模板的警告。有关模板说明，请参见 [smf_template\(5\)](#)。如果指定了 `-v` 选项，将无法导入违反已定义模板的清单。在交互式调用 `svccfg` 时，`-v` 是缺省行为。

如果要**应用**的参数是一个目录，则会按照以上所述应用在该目录树下找到的所有配置文件。如果指定文件或在指定目录下找到的任何文件不是服务配置文件，该子命令将失败。

extract [-a] [-l *layer*] [> *file*]

在系统上输出表示本地定制的服务配置文件。

如果提供 -l，则可以选择从中导出值的层列表。-l 选项需要层名称，并使用以下参数：`manifest`、`system-profile`、`site-profile`、`admin`、`current`、`all`。`current` 和 `all` 是同义词，表示最高层值。提供这两者以实现与 `listprop` 子命令的一致性。多个层可以以逗号分隔，或使用多个 -l 选项指定。

如果未提供 -l，则缺省值为 `-l admin,site-profile`。

如果在多个选定层中定义属性，则在配置文件中仅导出最高层。

如果没有 -a 选项，将提取包含受保护信息（由存在的 `read_authorization` 属性标识，请参见 [smf_security\(5\)](#)）的属性组，但不提取它们的属性值。指定 -a 选项时，将提取所有的值。如果没有读取这些值的足够特权，则会生成错误。

export [-a] *service_FMRI* [> *file*]

运行 `svccfg export` 等效于：

```
svccfg extract -l current [-a] service_FMRI [> file]
```

import [-V] [*file* | *directory*]

系统管理的文件系统位置（`/lib/svc/manifest` 和 `/var/svc/manifest` 的子目录）中的文件上的 `svccfg import` 调用：`svcadm restart manifest-import`。

推荐的做法是将清单放置在系统管理的位置，并调用 `svcadm restart manifest-import` 来导入它们。

位于其他位置的文件上的 `svccfg import` 将其属性作为管理定制导入到管理层。这等效于：

```
svccfg apply [file | directory]
```

inventory *file*

如果 *file* 确定为服务清单或配置文件，则会输出 *file* 说明的服务和实例的 FMRI。对于每个服务，在服务的 FMRI 之前会显示其实例的 FMRI。

validate [*file* | *fmri*]

`validate` 子命令可以针对清单文件、实例 FMRI、当前实例或快照实体选择运行。指定参数后，`svccfg` 将检查是否存在指定的文件。如果该文件存在，将对其进行验证。如果指定名称的文件不存在，则将该参数视为 FMRI 模式。如果文件名与 FMRI 之间存在冲突，请使用 `svc:` 和 `file:` 前缀指示 `svccfg` 如何解释参数。

指定文件时，将按照与 `import -V` 相似的方式处理该文件，但不对系统信息库执行任何更改。如果检测到任何错误，`svccfg` 会显示错误并以非零退出状态退出。

对于实例 *fmri*、实例实体选择或快照实体选择，将针对系统信息库中的模板数据验证复合格式的指定实例（请参见 [smf\(5\)](#) 中的“属性和属性组”）。实例 FMRI 和实例实体选择使用“正在运行的”快照进行验证。将发出所有违反模板的警告。有关模板说明，请参见 [smf_template\(5\)](#)。

如果服务或实例由清单或配置文件支持，则它被屏蔽。请参见 [smf\(5\)](#)。

实体选择、修改和 导航子命令

“实体”表示作用域、服务或服务实例。

`add name`

将创建具有指定名称的新实体作为当前选择的子实体。有关创建实体所需的特权，请参见 [smf_security\(5\)](#)。

`delete [-f] {name | fmri}`

删除当前选择的命名子实体或 *fmri* 指定的实体。除非指定 `-f` 标志，否则尝试删除“联机”或“降级”状态的服务实例会失败。如果服务或服务实例具有类型为“framework”的“dependents”属性组，则对于其中类型为“astring”或“fmri”的每个属性，当属性包含用于命名服务或服务实例的单个值时，将删除指定服务或服务实例中与该属性同名的相关项属性组。有关删除服务配置所需的特权，请参见 [smf_security\(5\)](#)。

`list [pattern]`

显示其名称与 `glob` 模式 *pattern* 匹配的当前选择的子实体（请参见 [fnmatch\(5\)](#)）。还为具有属性的实体（即服务和服务实例）列出 `'properties'`。

`select {name | fmri}`

如果参数指定当前所选内容的子级，它将成为当前所选内容。否则，参数将被解释为 FMRI，参数指定的实体成为当前所选内容。

`unselect`

当前所选内容的父级成为当前所选内容。

属性检查和修改子 命令

`addpg name type [flags]`

将具有指定 *name* 和类型的属性组添加到当前所选内容。*flags* 是一个字符串，用于指定创建属性组所使用的标志。`'P'` 表示 `SCF_PG_FLAG_NONPERSISTENT`（请参见 [scf_service_add_pg\(3SCF\)](#)）。有关创建属性组所需的特权，请参见 [smf_security\(5\)](#)。

`addpropvalue pg/name [type:] value`

将指定值添加到属性。如果 *type* 已指定并且属性存在，则当 *type* 与属性的 *type* 不一致时，子命令将失败。这些值可以用双引号引住。包含双引号或反斜杠的字符串值必须用双引号引住，包含的双引号和反斜杠必须用反斜杠引住。创建不存在的属性，在这种情况下必须提供 *type* 说明符。有关可用属性类型的列表，请参见 [scf_value_create\(3SCF\)](#)。有关修改属性所需的特权，请参见 [smf_security\(5\)](#)。新值将附加到与属性相关的属性值列表末尾处。

`delcust [-M] [pattern]`

删除当前选择的任何管理定制。如果提供参数，将作为 `glob` 模式获取并仅删除名称与参数匹配的属性组和属性。

如果提供 **-M**，则仅删除屏蔽的实体。

要查看 **delcust** 将删除的定制，请使用 **listcust** 和相同的选项。因为 **delcust** 可能删除所有管理定制，所以始终首先运行 **listcust** 来确定您要删除所需的内容。

delpg *name*

删除当前选择的属性组 *name*。有关删除属性组所需的特权，请参见 [smf_security\(5\)](#)。

如果属性组由清单或配置文件支持，则它被屏蔽。请参见 [smf\(5\)](#)。

delprop *pg[/name]*

删除当前选择的命名属性组或属性。有关删除属性所需的特权，请参见 [smf_security\(5\)](#)。

delpropvalue *pg/name globpattern*

删除与命名属性中的指定 *glob* 模式匹配的所有值。即使没有值与之匹配也会成功。有关修改属性所需的特权，请参见 [smf_security\(5\)](#)。

describe [-v] [-t] [*propertygroup/property*]

说明当前或可能的设置。

不使用参数调用时，**describe** 会提供当前所选实体及其所有当前设置的属性组和属性的基本说明（如果有）。通过指定属性组名称或者属性组名称和属性名称（用斜杠 (/) 隔开），可以将属性组或具体属性作为参数查询。

-v 选项提供所有的可用信息，包括当前设置、约束及其他可能的设置选择的说明。

-t 选项仅显示选择的模板数据（请参见 [smf_template\(5\)](#)），不显示属性组和属性的当前设置。

editprop

对于重现当前选择的属性组和属性的命令的注释，将放置在临时文件中，调用 **EDITOR** 环境变量指定的程序对其进行编辑。完成后执行临时文件中的命令。缺省编辑器是 [vi\(1\)](#)。有关创建、修改或删除属性所需的特权，请参见 [smf_security\(5\)](#)。

listpg [*pattern*]

显示当前选择的属性组的名称、类型和标志。如果指定了某个参数，将作为 *glob* 模式获取并仅列出名称与该参数匹配的属性组。

在交互模式下，还提供属性组的基本说明。

listprop [-l *layer...*] [-f | -o *format*] [*pattern*]

列出当前选择的属性组和属性。对于属性组，列出其名称、类型和标志。对于属性，列出其名称（前面会加上属性组名称和斜杠 (/)）、类型和值。有关可用属性类型的列表，请参见 [scf_value_create\(3SCF\)](#)。如果提供参数，将作为 *glob* 模式获取并仅列出名称与参数匹配的属性组和属性。

使用 `-l` 选项，输出值来自的层。`-l` 选项需要层，并使用以下参数：`manifest`、`system-profile`、`site-profile`、`admin`、`current`、`all`。`current` 输出与 `listprop` 相同的属性值（不使用 `-l`），以及在其中定义值的层。

`-f` 和 `-o` 选项是互斥的。`-f` 输出属性来自的文件（如果有）。`-o` 允许字段选择。可选择的字段包括：

<code>propname</code>	属性名称
<code>pgname</code>	属性组名称
<code>instname</code>	实例名称
<code>servicename</code>	服务名称
<code>layer</code>	层
<code>proptype</code>	属性类型
<code>value</code>	属性值
<code>file</code>	源文件
<code>masked</code>	属性组或属性当前是否被屏蔽
<code>time</code>	上次更改该属性的时间

`listcust [-L] [-M] [pattern]`

为当前选择输出任何站点配置文件和管理层定制以及屏蔽实体。如果提供参数，将作为 `glob` 模式获取并仅列出名称与参数匹配的属性组和属性。如果没有当前选择，则列出所有服务的所有定制。

如果提供 `-M`，则仅输出屏蔽的实体。

如果提供 `-L`，则显示所有本地定制，包括同时位于管理层和站点配置文件层内的定制。

`setenv [-i | -s] [-m method_name] envvar value`

如果 *method_name* 属性组具有类型 `"method"`，则通过更改该属性组中的 `"environment"` 属性设置服务或实例的方法环境变量。如果未指定 *method_name* 并且使用 `-i` 选项，则当前选择实例时将使用 `"method_context"` 属性组。如果使用 `-s` 选项并且当前选择了服务，则会使用其 `"method_context"` 属性组。如果使用 `-s` 选项并且当前选择了实例，则会使用其父级的 `"method_context"` 属性组。如果既不使用 `-i` 选项也不使用 `-s` 选项，将在当前选择的实体中搜索 `"start"` 属性组，如果当前选择了实例，则还会搜索其父级。如果未找到 `"inetd_start"` 属性组，将以类似方式进行搜索。

找到属性后，将删除所有以 *envvar* 开始并且后跟 `"=` 的值，同时添加值 `"envvar=value"`。有关修改属性所需的特权，请参见 [smf_security\(5\)](#)。

`setprop pg/name = [type:] value`

`setprop pg/name = [type:] ([values ...])`

将当前选择的 *pg* 属性组的 *name* 属性设置为类型 *type* 的指定值。有关可用属性类型的列表，请参见 [scf_value_create\(3SCF\)](#)。如果属性已经存在并且 *type* 与属性中的现有 *type* 不一致，子命令将失败。值可能用双引号引住。包含双引号或反斜杠的字符串值必须用双引号引住，包含的双引号和反斜杠必须用反斜杠引住。如果命名属性不存在，只要指定类型，即创建属性。有关创建或修改属性所需的特权，请参见 [smf_security\(5\)](#)。多个值将按照它们的指定顺序存储。

unsetenv [-i | -s] [-m *method_name*] *envvar* *value*

如果 *method_name* 属性组具有类型 "method"，则通过更改该属性组中的 "environment" 属性删除服务或实例的方法环境变量。如果未指定 *method_name* 并且使用 -i 选项，则当前选择实例时将使用 "method_context" 属性组。如果使用 -s 选项并且当前选择了服务，则会使用其 "method_context" 属性组。如果使用 -s 选项并且当前选择了实例，则会使用其父级的 "method_context" 属性组。如果既不使用 -i 选项也不使用 -s 选项，将在当前选择的实体中搜索 "start" 属性组，如果当前选择了实例，则还会搜索其父级。如果未找到 "inetd_start" 属性组，将以类似方式进行搜索。

找到属性后，将删除所有以 *envvar* 开始并且后跟 "=" 的值。有关修改属性所需的特权，请参见 [smf_security\(5\)](#)。

通知参数子命令

setnotify [{-g} *tset* | *class*] *notification_parameters*

为 SMF 系统信息库中的软件事件和故障管理问题生命周期事件设置通知参数。

-g

用于设置 SMF 状态转换的系统范围的通知参数。请参见 [smf\(5\)](#)。这些通知参数是在 `svc:/system/svc/global:default` 中设置的，与当前选择的 `svccfg` 无关。该子命令会刷新其所修改的全部实例。

class

FMA 事件类或别名的逗号分隔列表。请参见 [smf\(5\)](#) 通知参数。

tset

SMF 状态转换的逗号分隔列表。请参见 [smf\(5\)](#) 通知参数。

notification_parameters

实施的每种通知机制的 URI 格式：对于 SMTP，使用：

```
mailto:addr[?header1=value1[&header2=value2]]
```

...或者：

```
mailto: {[active]|inactive}
```

...对于 SNMP 陷阱，使用：

```
snmp: {[active]|inactive}
```

[sasinfo\(1M\)](#) 中定义的参数 `msg_template` 可以设置为 `mailto` URI 中的头值。例如：

```
mailto:root@localhost?msg_template=<path to template file>
```

SNMP 陷阱指向 `/etc/net-snmp/snmp/snmpd.conf` 中的 `trapsink` 指令定义的或 SNMP 陷阱通知守护进程指定的主机。请参见 [sasinfo\(1M\)](#)。

通知参数特定于指定的类或 `tset` 并覆盖预先存在的通知参数。活动/非活动表单不会覆盖以前的通知参数。它仅启用或禁用指定类或 `tset` 的通知机制。设置通知参数时可隐式将它们设置为活动状态。

`listnotify [-g] [tset] | class`

显示指定类或 `tset` 的现有通知参数。使用 `-g` 选项，会显示 `svc:/system/svc/global:default` 中的通知参数。如果省略 `tset`，表示使用 `all`。

`delnotify [-g] tset | class`

删除指定类或 `tset` 的现有通知参数。使用 `-g` 选项，会删除 `svc:/system/svc/global:default` 中的通知参数。

快照导航和选择子命令

`listsnap`

显示用于当前所选实例的快照。

`revert [snapshot]`

将当前所选实例及其服务的管理定制恢复到记录在命名快照中的属性。如果未提供参数，请使用当前所选的快照，然后在成功时取消选择该快照。可以通过 [svcadm\(1M\)](#) 的 `refresh` 子命令将已更改的属性值设置为活动状态。有关更改属性所需的特权，请参见 [smf_security\(5\)](#)。

`selectsnap [name]`

将当前快照更改为 `name` 命名的快照。如果未指定 `name`，取消选择当前选择的快照。快照是只读的。

实例子命令

`refresh`

将当前配置的值提交到正在运行的快照，并使这些值可供当前选择的实例使用。如果系统信息库子命令尚未用于选择系统信息库，则指示该实例的重启程序重新读取更新的配置。

示例

示例1 导入服务说明

以下示例导入命令行上指定的 XML 清单中 `seismic` 服务的说明。

```
# svccfg import /var/svc/manifest/site/seismic.xml
```

请注意清单必须遵守 [service_bundle\(4\)](#) 中指定的格式。

示例2 导出服务说明

要在本地系统中导出服务说明，请执行以下命令：

```
# svccfg export dumpadm >/tmp/dump.xml
```

示例3 删除服务实例

要删除服务实例，请执行以下命令：

```
# svccfg delete network/inetd-upgrade:default
```

示例4 检查备用系统信息库中的属性

要在装入备用系统信息库之后检查服务的属性状态，请使用如下所示的命令序列。例如，用户可以使用此类命令确定是否在特定的系统信息库备份中启用了服务。

```
# svccfg
svc:> repository /etc/svc/repository-boot
svc:> select telnet:default
svc:/network/telnet:default> listprop general/enabled
general/enabled boolean false
svc:/network/telnet:default> exit
```

示例5 启用调试

要修改启动方法的 LD_PRELOAD 并在调试功能处于活动状态的情况下启用 [libumem\(3LIB\)](#)，请执行以下命令：

```
$ svccfg -s system/service setenv LD_PRELOAD libumem.so
$ svccfg -s system/service setenv UMEM_DEBUG default
```

示例6 使用 describe 子命令

以下命令说明 describe 子命令的用法。

```
# svccfg -s console-login describe ttymon
ttymon                                application
ttymon/device                        astring /dev/console
    terminal device to be used for the console login prompt
ttymon/label                         astring console
    appropriate entry from /etc/ttydefs
...
```

示例7 配置通知首选项

以下命令为 SMF 服务状态转换事件配置通知首选项。

```
# svccfg setnotify -g from-online,to-maintenance \
mailto:admin@somehost.com
```

示例8 启用 SNMP 通知

以下命令对故障管理事件启用 SNMP 通知。

```
# svccfg setnotify problem-diagnosed,problem-updated \
mailto:admin@somehost.com snmp:
```

示例9 列出通知设置

以下命令列出故障管理事件的通知设置。

```
# svccfg listnotify problem-diagnosed,problem-updated
Event: problem-diagnosed
Notification Type: smtp
```

示例9 列出通知设置 (续)

```
active: true
to: admin@somehost.com
Notification Type: snmp
active: true

Event: problem-updated
Notification Type: smtp
active: true
to: admin@somehost.com
Notification Type: snmp
active: true
```

环境变量

EDITOR
使用 `editprop` 子命令时运行的命令。缺省编辑器是 [vi\(1\)](#)。

退出状态

将返回下列退出值：

- 0
成功执行。
- 1
一个或多个子命令导致失败。错误消息将写入标准错误流中。
- 2
指定的命令行选项无效。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	请参见下文。

交互式输出为 `Uncommitted`（未确定）。调用和非交互式输出为 `Committed`（已确定）。

另请参见

[svccprop\(1\)](#)、[svcs\(1\)](#)、[sasinfo\(1M\)](#)、[svcadm\(1M\)](#)、[svc.configd\(1M\)](#)、[libscf\(3LIB\)](#)、[libumem](#)

引用名 `svc.configd` – 服务管理工具系统信息库守护进程

用法概要 `/lib/svc/bin/svc.configd`

描述 `svc.configd` 是服务管理工具的系统信息库守护进程。`svc.configd` 是在系统启动期间自动调用的，并在发生故障后重新启动。决不可直接调用 `svc.configd`。

与 `svc.configd` 的交互是通过 [libscf\(3LIB\)](#) 以及以下命令行工具进行的：[svcs\(1\)](#)、[svcprop\(1\)](#)、[svcadm\(1M\)](#) 和 [svccfg\(1M\)](#)。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见 [svcs\(1\)](#)、[svcprop\(1\)](#)、[svcadm\(1M\)](#)、[svccfg\(1M\)](#)、[libscf\(3LIB\)](#)、[attributes\(5\)](#)

引用名	svc.ipfd – IP 过滤器防火墙监视守护进程
用法概要	<pre>/lib/svc/bin/svc.ipfd svc:/network/ipfilter:default</pre>
描述	svc.ipfd 守护进程监视使用防火墙配置的服务操作，并启动更新服务的 IP 过滤器配置。此守护进程使得系统能够在每个服务级别以递增方式对系统的防火墙配置的变化作出反应。 服务被启用时，将激活服务的防火墙策略；服务被禁用时，将停用服务的防火墙策略；服务的配置属性组被修改时，将更新该服务的防火墙策略。svc.ipfd 监视服务管理工具 (Service Management Facility, SMF) 系统信息库以注意这些操作，并调用 IP 过滤器规则生成进程来实施服务的防火墙策略。 此守护进程由 network/ipfilter 服务通过 start 或 refresh 方法启动。因此，此守护进程继承相应方法的环境变量和凭证，并以具有所有区域特权的超级用户身份运行。
防火墙静态配置	静态定义描述了服务用于生成特定于服务的 IPF 规则的网络资源配置。每个服务的 firewall_context 属性组包含服务的静态定义，类似于 inetd 管理的服务中的 inetd 属性组。该属性组支持： firewall_context/name 适用于非 inetd 服务。IANA 名称或 RPC 名称，等同于 inetd/name 属性。 firewall_context/isrpc 适用于非 inetd 服务。布尔属性，其中 true 值表示 RPC 服务，等同于 inetd/isrpc 属性。对于 RPC 服务，firewall_context/name 的值不是 IANA 名称，而是 RPC 程序编号或名称。请参见 rpc(4)。 另外，某些服务可能需要一种机制来生成和提供其自己的 IPF 规则。可选属性 ipf_method 提供了一种允许此类自定义规则生成的机制： firewall_context/ipf_method 一个命令。通常是一个用于为服务生成 IPF 规则的脚本。该框架不会为具有此属性定义的服务来生成规则。相反，该框架希望这些服务提供它们自己的规则。 服务的 ipf_method 指定一个命令，该命令接受一个附加参数、其自己的故障管理资源标识符 (fault management resource identifier, FMRI)，生成服务的防火墙规则并将这些规则输出到 stdout。要为具有 ipf_method 属性的服务生成规则，该框架通过将服务 FMRI 传递为附加参数来执行 ipf_method 中指定的命令，并通过将命令输出、规则重定向到服务的规则文件来存储服务的规则。因为 ipf_method 是从 network/ipfilter start 或 refresh 方法进程上下文中执行的，因此它继承执行上下文并以超级用户身份运行。 服务静态配置是由服务开发者提供的，用户不能修改。这些属性仅在安装更新的服务定义时进行修改。

防火墙策略配置

每个服务的属性组 `firewall_config` 存储服务的防火墙策略配置。因为 `network/ipfilter:default` 负责两个防火墙策略，即适用于整个系统范围的 `Global Default`（全局缺省值）和 `Global Override`（全局覆盖）策略（如 `ipfilter(5)` 中所述），因此它有两个属性组 `firewall_config_default` 和 `firewall_config_override`，分别存储适用于整个系统范围的这两个策略。

下面是属性、属性的可能值，以及相应的语义：

`policy`

`policy` 有以下模式：

`none` 策略模式

无访问限制。对于全局策略，此模式允许所有传入通信。对于服务策略，此模式允许发送到其服务的所有传入通信。

`deny` 策略模式

比 `none` 的限制性要强。此模式允许来自除 `apply_to` 中指定的源以外的所有源的传入通信。

`allow` 策略模式

限制性最强的模式。此模式阻止来自除 `apply_to` 中指定的源以外的所有源的传入通信。

`apply_to`

一个多值属性，它列出了要实施所选的策略模式的网络实体。如果策略为 `deny`，则将拒绝 `apply_to` 属性中列出的实体；如果策略为 `allow`，则将允许列出的实体。可能值的语法如下：

<code>host:</code>	<code>host:IP</code>	<code>"host:192.168.84.14"</code>
<code>subnet:</code>	<code>network:IP/netmask</code>	<code>"network:129.168.1.5/24"</code>
<code>ippool:</code>	<code>pool:pool number</code>	<code>"pool:77"</code>
<code>interface:</code>	<code>if:interface_name</code>	<code>"if:e1000g0"</code>

`exceptions`

一个多值属性，它列出了需要排除在 `apply_to` 列表之外的网络实体。例如，当向某个子网应用 `deny` 策略时，通过在 `exceptions` 属性中指定该子网中的某些主机，可对这些主机进行例外处理。此属性具有与 `apply_to` 属性相同的值语法。

仅适用于各个网络服务：

`firewall_config/policy`

还可以将服务的策略设置为 `use_global`。采用 `use_global` 策略模式的服务继承 `Global Default`（全局缺省值）防火墙策略。

仅适用于 `Global Default`（全局缺省值）：

`firewall_config_default/policy`

`Global Default`（全局缺省值）策略，`svc:/network/ipfilter:default` 中的 `firewall_config` 属性组也可以设置为 `custom`。用户可将 `policy` 设为 `custom` 以使用预填充的 IP 过滤器配置，例如，现有的 IP 过滤器配置或框架无法提供的自定义配

置。此仅适用于 Global Default（全局缺省值）的策略模式允许用户提供一个包含完整的 IPF 规则集的文本文件。当选择了 `custom` 模式时，指定的 IPF 规则集是**完整**的，且框架不会基于所配置的防火墙策略生成 IPF 规则。

firewall_config_default/custom_policy_file

当 Global Default（全局缺省值）策略设为 `custom` 时要使用的文件路径。该文件包含用来提供所需的 IP 过滤器配置的 IPF 规则集。例如，在 `/etc/ipf/ipf.conf` 中具有现有 IPF 规则的用户可执行下列命令来使用现有的规则：

1. 设置自定义策略：

```
# svccfg -s ipfilter:default setprop \
  firewall_config_default/policy = astring: "custom"
```

2. 指定自定义文件：

```
# svccfg -s ipfilter:default setprop \
  firewall_config_default/custom_policy_file = astring: \
  "/etc/ipf/ipf.conf"
```

3. 刷新配置：

```
# svcadm refresh ipfilter:default
```

firewall_config_default/open_ports

要求允许其传入通信的非服务程序可以请求防火墙允许将通信流量发往其通信端口。此多值属性包含协议和端口元组，格式如下：

```
"{tcp | udp}:{PORT | PORT-PORT}"
```

初始状态下，适用于整个系统范围的策略设为 `none`，网络服务的策略设为 `use_global`。由于适用于整个系统范围的策略为 `none` 且所有服务均继承该策略，因此启用 `network/ipfilter` 将使用一组空的 IP 过滤器规则来激活防火墙。要配置更具限制性的策略，请使用 [svccfg\(1M\)](#) 来修改网络服务的策略和适用于整个系统范围的策略。

用户可通过修改服务的 `firewall_config` 属性组来配置防火墙策略。已创建了一个新的授权 `solaris.smf.value.firewall.config`，用以允许将防火墙管理特权委托给用户。具有 Service Operator（服务操作员）特权的用户需要具有该新的授权才能配置防火墙策略。

防火墙可用性

在引导过程中，将在启动已启用的服务之前为其配置防火墙。因此，在引导时服务是受保护的。当系统在运行时，在为服务配置防火墙期间运行服务时，诸如服务重新启动、启用和刷新之类的管理操作可能会导致服务在短暂时间内存在漏洞。

`svc.ipfd` 监视服务的启动和停止事件，并在 SMF 启动或停止服务的同时配置或取消配置服务的防火墙。因为两个操作同步进行，如果在防火墙配置完成前启动服务，可能会暴露窗口（小于一秒）。RPC 服务通常在临时地址上侦听，在服务实际运行前这些地址是未知的。RPC 服务会面临同样的暴露，因为在服务运行前不会配置防火墙。

开发者文档

建议将提供远程功能的服务加入到防火墙框架中以控制对服务的网络访问。虽然该框架集成不是强制性的，但如果配置了适用于整个系统范围的策略，对未集成到框架中的服务的远程访问可能无法正常进行。

将服务集成到框架中非常简便，仅需在服务清单中定义两个附加的属性组及其相应的属性。当用户启用服务时将生成的 IP 过滤器规则。需要使用 shell 脚本生成非琐碎的自定义规则时，可以使用现有的一些脚本作为示例。

其他属性组（`firewall_config` 和 `firewall_context`）分别存储防火墙策略配置和提供静态防火墙定义。下面是新的属性组和属性及其相应的缺省值的摘要。

防火墙策略配置：

`firewall_config`

对系统的访问由一个新的授权定义和一个用户自定义属性类型予以保护。应当以如下方式将新的授权分配给属性组 `value_authorization` 属性：

```
<propval name='value_authorization' type='astring'
value='solaris.smf.value.firewall.config' />
```

第三方应按照服务符号名称空间约定来生成用户自定义类型。Sun 提供的服务可以使用 `com.sun,fw_configuration` 作为属性类型。

有关更多信息，请参见上文中的“防火墙策略配置”。

`firewall_config/policy`

此属性的初始值应为 `use_global`，因为缺省情况下服务会继承 Global Default（全局缺省值）防火墙策略。

`firewall_config/apply_to`

一个空属性，此属性没有初始值。

`firewall_config/exceptions`

一个空属性，此属性没有初始值。

防火墙静态定义：

`firewall_context`

第三方应按照服务符号名称空间约定来生成用户自定义类型，Oracle 提供的服务可以使用 `com.sun,fw_definition` 作为属性类型。

有关更多信息，请参见上文中的“防火墙静态配置”。

`firewall_context/name`

具有大家熟知的由 IANA 定义的端口的服务，可通过 `getservbyname(3SOCKET)` 获取。服务的 IANA 名称存储在此属性中。对于 RPC 服务，RPC 程序编号存储在此属性中。

`firewall_context/isrpc`

对于 RPC 服务，创建此属性时应将其值设为 `true`。

firewall_context/ipf_method
通常，指定的防火墙策略将用于生成服务的通信端口（派生自 firewall_context/name 属性）的 IP 过滤器规则。没有由 IANA 定义的端口且不是 RPC 服务的服务将需要生成其自己的 IP 过滤器规则。生成其自己的规则的服务可以选择不使用 firewall_context/name 和 firewall_context/isrpc 属性。有关指南，请参见以下服务：

svc:/network/ftp:default
svc:/network/nfs/server:default
svc:/network/ntp:default

...以及具有 ipf_method 的其他服务。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os、network/ipfilter
接口稳定性	Committed（已确定）

另请参见 [svcprop\(1\)](#)、[svcs\(1\)](#)、[ipf\(1M\)](#)、[svcadm\(1M\)](#)、[svccfg\(1M\)](#)、[getservbyname\(3SOCKET\)](#)、[rpc\(4\)](#)

引用名	svc.startd – 服务管理工具主重启程序
用法概要	<pre>/lib/svc/bin/svc.startd svc:/system/svc/restarter:default</pre>
描述	svc.startd 是服务管理工具 (Service Management Facility, SMF) 的主重启程序守护进程和所有服务的缺省重启程序。svc.startd 基于管理请求、系统故障或应用程序故障启动、停止和重启服务。 svc.startd 维护服务状态，并负责根据每个服务的依赖项来管理故障。 在系统启动期间会自动调用 svc.startd。如果发生任何故障，它都会重启。绝不当直接调用 svc.startd。 有关所有重启程序通用的配置和行为的信息，请参见 smf_restarter(5)。 svcs(1) 报告由服务管理工具管理的所有服务的状态。svcadm(1M) 允许处理与服务重启程序相关的服务实例。
环境变量	带 "SMF_" 前缀的环境变量是保留的，且可以被覆盖。 svc.startd 将 smf_method(5) 中指定的 "SMF_" 环境变量提供给方法。缺省情况下，PATH 设为 "/usr/sbin:/usr/bin"。缺省情况下，提供给 svc.startd 的所有其他环境变量均继承自 init(1M)。 重复的条目将缩减为单个条目。所使用的值是未定义的。将忽略未带 "<name>=" 前缀的环境条目。
重启程序选项	svc.startd 不是通过命令行选项配置的。其配置是从服务配置系统信息库中读取的。您可以使用 svccfg(1M) 来设置所有选项和属性。 options 属性组中的下列配置变量可供开发者和管理员使用： boot_messages 一个 <i>astring</i>（如 scf_value_is_type 中所定义；请参见 scf_value_create(3SCF)），描述了在引导期间要输出到控制台的消息的缺省级别。支持的消息选项包括 quiet 和 verbose。使用 quiet 选项时，在引导期间向控制台输出的消息最少。verbose 选项针对所启动的每个服务输出一条消息来指示是成功还是失败。您可以使用 boot -m 选项来覆盖引导时的 boot_messages 设置。请参见 kernel(1M)。 logging 控制 svc.startd 的全局服务日志记录级别。一个 <i>astring</i>（如 scf_value_is_type 中所定义；请参见 scf_value_create(3SCF)），描述了要记录到 syslog（请参见 syslog(3C) 和 svc.startd 的全局日志文件 <code>/var/svc/log/svc.startd.log</code>）中的消息的缺省级别。支持的消息选项包括 quiet、verbose 和 debug。quiet 选项将需要管理员干预的错误消息发送到控制台、syslog，以及 svc.startd 的全局日志文件。verbose 选项将需要管理员干预的错误消息发送到控制台、syslog，以及

svc.startd 的全局文件，并将不需要管理员干预的错误的相关信息发送到 svc.startd 的全局日志文件。还会针对所启动的每个服务向控制台发送一条消息。debug 选项将 svc.startd 调试消息发送到 svc.startd 的全局日志文件；将需要管理员干预的错误消息发送到控制台、syslog，以及 svc.startd 的全局日志文件；并针对所启动的每个服务向控制台发送一条消息。

milestone

一个 FMRI，用于确定用作缺省引导级别的里程碑。可接受的选项仅包括以下主要里程碑：

```
svc:/milestone/single-user:default
svc:/milestone/multi-user:default
svc:/milestone/multi-user-server:default
```

或特殊值 all 或 none。all 表示依赖于每个服务的一个理想化里程碑。none 是一个特殊的里程碑，除了主 svc:/system/svc/restarter:default 外，没有任何服务在运行。缺省情况下，svc.startd 使用 all 这个依赖于每个服务的综合性里程碑。如果指定了此属性，它将覆盖 [inittab\(4\)](#) 中的任何 initdefault 设置。

system/reconfigure

指示已请求了重新配置重新引导。包含必须切断重新配置重新引导的操作的服务可以通过检查此属性是否存在以及是否设为 1 来确认是否请求了重新配置引导。

此属性由 svc.startd 管理，且管理员不应对其进行修改。

配置错误（例如禁用 svc.startd）将记录在 syslog 中，但会被忽略。

服务状态

svc.startd 管理的服务可能会显示为 [smf\(5\)](#) 中描述的任一状态。此重启程序未修改状态定义。

服务报告

除了由受管的服务执行的日志记录外，svc.startd 还提供了一组通用的服务报告和日志记录机制。

报告属性 svc.startd 更新它管理的所有服务通用的一组属性。这些属性是可用来基于服务实例运行状况采取措施的通用接口。可以使用 [svcs\(1\)](#) 来方便地显示这些属性。

restarter/state

restarter/next_state

实例的当前和下一个（如果当前正在转换中）状态。

restarter/auxiliary_state

一个详细介绍当前实例状态的更多相关信息的标题。svc.startd 管理的服务的可用辅助状态为：

maintenance

```
fault_threshold_reached
stop_method_failed
administrative_request
```

`restarter/state_timestamp`

访问当前状态的时间。

`restarter/contract`

执行服务实例所依据的主进程合同 ID（如果有）。

日志

缺省情况下，`svc.startd` 会将服务的重要重启程序操作以及方法标准输出和标准错误文件描述符记录到 `/var/svc/log/service:instance.log` 中。在系统全局位置（例如 `/var/svc/log/svc.startd.log` 和 `syslog`）中执行的日志记录的级别是由 `options/logging` 属性控制的。

服务定义

开发或配置 `svc.startd` 管理的服务时，使用一组通用的属性来影响服务实例与重启程序之间的交互。

方法

[smf_method\(5\)](#) 中介绍了 `svc.startd` 提供的 `fork/exec` 模型的方法的一般形式。下列方法是 `svc.startd` 管理的服务支持的必需或可选方法。

refresh 在不中断服务的情况下从系统信息库或 `config` 文件重新装入任何适当的配置参数。这通常是通过对系统守护进程使用 `SIGHUP` 实现的。如果服务在不重新启动的情况下无法识别配置更改，则不提供 `refresh` 方法。

此方法是可选的。

start 开始服务。仅在应用程序可供使用者使用后返回成功。如果某个冲突的实例已在运行或服务无法启动，此方法将失败。

此方法是必需的。

stop 停止服务。在某些情况下，当某些或全部服务已停止时可调用 `stop` 方法。只有在方法返回时服务未完全停止的情况下才返回错误。

此方法是必需的。

如果服务不需要在某个必需的方法中执行任何操作，则它必须为该方法指定 `:true` 标记。

`svc.startd` 会接受为服务或任何特定方法指定的任何方法上下文。[smf_method\(5\)](#) 中描述的方法扩展标记可在 `svc.startd` 调用的所有方法中使用。

属性

[smf\(5\)](#) 中提供了常规属性的概述。这些常规属性与 `svc.startd` 进行交互的具体方式如下所述：

general/enabled

如果将 **enabled** 设为 **true**，重启程序会尝试在服务的所有依赖项都满足后启动该服务。如果设为 **false**，服务将保持为禁用状态，不会运行。

general/restarter

如果此 FMRI 属性为空或设为 **svc:/system/svc/restarter:default**，则服务由 **svc.startd** 进行管理。其他情况下，由指定的重启程序负责（在它可用后）管理该服务。

general/single_instance

如果 **single_instance** 设为 **true**，在任意时间，**svc.startd** 仅允许该服务的一个实例转换为 **online**（联机）或 **degraded**（降级）状态。

此外，**svc.startd** 管理的服务可以在 **startd** 属性组中定义下面列出的可选属性。

startd/duration

duration 属性定义服务的模型。可将其设为 **transient**、**child**（这两者也称为 "wait" 模型服务）或 **contract**（缺省值）。

startd/ignore_error

ignore_error 属性（如果已设置）指定忽略的事件的逗号分隔列表。列表中的合法字符串值为 **core** 和 **signal**。缺省情况下，发生任何错误都会重新启动。

startd/need_session

need_session 属性（如果设为 **true**）指示此实例应在其自己的会话中启动。缺省情况下不是这样。

startd/utmpx_prefix

utmpx_prefix 字符串属性规定实例在启动方法执行前需要一个有效的 **utmpx** 条目。缺省情况下，不创建 **utmpx** 条目。

服务故障

如果一个方法返回非零退出代码或在指定的超时时间到期前没有完成，**svc.startd** 将认为该方法失败。如果返回了 **\$SMF_EXIT_ERR_CONFIG** 或 **\$SMF_EXIT_ERR_FATAL**，**svc.startd** 会立即将服务置于 **maintenance**（维护）状态。对于所有其他故障，**svc.startd** 将服务置于 **offline**（脱机）状态。如果某个服务处于 **offline**（脱机）状态且满足了其依赖项，**svc.startd** 会重新尝试启动该服务（请参见 [smf\(5\)](#)）。

如果在定义的超时时间到期前，合同或瞬态服务没有从其 **start** 方法中返回，**svc.startd** 将向该方法发送 **SIGKILL**，并使该服务返回至 **offline**（脱机）状态。

如果连续发生五次故障，或者服务由于一个错误而每隔 10 分钟重新启动了多次，则 **svc.startd** 会将该服务置于 **maintenance**（维护）状态。

服务故障的状况是由服务模型（由 **startd/duration** 属性定义）和 **startd/ignore_error** 属性值的组合定义的。

如果发生下列任一状况，合同模型服务将失败：

- 服务中的所有进程都退出

- 服务中的任意进程生成核心转储
- 服务外的进程向某个服务进程发送了致命的信号（例如，管理员使用 `pkill` 命令来终止某个服务进程）

可以通过在 `startd/ignore_error` 中指定 `core` 和/或 `signal` 让服务忽略后两种状况。

将某个服务定义为瞬态的意味着 `svc.startd` 不会跟踪该服务的进程。因此，合同模型服务的上述潜在故障不会被认为是瞬态服务的故障。瞬态服务仅在出现某种方法故障状况时进入 `maintenance`（维护）状态。

每当与服务相关联的子进程退出时，“wait”模型的服务都将重新启动。退出的子进程不会被视为“wait”模型的服务的错误，且重复发生的故障不会导致转换为 `maintenance`（维护）状态。

传统服务

`svc.startd` 继续为在启动运行级转换期间调用的服务提供支持。每个 `/etc/rc?.d` 目录是在构成等效的运行级里程碑的所有受管理服务转换为 `online`（联机）状态后处理的。标准 `init` 脚本放置在 `/etc/rc?.d` 目录中的脚本按照其序列号顺序运行。

里程碑到运行级的映射为：

<code>milestone/single-user</code>	单用户 (S)
<code>milestone/multi-user</code>	多用户 (2)
<code>milestone/multi-user-server</code>	使用网络服务的多用户 (3)

此外，通过将每个脚本的一个实例插入到系统信息库中，`svc.startd` 使得这些传统服务在 SMF 中可见。可以使用标准 SMF 接口（例如 `svcs(1)`）来访问这些传统实例，并且它们始终显示为 `LEGACY-RUN` 状态，无法修改，且无法指定为其他服务的依赖项。为了给管理员提供便利，捕获了传统服务的初始启动时间。

文件

<code>/var/svc/log</code>	<code>svc.startd</code> 用于存储日志文件的目录。
<code>/etc/svc/volatile</code>	在引导过程的早期阶段，将 <code>/var</code> 挂载为可读写之前， <code>svc.startd</code> 用于存储日志文件的目录。

示例

示例1 启用详细日志记录

要启用详细日志记录，请键入以下命令：

```
# /usr/sbin/svccfg -s system/svc/restarter:default
svc:/system/svc/restarter:default> addpg options application
svc:/system/svc/restarter:default> setprop options/logging = \
astring: verbose
svc:/system/svc/restarter:default> exit
```

此请求将在 `svc.startd` 下次重新启动时生效。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见 [svcs\(1\)](#)、[svcprop\(1\)](#)、[kernel\(1M\)](#)、[init\(1M\)](#)、[svcadm\(1M\)](#)、[svccfg\(1M\)](#)、[svc.configd\(1M\)](#)、

引用名 swap – 交换管理接口

用法概要

```
/usr/sbin/swap -a swapname [swaplow] [swaplen]
/usr/sbin/swap -d swapname [swaplow]
/usr/sbin/swap -l [-h | -k]
/usr/sbin/swap -s [-h]
```

描述 swap 实用程序提供了一种添加、删除和监视内存管理器使用的系统交换区域的方法。

选项 所支持的选项如下：

-a swapname [swaplow] [swaplen]

添加指定的交换区域。只有超级用户或承担主管理员角色的用户可以使用此选项。*swapname* 是交换区域或常规文件的名称。例如，在运行 UFS 根文件系统的系统上，可以指定分片（如 `/dev/dsk/c0t0d0s1`）或常规文件作为交换区域。在运行 ZFS 文件系统的系统上，可以指定 ZFS 卷（如 `/dev/zvol/dsk/rpool/swap`）作为交换区域。在 ZFS 文件系统上不支持使用常规文件作为交换区域。此外，当系统运行 ZFS 根文件系统时，不能将同一个 ZFS 卷同时用作交换区域和转储设备。

swaplow 是文件内以 512 字节块为单位的偏移量，交换区域应从此偏移量位置开始。*swaplen* 是交换区域所需的长度，以 512 字节块为单位。*swaplen* 的值不能小于 16。例如，如果指定了 *n* 个块，则 $(n-1)$ 个块是实际的交换长度。*swaplen* 的长度必须至少为一个页面。可以使用 `pagesize` 命令确定内存页面的大小。请参见 [pagesize\(1\)](#)。由于交换文件的第一个页面会被自动跳过，并且交换文件的长度至少需要为一个页面，因此，其最小大小应该是 2 个页面大小字节的倍数。内存页面的大小取决于计算机。

swaplow + *swaplen* 必须小于等于交换文件的大小。如果未指定 *swaplen*，将添加一个始于 *swaplow* 并扩展到指定文件的末尾的区域。如果 *swaplow* 和 *swaplen* 都未指定，将使用除第一个页面外的整个文件。交换区域通常在启动系统的过程中通过 `/usr/sbin/swapadd` 脚本自动添加。该脚本会添加 `/etc/vfstab` 文件中指定的所有交换区域；有关这些规范的语法，请参见 [vfstab\(4\)](#)。

通过在 [vfstab\(4\)](#) 中指定 `encrypted` 选项，以及指定 ZFS 卷的 `encryption` 属性，可以对用作交换设备的 ZFS 卷进行加密。请参见 [zfs\(1M\)](#)。

要使用 NFS 或本地文件系统 *swapname*，首先应该使用 [mkfile\(1M\)](#) 来创建一个文件。现在仅通过运行 `swap -a` 命令即可将一个本地文件系统交换文件添加到正在运行的系统中。对于已挂载的 NFS 交换文件，服务器需要导出该文件。要进行导出，请执行以下步骤：

1. 向 `/etc/dfs/dfstab` 中添加以下行：

```
share -F nfs -o \
    rw=clientname,root=clientname path-to-swap-file
```

2. 运行 [shareall\(1M\)](#)。

3. 使客户机向 `/etc/vfstab` 中添加以下行：

```
server:path-to-swap-file - local-path-to-swap-file nfs \
- - - local-path-to-swap-file - - swap - - -
```

4. 使客户机运行 mount :

```
# mount local-path-to-swap-file
```

5. 然后, 客户机可以运行 swap -a 来添加交换空间 :

```
# swap -a local-path-to-swap-file
```

-d *swapname*

删除指定的交换区域。只有超级用户可以使用此选项。*swapname* 是交换文件（如 */dev/dsk/c0t0d0s1*）或常规文件的名称。*swaplow* 是要删除的交换区域的偏移量, 以 512 字节块为单位。如果未指定 *swaplow*, 将从第二个页面开始删除区域。命令完成后, 将无法再分配该区域中的交换块, 并且该交换区域中之前正在使用的所有交换块已移动到其他交换区域。

-h

所有大小将调整为用户可阅读的格式。调整是通过反复除以 1024 完成的。

-k

以 1024 字节为单位写出文件大小。

-l

列出所有交换区域的状态。输出包含五列:

path

交换区域的路径名。

dev

如果设备不是块特殊设备, 则此列为十进制的主/从设备号; 否则为零。

swaplo

区域的 *swaplow* 值, 以 512 字节块为单位。

blocks

区域的 *swaplen* 值, 以 512 字节块为单位。

free

此区域中当前没有分配的 512 字节块的数量。

此列表不包括物理内存形式的交换空间, 因为该空间不与特定的交换区域相关。

如果在删除 *swapname* (通过 *swap -d*) 的过程中运行 *swap -l*, 则会在交换状态的第 6 列显示字符串 INDEL。

-s

输出有关总交换空间的使用情况和可用性的摘要信息:

allocated

当前分配的用作后备存储的交换空间总量, 以字节为单位。

reserved
当前未分配但内存映射声明将来可能使用的交换空间总量，以字节为单位。

used
分配的或保留的交换空间总量，以字节为单位。

available
当前可用于将来保留和分配的交换空间总量，以字节为单位。

这些数字包括由 `-l` 选项列出的所有已配置的交换区域中的交换空间，以及处于物理内存形式的交换空间。

用法
在 32 位操作系统上，对于大于等于 2 GB 的交换设备，只会使用前 2 GB -1 的空间。在 64 位操作系统上，对于大于 2 GB 的块设备，最多可以将其 $2^{63}-1$ 个字节完全用作交换空间。

环境变量
有关影响 `swap` 执行的环境变量 `LC_CTYPE` 和 `LC_MESSAGE` 的描述，请参见 [environ\(5\)](#)。

属性
有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os

另请参见
[pagesize\(1\)](#)、[mkfile\(1M\)](#)、[shareall\(1M\)](#)、[zfs\(1M\)](#)、[getpagesize\(3C\)](#)、[vfstab\(4\)](#)、[attributes\(5\)](#)

附注
有关使用 ZFS 设置交换区域的信息，请参见“ZFS 管理指南”。

警告
不会进行检查以确定要添加的交换区域是否与现有文件系统重叠。

引用名	sync – update the super block
用法概要	sync
描述	sync executes the sync system primitive. If the system is to be stopped, sync must be called to ensure file system integrity. It will flush all previously unwritten system buffers out to disk, thus assuring that all file modifications up to that point will be saved. See sync(2) for details.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [sync\(2\)](#), [attributes\(5\)](#)

引用名

syncinit – set serial line interface operating parameters

用法概要

```
/usr/sbin/syncinit device
[ [baud_rate] | [keyword=value,...] | [single-word option]]
```

描述

The syncinit utility allows the user to modify some of the hardware operating modes common to synchronous serial lines. This can be useful in troubleshooting a link, or necessary to the operation of a communications package.

If run without options, syncinit reports the options as presently set on the port. If options are specified, the new settings are reported after they have been made.

选项

Options to syncinit normally take the form of a keyword, followed by an equal sign and a value. The exception is that a baud rate may be specified as a decimal integer by itself. Keywords must begin with the value shown in the options table, but may contain additional letters up to the equal sign. For example, loop= and loopback= are equivalent.

The following options are supported:

Keyword	Value	Effect
loop	yes	Set the port to operate in internal loopback mode. The receiver is electrically disconnected from the DCE receive data input and tied to the outgoing transmit data line. Transmit data is available to the DCE. The Digital Phase-Locked Loop (DPLL) may not be used as a clock source in this mode. If no other clocking options have been specified, perform the equivalent of txc=baud and rxc=baud.
	no	Disable internal loopback mode. If no other clocking options have been specified, perform the equivalent of txc=txc and rxc=rxc.
echo	yes	Set the port to operate in auto-echo mode. The transmit data output is electrically disconnected from the transmitter and tied to the receive data input. Incoming receive data is still visible. Use of this mode in combination with local loopback mode has no value, and should be rejected by the device driver. The auto-echo mode is useful to make a system become the endpoint of a remote loopback test.
	no	Disable auto-echo mode.
nrzi	yes	Set the port to operate with NRZI data encoding.
	no	Set the port to operate with NRZ data encoding.
txc	txc	Transmit clock source will be the TxC signal (pin 15).
	rxc	Transmit clock source will be the RxC signal (pin 17).
	baud	Transmit clock source will be the internal baud rate generator.
	pll	Transmit clock source will be the output of the DPLL circuit.

rx	rx	Receive clock source will be the Rx signal (pin 17).
	tx	Receive clock source will be the Tx signal (pin 15).
	baud	Receive clock source will be the internal baud rate generator.
	pll	Receive clock source will be the output of the DPLL circuit.
speed	integer	Set the baud rate to <i>integer</i> bits per second.

There are also several single-word options that set one or more parameters at a time:

Keyword	Equivalent to Options:
external	tx=tx rx=rx loop=no
sender	tx=baud rx=rx loop=no
internal	tx=pll rx=pll loop=no
stop	speed=0

示例

示例1 Using syncinit

The following command sets the first CPU port to loop internally, using internal clocking and operating at 38400 baud:

```
example# syncinit zsh0 38400 loop=yes
device: /dev/zsh ppa: 0
speed=38400, loopback=yes, echo=no, nrzi=no, tx=baud, rx=baud
```

The following command sets the same port's clocking, local loopback and baud rate settings to their default values:

```
example# syncinit zsh0 stop loop=no
device: /dev/zsh ppa: 0
speed=0, loopback=no, echo=no, nrzi=no, tx=tx, rx=rx
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见
诊断

[syncloop\(1M\)](#), [syncstat\(1M\)](#), [Intro\(2\)](#), [ioctl\(2\)](#), [attributes\(5\)](#), [zsh\(7D\)](#)

device missing minor device number The name *device* does not end in a decimal number that can be used as a minor device number.

bad speed: <i>arg</i>	The string <i>arg</i> that accompanied the speed= option could not be interpreted as a decimal integer.
Bad arg: <i>arg</i>	The string <i>arg</i> did not make sense as an option.
ioctl failure code = <i>errno</i>	An ioctl(2) system call failed. The meaning of the value of <i>errno</i> may be found in Intro(2) .

警告 Do not use syncinit on an active serial link, unless needed to resolve an error condition. Do not use this command casually or without being aware of the consequences.

引用名 syncloop – synchronous serial loopback test program

用法概要 /usr/sbin/syncloop [-cdlstv] *device*

描述 The syncloop command performs several loopback tests that are useful in exercising the various components of a serial communications link.

Before running a test, syncloop opens the designated port and configures it according to command line options and the specified test type. It announces the names of the devices being used to control the hardware channel, the channel number (ppa) corresponding to the *device* argument, and the parameters it has set for that channel. It then runs the loopback test in three phases.

The first phase is to listen on the port for any activity. If no activity is seen for at least four seconds, syncloop proceeds to the next phase. Otherwise, the user is informed that the line is active and that the test cannot proceed, and the program exits.

In the second phase, called the "first-packet" phase, syncloop attempts to send and receive one packet. The program will wait for up to four seconds for the returned packet. If no packets are seen after five attempts, the test fails with an exhorting message. If a packet is returned, the result is compared with the original. If the length and content do not match exactly, the test fails.

The final phase, known as the "multiple-packet" phase, attempts to send many packets through the loop. Because the program has verified the integrity of the link in the first-packet phase, the test will not fail after a particular number of timeouts. If a packet is not seen after four seconds, a message is displayed. Otherwise, a count of the number of packets received is updated on the display once per second. If it becomes obvious that the test is not receiving packets during this phase, the user may wish to stop the program manually. The number and size of the packets sent during this phase is determined by default values, or by command line options. Each returned packet is compared with its original for length and content. If a mismatch is detected, the test fails. The test completes when the required number of packets have been sent, regardless of errors.

After the multiple-packet phase has completed, the program displays a summary of the hardware event statistics for the channel that was tested. The display takes the following form:

CRC errors	Aborts	Overruns	Underruns	In<-Drops->	Out
0	0	0	0	0	0

This is followed by an estimated line speed, which is an approximation of the bit rate of the line, based on the number of bytes sent and the actual time that it took to send them.

选项 The options for syncloop are described in the following table:

Option	Parameter	Default	Description
-c	<i>packet_count</i>	100	Specifies the number of packets to be sent in the multiple-packet phase.
-d	<i>hex_data_byte</i>	<i>random</i>	Specifies that each packet will be filled with bytes with the value of <i>hex_data_byte</i> .
-l	<i>packet_length</i>	100	Specifies the length of each packet in bytes.
-s	<i>line_speed</i>	9600	Bit rate in bits per second.
-v			Sets verbose mode. If data errors occur, the expected and received data is displayed.
-t	<i>test_type</i>	<i>none</i>	A number, from 1 to 4, that specifies which test to perform. The values for <i>test_type</i> are as follows: 1: Internal loopback test. Port loopback is on. Transmit and receive clock sources are internal (baud rate generator). 2: External loopback test. Port loopback is off. Transmit and receive clock sources are internal. Requires a loopback plug suitable to the port under test. 3: External loopback test. Port loopback is off. Transmit and receive clock sources are external (modem). Requires that one of the local modem, the remote modem, or the remote system be set in a loopback configuration. 4: Test using predefined parameters. User defines hardware configuration and may select port parameters using the syncinit(1M) command.

All numeric options except -d are entered as decimal numbers (for example, -s 19200). If you do not provide the -t *test_type* option, syncloop prompts for it.

示例

示例 1 A sample display of using the syncloop command.

In the following command syncloop uses a packet length of 512 bytes over the first CPU port:

```
example# syncloop -l 512 zsh0
```

In response to the above command, syncloop prompts you for the test option you want.

The following command performs an internal loopback test on the first CPU port, using 5000 packets and a bit rate of 56Kbps:

```
example# syncloop -t 1 -s 56000 -c 5000 zsh0
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见
诊断

syncinit(1M) , syncstat(1M) , attributes(5) , zsh(7D)	
<i>device</i> missing minor device number	The name <i>device</i> does not end in a decimal number that can be used as a minor device number.
invalid packet length: <i>nnn</i>	The packet length was specified to be less than zero or greater than 4096.
poll: nothing to read	
poll: nothing to read or write.	The poll(2) system call indicates that there is no input pending and/or that output would be blocked if attempted.
len <i>xxx</i> should be <i>yyy</i>	The packet that was sent had a length of <i>yyy</i> , but was received with a length of <i>xxx</i> .
<i>nnn</i> packets lost in outbound queueing	
<i>nnn</i> packets lost in inbound queueing	A discrepancy has been found between the number of packets sent by syncloop and the number of packets the driver counted as transmitted, or between the number counted as received and the number read by the program.

警告

To allow its tests to run properly, as well as prevent disturbance of normal operations, syncloop should only be run on a port that is not being used for any other purpose at that time.

引用名 syncstat – report driver statistics from a synchronous serial link

用法概要 /usr/sbin/syncstat [-c] device [interval]

描述 The syncstat command reports the event statistics maintained by a synchronous serial device driver. The report may be a single snapshot of the accumulated totals, or a series of samples showing incremental changes. Prior to these it prints the device name being used to query a particular device driver, along with a number indicating the channel number (ppa) under control of that driver.

Event statistics are maintained by a driver for each physical channel that it supports. They are initialized to zero at the time the driver module is loaded into the system, which may be either at boot time or when one of the driver's entry points is first called.

The *device* argument is the name of the serial device as it appears in the /dev directory. For example, *zsh0* specifies the first on-board serial device.

The following is a breakdown of syncstat output:

speed	The line speed the device has been set to operate at. It is the user's responsibility to make this value correspond to the modem clocking speed when clocking is provided by the modem.
ipkts	The total number of input packets.
opkts	The total number of output packets.
undrun	The number of transmitter underrun errors.
ovrrun	The number of receiver overrun errors.
abort	The number of aborted received frames.
crc	The number of received frames with CRC errors.
isize	The average size (in bytes) of input packets.
osize	The average size (in bytes) of output packets.

选项 -c Clear the accumulated statistics for the device specified. This may be useful when it is not desirable to unload a particular driver, or when the driver is not capable of being unloaded.

interval syncstat samples the statistics every *interval* seconds and reports incremental changes. The output reports line utilization for input and output in place of average packet sizes. These are the relationships between bytes transferred and the baud rate, expressed as percentages. The loop repeats indefinitely, with a column heading printed every twenty lines for convenience.

示例

示例 1 Sample output from the syncstat command:

```
example# syncstat zsh0

speed ipkts opkts undrun ovrrun abort crc isize osize
9600 15716 17121 0 0 1 3 98 89

example# syncstat -c zsh0

speed ipkts opkts undrun ovrrun abort crc isize osize
9600 0 0 0 0 0 0 0 0
```

In the following sample output a new line of output is generated every five seconds:

```
example# syncstat zsh0 5

ipkts opkts undrun ovrrun abort crc iutil outil
12 10 0 0 0 0 5% 4%
22 60 0 0 0 0 3% 90%
36 14 0 0 0 1 51% 2%
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[syncinit\(1M\)](#), [syncloop\(1M\)](#), [attributes\(5\)](#), [zsh\(7D\)](#)

诊断

- bad interval: *arg*

The argument *arg* is expected to be an interval and could not be understood.
- device* missing minor device number

The name *device* does not end in a decimal number that can be used as a minor device number.
- baud rate not set

The *interval* option is being used and the baud rate on the device is zero. This would cause a divide-by-zero error when computing the line utilization statistics.

警告

Underrun, overrun, frame-abort, and CRC errors have a variety of causes. Communication protocols are typically able to handle such errors and initiate recovery of the transmission in which the error occurred. Small numbers of such errors are not a significant problem for most protocols. However, because the overhead involved in recovering from a link error can be much greater than that of normal operation, high error rates can greatly degrade overall link throughput. High error rates are often caused by problems in the link hardware, such as

cables, connectors, interface electronics or telephone lines. They may also be related to excessive load on the link or the supporting system.

The percentages for input and output line utilization reported when using the *interval* option may occasionally be reported as slightly greater than 100% because of inexact sampling times and differences in the accuracy between the system clock and the modem clock. If the percentage of use greatly exceeds 100%, or never exceeds 50%, then the baud rate set for the device probably does not reflect the speed of the modem.

引用名	sysconfig – unconfigure or reconfigure a Solaris instance
用法概要	<pre> /usr/sbin/sysconfig /usr/sbin/sysconfig configure [-s] [-c <i>config_profile.xml</i> <i>dir</i>] [--destructive] [-g <i>system</i>] /usr/sbin/sysconfig unconfigure [-s] [--destructive] [-g <i>system</i>] /usr/sbin/sysconfig create-profile [-o <i>output_file</i> [-l <i>logfile</i>] [-v <i>verbosity</i>] [-b] [-g <i>system</i>] </pre>
描述	The <code>sysconfig</code> utility is the interface for unconfiguring and reconfiguring a Solaris instance. A Solaris instance is defined as a boot environment in either a global or a non-global zone. There are three operations that are performed using the <code>sysconfig</code> utility: unconfiguration, configuration, and profile creation. When <code>sysconfig</code> is called with the <code>unconfigure</code> subcommand, the system is unconfigured and left in an unconfigured state. System configuration can occur either interactively or non-interactively. If the <code>configure</code> sub-command is invoked without a profile, the SCI Tool is activated and walks the user through the system configuration process. If the <code>configure</code> subcommand is invoked with a profile, then the configuration reads the profile and the configuration occurs non-interactively. The result in either case is a new system configuration. The <code>sysconfig</code> command can also be used to generate a system configuration profile using the <code>create-profile</code> subcommand. The resulting profile is used with the <code>sysconfig configure</code> subcommand to configure systems non-interactively. Valid profile names include an <code>.xml</code> extension. Configuration of a system can be performed either interactively, using the System Configuration Interactive (SCI) Tool, or non-interactively, using a system configuration profile. The SCI tool configures the target system in an interactive way using a text user interface. It can also be used to collect information generated by the user that describes the desired configuration of the target system. The tool then generates a system configuration profile containing the desired system configuration. The SCI tool supports configuration of freshly installed or unconfigured systems. It is designed to provide system configuration for newly created non-global zones and during text installation. If there is a need to modify the configuration of an already configured system utilizing SCI tool, such a system has to be unconfigured first before SCI tool can run.

The functional groupings that can be configured on a system are network, location, users, identity, and kbd_layout. Groupings can also be unconfigured and left in an unconfigured state. The default values for unconfigured groupings are shown below.

The following groupings are configurable.

Grouping	Components	Unconfigured State
identity	system nodename	unknown
kbd_layout	Keyboard	U.S. English
network	network	No network
location	timezone locale	UTC C locale
users	root initial user account	Empty root password Remove user account
naming_services	DNS, NIS and LDAP clients, nsswitch	No network naming services

子命令

This section describes supported subcommands and their associated options.

`unconfigure [-s] [-g system] [--destructive]`

Unconfigure a system and leave it in the unconfigured state.

`-s`

Shut the system down after the unconfiguration completes.

`-g system`

If `-g` is not specified, the user will be queried for confirmation before *system* configuration occurs.

`--destructive`

Do not preserve system data that is normally preserved during unconfiguration. By specifying this flag, the user indicates to any groupings unconfigured that data they would ordinarily preserve might be deleted.

`configure [-s] [-g system] [-c config_profile.xml | dir] [--destructive]`

Configure or reconfigure a system. The configure subcommand has access to the same options as the unconfigure subcommand. It also includes the following additional option.

`-c config_profile.xml | dir`

Provides a profile or a directory of profiles to apply during configuration. If a profile is applied, the configuration step occurs non-interactively. If no profile is provided, the interactive system configuration tool is used for the configuration of the system.

All profiles must have an `.xml` file extension.

If you supply a directory to `-c`, all profiles in that directory must be valid (correctly formed) configuration profiles.

`create-profile [-o output_file [-l logfile] [-v verbosity] [-b] [-g grouping,...]`

Run the SCI tool and create a system configuration profile. The default location for the profile is `/system/volatile/profile/sc_profile.xml`. The configuration generated is not applied to the system.

`-o output_file`

Replace the default profile location with *output_file* for the configuration profile.

`-l`

Location of the log file. Used only with `-o` option.

`-v`

Verbosity level. Used only with `-o` option.

`-b`

Black and white version of SCI tool. Used only with `-o` option.

示例

示例 1 Unconfiguring and Shutting Down

The following command unconfigures the system and leaves it in an unconfigured state. By default, if no grouping is specified, the groupings for the whole system are unconfigured.

```
# sysconfig unconfigure -s
```

示例 2 Unconfiguring the System

The following command unconfigures the system groupings and leaves the system unconfigured.

```
# sysconfig unconfigure -g system
```

示例 3 Reconfiguring System Using SCI Tool

The following command brings up the SCI Tool to reconfigure a system.

```
# sysconfig configure
```

示例 4 Reconfiguring Using a Profile

The following command reconfigures a system using a profile.

```
# sysconfig configure -c some_profile.xml
```

示例 5 Creating and Using a Profile

The following sequence of commands creates a profile, then uses it to reconfigure a system.

```
# sysconfig create-profile -o /tmp/myprofile.xml
```

```
# sysconfig configure -g system -c /tmp/myprofile.xml
```

示例 6 Configuring the System in a Zone

The following command configures the system in a zone.

```
# zlogin ZONENAME
root@ZONENAME# sysconfig configure -g system
```

退出状态

```
0
    Success.

>0
    Failure.
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
	system/install
	system/install/configuration
	system/library/install
Interface Stability	Committed

另请参见

[svcprop\(1\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [attributes\(5\)](#), [attributes\(5\)](#)

引用名	sysdef – output system definition
用法概要	<pre>/usr/sbin/sysdef [-i] [-n <i>namelist</i>] /usr/sbin/sysdef [-h] [-d] [-i] [-D]</pre>
描述	The <code>sysdef</code> utility outputs the current system definition in tabular form. It lists all hardware devices, as well as pseudo devices, system devices, loadable modules, and the values of selected kernel tunable parameters. It generates the output by analyzing the named bootable operating system file (<i>namelist</i>) and extracting the configuration information from it. The default system <i>namelist</i> is <code>/dev/kmem</code>.
选项	-i Prints the configuration information from <code>/dev/kmem</code>. This is the default and only needs to be specified if the configuration information from both <code>/dev/kmem</code> and the system file specified with the “-n <i>namelist</i>” option is needed. -n <i>namelist</i> Specifies a <i>namelist</i> other than the default (<code>/dev/kmem</code>). The <i>namelist</i> specified must be a valid bootable operating system. -h Prints the identifier of the current host in hexadecimal. If <code>sysdef -h</code> is executed within a non-global zone and the zone emulates a host identifier, then the zone's host identifier is printed. This numeric value is not guaranteed to be unique. -d The output includes the configuration of system peripherals formatted as a device tree. -D For each system peripheral in the device tree, display the name of the device driver used to manage the peripheral.

示例

示例1 Sample output format

The following example displays the format of the `sysdef -d` output:

```
example% sysdef -d
Node 'SUNW,Ultra-5_10', unit #-1
  Node 'packages', unit #-1 (no driver)
    Node 'terminal-emulator', unit #-1 (no driver)
    Node 'deblocker', unit #-1 (no driver)
    Node 'obp-tftp', unit #-1 (no driver)
    Node 'disk-label', unit #-1 (no driver)
    Node 'SUNW,builtin-drivers', unit #-1 (no driver)
    Node 'sun-keyboard', unit #-1 (no driver)
    Node 'ufs-file-system', unit #-1 (no driver)
  Node 'chosen', unit #-1 (no driver)
  Node 'openprom', unit #-1 (no driver)
  Node 'client-services', unit #-1 (no driver)
```

示例1 Sample output format (续)

```
Node 'options', unit #0
Node 'aliases', unit #-1 (no driver)
Node 'memory', unit #-1 (no driver)
Node 'virtual-memory', unit #-1 (no driver)
Node 'pci', unit #0
  Node 'pci', unit #0
    Node 'ebus', unit #0
      Node 'auxio', unit #-1 (no driver)
      Node 'power', unit #0
      Node 'SUNW,pll', unit #-1 (no driver)
      Node 'se', unit #0 (no driver)
      Node 'su', unit #0
      Node 'su', unit #1
      Node 'ecpp', unit #-1 (no driver)
      Node 'fdthree', unit #0
      Node 'eeprom', unit #-1 (no driver)
      Node 'flashprom', unit #-1 (no driver)
      Node 'SUNW,CS4231', unit #0 (no driver)
    Node 'network', unit #0
    Node 'SUNW,m64B', unit #0
    Node 'ide', unit #0
      Node 'disk', unit #-1 (no driver)
      Node 'cdrom', unit #-1 (no driver)
      Node 'sd', unit #1
      Node 'dad', unit #1
    Node 'pci', unit #-1 (no driver)
  Node 'SUNW,UltraSPARC-IIi', unit #-1 (no driver)
Node 'pseudo', unit #0
```

[output truncated]

文件 /dev/kmem default operating system image

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	system/core-os

另请参见 [hostid\(1\)](#), [prtconf\(1M\)](#), [nlist\(3ELF\)](#), [attributes\(5\)](#), [zones\(5\)](#)

引用名

syseventadm – sysevent event specification administration

用法概要

```
syseventadm add [-R rootdir] [-v vendor] [-p publisher]
                [-c class] [-s subclass] [-u username] path [args]

syseventadm remove [-R rootdir] [-v vendor] [-p publisher]
                   [-c class] [-s subclass] [-u username] [path [args]]

syseventadm list [-R rootdir] [-v vendor] [-p publisher]
                 [-c class] [-s subclass] [-u username] [path [args]]

syseventadm restart
```

描述

The `syseventadm` command is an administrative front-end to add, remove and list sysevent event handlers. You can also restart the sysevent daemon by use of the restart command. `syseventadm` can only be run by root.

The `syseventadm add` command adds a handler for a sysevent event specified by at least one of *vendor*, *publisher* or *class*. If *class* is specified, it may be qualified with a *sub-class*. Only the values specified for *vendor*, *publisher*, *class* and *sub-class* when adding the handler are matched against sysevent events to determine if the specification matches the event and the handler should be run. *path* is the full pathname of the command to be run in response to matching events, with optional arguments (*args*). If *username* is specified, the command is invoked as user *username*, otherwise as root.

The `syseventadm remove` command removes handlers for matching sysevent event specifications. Event specifications may be matched by specifying at least one of *vendor*, *publisher*, *class*, *username* or *path*. If *class* is specified, it may be qualified with a *sub-class*. Any of *vendor*, *publisher*, *class*, *sub-class*, *username*, *path* or *args* not specified match the corresponding fields of all events. Handlers for all matching specifications are removed.

The `syseventadm list` command lists the handlers for matching sysevent event specifications using the same match criteria as the remove command but without the requirement that at least one of *vendor*, *publisher*, *class*, *username* or *path* be specified. With no match criteria, all specifications are listed. The list command output format is: `[vendor=vendor] [publisher=publisher] [class=class] [subclass=subclass] [username=username] path [args]` where each of *class*, *sub-class*, *vendor*, *publisher* and *username* is listed only if part of the match criteria for the listed specification.

The `syseventadm restart` command informs the syseventd daemon to reread the sysevent registry after a change has been made by adding or removing one or more sysevent handler specifications.

Argument Macro Substitution

The sysevent handling facility provides extensive macro capability for constructing the command line arguments to be executed in response to an event. Macro expansion applies only to the command line *args* specified for an event handler, with macros expanded with data from the event itself. Pre-defined macros are provided for the event *class*, *subclass*, *publisher* and *vendor* information. Macros not matching one of the pre-defined macro names cause the

attribute list attached to the event to be searched for an attribute of that name, with the value of the matching attribute substituted on the command line.

Macros are introduced by the `$` character, with the macro name being the following token separated by a SPACE or TAB character. If the macro name is embedded in text, it may be delineated by `${` and `}`. A `\` before the `$` causes macro expansion not to occur.

<code>\$class</code>	The class string defining the event
<code>\$publisher</code>	The publisher string defining the event
<code>\$sequence</code>	The sequence number of the event.
<code>\$subclass</code>	The subclass string defining the event
<code>\$timestamp</code>	The timestamp of the event.
<code>\$vendor</code>	The vendor string defining the event

Macro names other than those pre-defined are compared against the attribute list provided with the event. An attribute with name matching the macro name causes the value of the attribute to be substituted as ASCII text on the generated command line.

Use of a macro for which no attribute with that name is defined, or for which multiple attributes with that name are provided, cause an error and the command is not invoked.

Attributes with signed data types (`DATA_TYPE_INT16`, `DATA_TYPE_INT32` and `DATA_TYPE_INT64`) are expanded as decimal digits.

Attributes with unsigned data types (`DATA_TYPE_BYTE`, `DATA_TYPE_UINT16`, `DATA_TYPE_UINT32`, `DATA_TYPE_UINT64` and `DATA_TYPE_HTTPTIME`) are expanded as hexadecimal digits with a `0x` prefix.

Attributes with string data type (`DATA_TYPE_STRING`) are expanded with the string data. The data is not quoted. If it is desired that the quoted strings be generated on the command line, put quotes around the macro call in the arguments.

Array types are expanded with each element expanded as defined for that scalar type, with a space separating each element substitution.

选项

The `add`, `list` and `remove` subcommands support the following options:

- `-c class` Specify the event class, *class*.
- `-p publisher` Specify the event publisher, *publisher*.
- `-R rootdir` Specify an alternate root path, *rootdir*.

注 – The root file system of any non-global zones must not be referenced with the `-R` option. Doing so might damage the global zone's file system, might

compromise the security of the global zone, and might damage the non-global zone's file system. See [zones\(5\)](#).

- s *subclass* Specify the event subclass, *subclass*.
- u *username* Specify the username (*username*) to invoke the command.
- v *vendor* Specify the vendor (*vendor*) that defines the event. Events defined by third-party software should specify the company's stock symbol as *vendor*. Sun-defined events use SUNW.

操作数

The add, list and remove subcommands support the following options:

args Command arguments

path Full path of command to be run in response to event

示例

示例 1 Adding an Event Handler

The following example adds an event handler for an event defined by vendor MYCO ("My Company"), class EC_ENV and sub-class ESC_ENV_TEMP. The command to be run is /opt/MYCOenv/bin/ec_env_temp, with arguments being the class name, sub-class name and pathname derived from the event attributes. The \$ characters are preceded by a backslash to circumvent shell interpretation. There is no need to restart the service after the change since the registry is maintained on \$ALTRoot.

```
# syseventadm add -R \${ALTRoot} -v MYCO -c EC_ENV -s ESC_ENV_TEMP \
/opt/MYCOenv/bin/ec_env_temp \${class} \${subclass} \${pathname}
```

Note the caveat on the use of the -R option in the description of that option, above.

示例 2 Removing an Event Handler

The following example removes the event handler added in Example 1.

```
# syseventadm remove -R \${ALTRoot} -v MYCO -c EC_ENV -s ESC_ENV_TEMP \
/opt/MYCOenv/bin/ec_env_temp \${class} \${subclass} \${pathname}
```

Note the caveat on the use of the -R option in the description of that option, above.

示例 3 Listing Event Handlers

The following example lists all event handlers for events of class EC_ENV, subclass ESC_ENV_TEMP, as defined by vendor MYCO:

```
# syseventadm list -v MYCO -c EC_ENV -s ESC_ENV_TEMP \
vendor=MYCO class=EC_ENV subclass=ESC_ENV_TEMP \
/opt/MYCOenv/bin/ec_env_temp \${class} \${subclass} \${pathname}
```

示例 4 Listing Event Handlers

The following example lists all event handlers defined by vendor VRTS.

```
# syseventadm list -v VRTS
```

示例 5 Removing Event Handlers

The following example removes all event handlers defined by vendor VRTS, and restarts service.

```
# syseventadm remove -v VRTS
# syseventadm restart
```

示例 6 Listing All Event Handlers Specified to Run a Command

The following example lists all event handlers specified to run the command /opt/MYCOenv/bin/ec_env_temp:

```
# syseventadm list /opt/MYCOenv/bin/ec_env_temp
```

示例 7 Removing Event Handlers and Restarting Service

The following example removes all event handlers specified to run the command /opt/MYCOenv/bin/ec_env_temp, and restarts service:

```
# syseventadm remove /opt/MYCOenv/bin/ec_env_temp
# syseventadm restart
```

退出状态

The following exit values are returned:

- 0 Successful completion.
- 1 No matching event specification found (remove or list commands only).
- 2 Incorrect command usage.
- 3 Permission denied.
- 4 Command failed.
- 5 Out of memory.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[syseventd\(1M\)](#), [sysevent_post_event\(3SYSEVENT\)](#), [attributes\(5\)](#), [ddi_log_sysevent\(9F\)](#)

附注

To avoid upgrade problems, packages delivering a sysevent event handler should install the event handler by running `syseventadm` from the package's `postinstall` script. The event handler can then be removed by running `syseventadm` from the package's `preremove` script using the same arguments as when added.

引用名	syseventconfd – kernel system event command invocation daemon				
用法概要	/usr/lib/sysevent/syseventconfd [-r <i>rootdir</i>]				
描述	syseventconfd is the user-level daemon that invokes user-level commands in response to kernel system events received from syseventd(1M) .				
选项	The following options are supported: -r <i>rootdir</i> Cause syseventconfd to use an alternate root path when creating its door. The root path must match the root path used to invoke syseventd.				
文件	/etc/sysevent/syseventconfd_event_service syseventconfd event service door file /usr/lib/sysevent/modules/sysevent_conf_mod.so syseventd loadable module (SLM) managing sysevent.conf files				
属性	See attributes(5) for descriptions of the following attributes:				
<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>system/core-os</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	system/core-os
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	system/core-os				
另请参见	svcs(1) , svcadm(1M) , syseventd(1M) , attributes(5) , smf(5)				
附注	The syseventconfd service is managed by the service management facility, smf(5), under the service identifier: svc:/system/sysevent:default Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using svcadm(1M). The service's status can be queried using the svcs(1) command.				

引用名	syseventd – kernel system event notification daemon
用法概要	<code>/usr/lib/sysevent/syseventd [-d <i>debug_level</i>] [-r <i>rootdir</i>]</code>
描述	syseventd is a user-level daemon that accepts delivery of system event buffers from the kernel. Once an event buffer has been delivered to syseventd, it, in turn, attempts to propagate the event to all interested end event subscribers. Event subscribers take the form of a syseventd loadable module (SLM). syseventd passes the event buffer to each of its subscribers and in return expects a notification as to the successful or unsuccessful delivery attempt. Upon successful delivery of the event buffer to all interested event subscribers, syseventd frees the event buffer from the kernel event queue.
选项	The following option is supported: <code>-d <i>debug_level</i></code> Enable debug mode. Messages are printed to the invoking user's terminal.
退出状态	The following exit values are returned: 0 Successful completion. non-zero An error occurred.
文件	<code>/etc/sysevent/syseventd_daemon.lock</code> daemon lock file <code>/etc/sysevent/sysevent_door</code> kernel to syseventd door file <code>/usr/lib/sysevent/modules</code> SLM directory repository <code>/usr/platform/‘uname -i’/lib/sysevent/modules</code> SLM directory repository <code>/usr/platform/‘uname -m’/lib/sysevent/modules</code> SLM directory repository
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [svcs\(1\)](#), [svcadm\(1M\)](#), [syseventconfd\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注

The syseventd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/system/sysevent:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	syslogd – log system messages
用法概要	<pre>/usr/sbin/syslogd [-d] [-f configfile] [-m markinterval] [-p path] [-t -T]</pre>
描述	syslogd reads and forwards system messages to the appropriate log files or users, depending upon the priority of a message and the system facility from which it originates. The configuration file <code>/etc/syslog.conf</code> (see syslog.conf(4)) controls where messages are forwarded. syslogd logs a mark (timestamp) message every <i>markinterval</i> minutes (default 20) at priority LOG_INFO to the facility whose name is given as mark in the <code>syslog.conf</code> file. A system message consists of a single line of text, which may be prefixed with a priority code number enclosed in angle-brackets (< >); priorities are defined in <code><sys/syslog.h></code>. syslogd reads from the STREAMS log driver, <code>/dev/log</code>, and from any transport provider specified in <code>/etc/netconfig</code>, <code>/etc/net/transport/hosts</code>, and <code>/etc/net/transport/services</code>. syslogd reads the configuration file when it starts up, and again whenever it receives a HUP signal (see signal.h(3HEAD)), at which time it also closes all files it has open, re-reads its configuration file, and then opens only the log files that are listed in that file. syslogd exits when it receives a TERM signal. As it starts up, syslogd creates the file <code>/var/run/syslog.pid</code>, if possible, containing its process identifier (PID). If message ID generation is enabled (see log(7D)), each message will be preceded by an identifier in the following format: <code>[ID msgid facility . priority]</code>. <i>msgid</i> is the message's numeric identifier described in msgid(1M). <i>facility</i> and <i>priority</i> are described in syslog.conf(4). <code>[ID 123456 kern.notice]</code> is an example of an identifier when message ID generation is enabled. If the message originated in a loadable kernel module or driver, the kernel module's name (for example, <code>ufs</code>) will be displayed instead of <code>unix</code>. See EXAMPLES for sample output from syslogd with and without message ID generation enabled. In an effort to reduce visual clutter, message IDs are not displayed when writing to the console; message IDs are only written to the log file. See 示例. The <code>/etc/default/syslogd</code> file contains the following default parameter settings, which are in effect if neither the <code>-t</code> nor <code>-T</code> option is selected. See FILES. The recommended way to allow or disallow message logging is through the use of the service management facility (smf(5)) property: <pre>svc:/system/system-log/config/log_from_remote</pre> This property specifies whether remote messages are logged. <code>log_from_remote=true</code> is equivalent to the <code>-t</code> command-line option and <code>false</code> is equivalent to the <code>-T</code> command-line option. The default value for <code>-log_from_remote</code> is <code>false</code>. See NOTES, below.

LOG_FROM_REMOTE

Specifies whether remote messages are logged. LOG_FROM_REMOTE=NO is equivalent to the `-t` command-line option. The default value for LOG_FROM_REMOTE is YES.

选项

The following options are supported:

- d**
Turn on debugging. This option should only be used interactively in a root shell once the system is in multi-user mode. It should *not* be used in the system start-up scripts, as this will cause the system to hang at the point where `syslogd` is started.
- f *configfile***
Specify an alternate configuration file.
- m *markinterval***
Specify an interval, in minutes, between mark messages.
- p *path***
Specify an alternative log device name. The default is `/dev/log`.
- T**
Enable the `syslogd` UDP port to turn on logging of remote messages. This is the default behavior. See [文件](#).
- t**
Disable the `syslogd` UDP port to turn off logging of remote messages. See [文件](#).

示例**示例 1** syslogd Output Without Message ID Generation Enabled

The following example shows the output from `syslogd` when message ID generation *is not* enabled:

```
Sep 29 21:41:18 cathy unix: alloc /: file system full
```

示例 2 syslogd Output with ID generation Enabled

The following example shows the output from `syslogd` when message ID generation *is* enabled. The message ID is displayed when writing to log file `/var/adm/messages`.

```
Sep 29 21:41:18 cathy ufs: [ID 845546 kern.notice]
                        alloc /: file system full
```

示例 3 syslogd Output with ID Generation Enabled

The following example shows the output from `syslogd` when message ID generation *is* enabled when writing to the console. Even though message ID is enabled, the message ID is not displayed at the console.

```
Sep 29 21:41:18 cathy ufs: alloc /: file system full
```

示例 4 Enabling Acceptance of UDP Messages from Remote Systems

The following commands enable `syslogd` to accept entries from remote systems.

```
# svccfg -s svc:/system/system-log setprop config/log_from_remote = true
# svcadm restart svc:/system/system-log
```

文件

- `/etc/syslog.conf`
Configuration file
- `/var/run/syslog.pid`
Process ID
- `/etc/default/syslogd`
Contains default settings. You can override some of the settings by command-line options.
- `/dev/log`
STREAMS log driver
- `/etc/netconfig`
Transport providers available on the system
- `/etc/net/transport/hosts`
Network hosts for each transport
- `/etc/net/transport/services`
Network services for each transport

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[logger\(1\)](#), [svcs\(1\)](#), [msgid\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [syslog\(3C\)](#), [syslog.conf\(4\)](#), [attributes\(5\)](#), [signal.h\(3HEAD\)](#), [smf\(5\)](#), [log\(7D\)](#)

附注

The mark message is a system time stamp, and so it is only defined for the system on which `syslogd` is running. It can not be forwarded to other systems.

When `syslogd` receives a HUP signal, it attempts to complete outputting pending messages, and close all log files to which it is currently logging messages. If, for some reason, one (or more) of these files does not close within a generous grace period, `syslogd` discards the pending messages, forcibly closes these files, and starts reconfiguration. If this shutdown procedure is disturbed by an unexpected error and `syslogd` cannot complete reconfiguration, `syslogd` sends a mail message to the superuser on the current system stating that it has shut down, and exits.

Care should be taken to ensure that each window displaying messages forwarded by `syslogd` (especially console windows) is run in the system default locale (which is `syslogd`'s locale). If

this advice is not followed, it is possible for a sys log message to alter the terminal settings for that window, possibly even allowing remote execution of arbitrary commands from that window.

The syslogd service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/system/system-log:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

When syslogd is started by means of [svcadm\(1M\)](#), if a value is specified for LOG_FROM_REMOTE in the /etc/defaults/syslogd file, the SMF property

svc:/system/system-log/config/log_from_remote is set to correspond to the LOG_FROM_REMOTE value and the /etc/default/syslogd file is modified to replace the LOG_FROM_REMOTE specification with the following comment:

```
# LOG_FROM_REMOTE is now set using svccfg(1m), see syslogd(1m).
```

If neither LOG_FROM_REMOTE nor svc:/system/system-log/config/log_from_remote are defined, the default is to log remote messages.

On installation, the initial value of svc:/system/system-log/config/log_from_remote is false.

引用名	tapes – creates /dev entries for tape drives attached to the system
用法概要	<code>/usr/sbin/tapes [-r root_dir]</code>
描述	devfsadm(1M) is now the preferred command for /dev and /devices and should be used instead of tapes. tapes creates symbolic links in the /dev/rmt directory to the actual tape device special files under the /devices directory tree. tapes searches the kernel device tree to see what tape devices are attached to the system. For each equipped tape drive, the following steps are performed: 1. The /dev/rmt directory is searched for a /dev/rmt/<i>n</i> entry that is a symbolic link to the /devices special node of the current tape drive. If one is found, this determines the logical controller number of the tape drive. 2. The rest of the special devices associated with the drive are checked, and incorrect symbolic links are removed and necessary ones added. 3. If none are found, a new logical controller number is assigned (the lowest-unused number), and new symbolic links are created for all the special devices associated with the drive. tapes does not remove links to non-existent devices; these must be removed by hand. tapes is run each time a reconfiguration-boot is performed, or when add_drv(1M) is executed.
Notice to Driver Writers	tapes(1M) considers all devices with the node type DDI_NT_TAPE to be tape devices; these devices must have their minor name created with a specific format. The minor name encodes operational modes for the tape device and consists of an ASCII string of the form [<i>l,m,h,c,u</i>][<i>b</i>][<i>n</i>]. The first character set is used to specify the tape density of the device, and are named low (<i>l</i>), medium (<i>m</i>), high (<i>h</i>), compressed (<i>c</i>), and ultra (<i>u</i>). These specifiers only express a relative density; it is up to the driver to assign specific meanings as needed. For example, 9 track tape devices interpret these as actual bits-per-inch densities, where <i>l</i> means 800 BPI, <i>m</i> means 1600 BPI, and <i>h</i> means 6250 BPI, whereas 4mm DAT tapes defines <i>l</i> as standard format, and <i>m</i>, <i>h</i>, <i>c</i> and <i>u</i> as compressed format. Drivers may choose to implement any or all of these format types. During normal tape operation (non-BSD behavior), once an EOF mark has been reached, subsequent reads from the tape device return an error. An explicit IOCTL must be issued to space over the EOF mark before the next file can be read. <i>b</i> instructs the device to observe BSD behavior, where reading at EOF will cause the tape device to automatically space over the EOF mark and begin reading from the next file. <i>n</i> or no-rewind-on-close instructs the driver to not rewind to the beginning of tape when the device is closed. Normal behavior for tape devices is to reposition to BOT when closing. See mtio(7I).

The minor number for tape devices should be created by encoding the device's instance number using the tape macro `MTMINOR` and ORing in the proper combination of density, BSD behavior, and no-rewind flags. See [mtio\(7I\)](#).

To prevent tapes from attempting to automatically generate links for a device, drivers must specify a private node type and refrain from using the node type string `DDI_NT_TAPE` when calling [ddi_create_minor_node\(9F\)](#).

选项

The following options are supported:

`-r root_dir` Causes tapes to presume that the `/dev/rmt` directory tree is found under `root_dir`, not directly under `/`.

错误

If tapes finds entries of a particular logical controller linked to different physical controllers, it prints an error message and exits without making any changes to the `/dev` directory, since it cannot determine which of the two alternative logical to physical mappings is correct. The links should be manually corrected or removed before another reconfiguration boot is performed.

示例

示例 1 Creating Tape Device Nodes From Within the Driver's `attach()` Function

This example demonstrates creating tape device nodes from within the `xktape` driver's [attach\(9E\)](#) function.

```
#include <sys/mtio.h>
struct tape_minor_info {
    char *minor_name;
    int  minor_mode;
};
/*
 * create all combinations of logical tapes
 */
static struct tape_minor_info example_tape[] = {
    {"", 0}, /* default tape */
    {"l", MT_DENSITY1},
    {"lb", MT_DENSITY1 | MT_BSD},
    {"lbn", MT_DENSITY1 | MT_BSD | MT_NOREWIND},
    {"m", MT_DENSITY2},
    {"mb", MT_DENSITY2 | MT_BSD},
    {"mbn", MT_DENSITY2 | MT_BSD | MT_NOREWIND},
    {"h", MT_DENSITY3},
    {"hb", MT_DENSITY3 | MT_BSD},
    {"hbn", MT_DENSITY3 | MT_BSD | MT_NOREWIND},
    {"c", MT_DENSITY4},
    {"cb", MT_DENSITY4 | MT_BSD},
    {"cbn", MT_DENSITY4 | MT_BSD | MT_NOREWIND},
    {NULL, 0},
};
```

示例1 Creating Tape Device Nodes From Within the Driver's attach() Function (续)

```

int
xktapeattach(dev_info_t *dip, ddi_attach_cmd_t cmd)
{
    int instance;
    struct tape_minor_info *mdp;
    /* other stuff in attach... */
    instance = ddi_get_instance(dip);

    for (mdp = example_tape; mdp->minor_name != NULL; mdp++) {
        ddi_create_minor_node(dip, mdp->minor_name, S_IFCHR,
            (MTMINOR(instance) | mdp->minor_mode), DDI_NT_TAPE, 0);
    }
}

```

Installing the xktape driver on a Sun Fire 4800, with the driver controlling a SCSI tape (target 4 attached to an [isp\(7D\)](#) SCSI HBA) and performing a reconfiguration-boot creates the following special files in /devices.

```

# ls -l /devices/ssm@0,0/pci@18,700000/pci@1/SUNW,isp@two@4
crw-rw-rw-  1 root sys   33,136 Aug 29 00:02 xktape@4,0:
crw-rw-rw-  1 root sys   33,200 Aug 29 00:02 xktape@4,0:b
crw-rw-rw-  1 root sys   33,204 Aug 29 00:02 xktape@4,0:bn
crw-rw-rw-  1 root sys   33,152 Aug 29 00:02 xktape@4,0:c
crw-rw-rw-  1 root sys   33,216 Aug 29 00:02 xktape@4,0:cb
crw-rw-rw-  1 root sys   33,220 Aug 29 00:02 xktape@4,0:cbn
crw-rw-rw-  1 root sys   33,156 Aug 29 00:02 xktape@4,0:cn
crw-rw-rw-  1 root sys   33,144 Aug 29 00:02 xktape@4,0:h
crw-rw-rw-  1 root sys   33,208 Aug 29 00:02 xktape@4,0:hb
crw-rw-rw-  1 root sys   33,212 Aug 29 00:02 xktape@4,0:hbn
crw-rw-rw-  1 root sys   33,148 Aug 29 00:02 xktape@4,0:hn
crw-rw-rw-  1 root sys   33,128 Aug 29 00:02 xktape@4,0:l
crw-rw-rw-  1 root sys   33,192 Aug 29 00:02 xktape@4,0:lb
crw-rw-rw-  1 root sys   33,196 Aug 29 00:02 xktape@4,0:lbh
crw-rw-rw-  1 root sys   33,132 Aug 29 00:02 xktape@4,0:ln
crw-rw-rw-  1 root sys   33,136 Aug 29 00:02 xktape@4,0:m
crw-rw-rw-  1 root sys   33,200 Aug 29 00:02 xktape@4,0:mb
crw-rw-rw-  1 root sys   33,204 Aug 29 00:02 xktape@4,0:mbn
crw-rw-rw-  1 root sys   33,140 Aug 29 00:02 xktape@4,0:mn
crw-rw-rw-  1 root sys   33,140 Aug 29 00:02 xktape@4,0:n

```

/dev/rmt will contain the logical tape devices (symbolic links to tape devices in /devices).

```

# ls -l /dev/rmt
/dev/rmt/0 -> ../../devices/[...]/xktape@4,0:
/dev/rmt/0b -> ../../devices/[...]/xktape@4,0:b
/dev/rmt/0bn -> ../../devices/[...]/xktape@4,0:bn

```

示例 1 Creating Tape Device Nodes From Within the Driver's attach() Function (续)

```
/dev/rmt/0c -> ../../devices/[...]/xktape@4,0:c
/dev/rmt/0cb -> ../../devices/[...]/xktape@4,0:cb
/dev/rmt/0cbn -> ../../devices/[...]/xktape@4,0:cbn
/dev/rmt/0cn -> ../../devices/[...]/xktape@4,0:cn
/dev/rmt/0h -> ../../devices/[...]/xktape@4,0:h
/dev/rmt/0hb -> ../../devices/[...]/xktape@4,0:hb
/dev/rmt/0hbn -> ../../devices/[...]/xktape@4,0:hbn
/dev/rmt/0hn -> ../../devices/[...]/xktape@4,0:hn
/dev/rmt/0l -> ../../devices/[...]/xktape@4,0:l
/dev/rmt/0lb -> ../../devices/[...]/xktape@4,0:lb
/dev/rmt/0lbn -> ../../devices/[...]/xktape@4,0:lbn
/dev/rmt/0ln -> ../../devices/[...]/xktape@4,0:ln
/dev/rmt/0m -> ../../devices/[...]/xktape@4,0:m
/dev/rmt/0mb -> ../../devices/[...]/xktape@4,0:mb
/dev/rmt/0mbn -> ../../devices/[...]/xktape@4,0:mbn
/dev/rmt/0mn -> ../../devices/[...]/xktape@4,0:mn
/dev/rmt/0n -> ../../devices/[...]/xktape@4,0:n
```

文件 /dev/rmt/* logical tape devices
/devices/* tape device nodes

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [add_drv\(1M\)](#), [devfsadm\(1M\)](#), [attributes\(5\)](#), [isp\(7D\)](#), [devfs\(7FS\)](#), [mtio\(7I\)](#), [attach\(9E\)](#), [ddi_create_minor_node\(9F\)](#)

《Writing Device Drivers》

已知问题 tapes silently ignores malformed minor device names.

引用名	tcxconfig – configure default linearity of 24-bit TrueColor Visual for OpenWindows				
用法概要	<code>/usr/sbin/tcxconfig [linear nonlinear]</code>				
描述	The <code>tcxconfig</code> script changes the default linearity of a 24-bit TrueColor Visual for OpenWindows on a system with an S24 frame buffer. When the S24 graphics driver for OpenWindows is installed, the default 24-bit TrueColor Visual is nonlinear. You can run <code>tcxconfig</code> with an argument that specifies the setting you want. OpenWindows should not be running when you execute the <code>tcxconfig</code> script with an option. Start OpenWindows after <code>tcxconfig</code> has set the linearity you desire.				
选项	If you specify no option, <code>tcxconfig</code> displays the current default setting. You must become superuser before you can execute <code>tcxconfig</code> with one of the following options. <table><tr><td><code>linear</code></td><td>Set linear visual to be the default 24-bit TrueColor Visual. This means colors will be gamma-corrected.</td></tr><tr><td><code>nonlinear</code></td><td>Set nonlinear visual to be the default 24-bit TrueColor Visual.</td></tr></table>	<code>linear</code>	Set linear visual to be the default 24-bit TrueColor Visual. This means colors will be gamma-corrected.	<code>nonlinear</code>	Set nonlinear visual to be the default 24-bit TrueColor Visual.
<code>linear</code>	Set linear visual to be the default 24-bit TrueColor Visual. This means colors will be gamma-corrected.				
<code>nonlinear</code>	Set nonlinear visual to be the default 24-bit TrueColor Visual.				
退出状态	The following exit values are returned: <table><tr><td>0</td><td>success</td></tr><tr><td>1</td><td>an error has occurred.</td></tr></table>	0	success	1	an error has occurred.
0	success				
1	an error has occurred.				
属性	See attributes(5) for descriptions of the following attributes:				
另请参见	attributes(5)				

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWtcxow

引用名

th_define – create fault injection test harness error specifications

用法概要

```
th_define [-n name -i instance| -P path] [-a acc_types]
        [-r reg_number] [-l offset [length]]
        [-c count [failcount]] [-o operator [operand]]
        [-f acc_chk] [-w max_wait_period [report_interval]]

or

th_define [-n name -i instance| -P path]
        [-a log [acc_types] [-r reg_number] [-l offset [length]]]
        [-c count [failcount]] [-s collect_time] [-p policy]
        [-x flags] [-C comment_string]
        [-e fixup_script [args]]

or

th_define [-h]
```

描述

The `th_define` utility provides an interface to the `bus_ops` fault injection `bofi` device driver for defining error injection specifications (referred to as `errdefs`). An `errdef` corresponds to a specification of how to corrupt a device driver's accesses to its hardware. The command line arguments determine the precise nature of the fault to be injected. If the supplied arguments define a consistent `errdef`, the `th_define` process will store the `errdef` with the `bofi` driver and suspend itself until the criteria given by the `errdef` become satisfied (in practice, this will occur when the access counts go to zero).

You use the `th_manage(1M)` command with the `start` option to activate the resulting `errdef`. The effect of `th_manage` with the `start` option is that the `bofi` driver acts upon the `errdef` by matching the number of hardware accesses—specified in *count*, that are of the type specified in *acc_types*, made by instance number *instance*—of the driver whose name is *name*, (or by the driver instance specified by *path*) to the register set (or DMA handle) specified by *reg_number*, that lie within the range *offset* to *offset* + *length* from the beginning of the register set or DMA handle. It then applies *operator* and *operand* to the next *failcount* matching accesses.

If *acc_types* includes `log`, `th_define` runs in automatic test script generation mode, and a set of test scripts (written in the Korn shell) is created and placed in a sub-directory of the current directory with the name `<driver>.test.<id>` (for example, `glm.test.978177106`). A separate, executable script is generated for each access handle that matches the logging criteria. The log of accesses is placed at the top of each script as a record of the session. If the current directory is not writable, file output is written to standard output. The base name of each test file is the driver name, and the extension is a number that discriminates between different access handles. A control script (with the same name as the created test directory) is generated that will run all the test scripts sequentially.

Executing the scripts will install, and then activate, the resulting error definitions. Error definitions are activated sequentially and the driver instance under test is taken offline and brought back online before each test (refer to the `-e` option for more information). By default, logging applies to all PIO accesses, all interrupts, and all DMA accesses to and from areas

mapped for both reading and writing. You can constrain logging by specifying additional *acc_types*, *reg_number*, *offset* and *length*. Logging will continue for *count* matching accesses, with an optional time limit of *collect_time* seconds.

Either the *-n* or *-P* option must be provided. The other options are optional. If an option (other than *-a*) is specified multiple times, only the final value for the option is used. If an option is not specified, its associated value is set to an appropriate default, which will provide maximal error coverage as described below.

选项

The following options are available:

<i>-n name</i>	Specify the name of the driver to test. (String)
<i>-i instance</i>	Test only the specified driver instance (<i>-1</i> matches all instances of driver). (Numeric)
<i>-P path</i>	Specify the full device path of the driver to test. (String)
<i>-r reg_number</i>	Test only the given register set or DMA handle (<i>-1</i> matches all register sets and DMA handles). (Numeric)
<i>-a acc_types</i>	Only the specified access types will be matched. Valid values for the <i>acc_types</i> argument are <i>log</i> , <i>pio</i> , <i>pio_r</i> , <i>pio_w</i> , <i>dma</i> , <i>dma_r</i> , <i>dma_w</i> and <i>intr</i> . Multiple access types, separated by spaces, can be specified. The default is to match all hardware accesses. If <i>acc_types</i> is set to <i>log</i> , logging will match all PIO accesses, interrupts and DMA accesses to and from areas mapped for both reading and writing. <i>log</i> can be combined with other <i>acc_types</i> , in which case the matching condition for logging will be restricted to the specified additional <i>acc_types</i> . Note that <i>dma_r</i> will match only DMA handles mapped for reading only; <i>dma_w</i> will match only DMA handles mapped for writing only; <i>dma</i> will match only DMA handles mapped for both reading and writing.
<i>-l offset [length]</i>	Constrain the range of qualifying accesses. The <i>offset</i> and <i>length</i> arguments indicate that any access of the type specified with the <i>-a</i> option, to the register set or DMA handle specified with the <i>-r</i> option, lie at least <i>offset</i> bytes into the register set or DMA handle and at most <i>offset + length</i>

-c *count* [*failcount*]

bytes into it. The default for *offset* is 0. The default for *length* is the maximum value that can be placed in an `offset_t` C data type (see `types.h`). Negative values are converted into unsigned quantities. Thus, `th_define -l 0 -1` is maximal.

Wait for *count* number of matching accesses, then apply an operator and operand (see the `-o` option) to the next *failcount* number of matching accesses. If the access type (see the `-a` option) includes logging, the number of logged accesses is given by *count* + *failcount* - 1. The -1 is required because the last access coincides with the first faulting access.

Note that access logging may be combined with error injection if *failcount* and *operator* are nonzero and if the access type includes logging and any of the other access types (`pio`, `dma` and `intr`) See the description of access types in the definition of the `-a` option, above.

When the *count* and *failcount* fields reach zero, the status of the `errdef` is reported to standard output. When all active `errdefs` created by the `th_define` process complete, the process exits. If *acc_types* includes `log`, *count* determines how many accesses to log. If *count* is not specified, a default value is used. If *failcount* is set in this mode, it will simply increase the number of accesses logged by a further *failcount* - 1.

-o *operator* [*operand*]

For qualifying PIO read and write accesses, the value read from or written to the hardware is corrupted according to the value of *operator*:

- EQ *operand* is returned to the driver.
- OR *operand* is bitwise ORed with the real value.
- AND *operand* is bitwise ANDed with the real value.
- XOR *operand* is bitwise XORed with the real value.

For PIO write accesses, the following operator is allowed:

NO Simply ignore the driver's attempt to write to the hardware.

Note that a driver performs PIO via the `ddi_getX()`, `ddi_putX()`, `ddi_rep_getX()` and `ddi_rep_putX()` routines (where *X* is 8, 16, 32 or 64). Accesses made using `ddi_getX()` and `ddi_putX()` are treated as a single access, whereas an access made using the `ddi_rep_*(9F)` routines are broken down into their respective number of accesses, as given by the *repcount* parameter to these DDI calls. If the access is performed via a DMA handle, *operator* and *value* are applied to every access that comprises the DMA request. If interference with interrupts has been requested then the operator may take any of the following values:

DELAY	After <i>count</i> accesses (see the <code>-c</code> option), delay delivery of the next <i>failcount</i> number of interrupts for <i>operand</i> number of microseconds.
LOSE	After <i>count</i> number of interrupts, fail to deliver the next <i>failcount</i> number of real interrupts to the driver.
EXTRA	After <i>count</i> number of interrupts, start delivering <i>operand</i> number of extra interrupts for the next <i>failcount</i> number of real interrupts.

The default value for *operand* and *operator* is to corrupt the data access by flipping each bit (XOR with -1).

`-f acc_chk`

If the *acc_chk* parameter is set to 1 or `pio`, then the driver's calls to `ddi_check_acc_handle(9F)` return `DDI_FAILURE` when the access count goes to 1. If the *acc_chk* parameter is set to 2 or `dma`, then the driver's calls to `ddi_check_dma_handle(9F)` return `DDI_FAILURE` when the access count goes to 1.

<code>-w max_wait_period [report_interval]</code>	Constrain the period for which an error definition will remain active. The option applies only to non-logging errdefs. If an error definition remains active for <i>max_wait_period</i> seconds, the test will be aborted. If <i>report_interval</i> is set to a nonzero value, the current status of the error definition is reported to standard output every <i>report_interval</i> seconds. The default value is zero. The status of the errdef is reported in parsable format (eight fields, each separated by a colon (:) character, the last of which is a string enclosed by double quotes and the remaining seven fields are integers): <i>ft:mt:ac:fc:chk:ec:s:"message"</i> which are defined as follows: <table><tr><td><i>ft</i></td><td>The UTC time when the fault was injected.</td></tr><tr><td><i>mt</i></td><td>The UTC time when the driver reported the fault.</td></tr><tr><td><i>ac</i></td><td>The number of remaining non-faulting accesses.</td></tr><tr><td><i>fc</i></td><td>The number of remaining faulting accesses.</td></tr><tr><td><i>chk</i></td><td>The value of the <i>acc_chk</i> field of the errdef.</td></tr><tr><td><i>ec</i></td><td>The number of fault reports issued by the driver against this errdef (<i>mt</i> holds the time of the initial report).</td></tr><tr><td><i>s</i></td><td>The severity level reported by the driver.</td></tr><tr><td><i>"message"</i></td><td>Textual reason why the driver has reported a fault.</td></tr></table>	<i>ft</i>	The UTC time when the fault was injected.	<i>mt</i>	The UTC time when the driver reported the fault.	<i>ac</i>	The number of remaining non-faulting accesses.	<i>fc</i>	The number of remaining faulting accesses.	<i>chk</i>	The value of the <i>acc_chk</i> field of the errdef.	<i>ec</i>	The number of fault reports issued by the driver against this errdef (<i>mt</i> holds the time of the initial report).	<i>s</i>	The severity level reported by the driver.	<i>"message"</i>	Textual reason why the driver has reported a fault.
<i>ft</i>	The UTC time when the fault was injected.																
<i>mt</i>	The UTC time when the driver reported the fault.																
<i>ac</i>	The number of remaining non-faulting accesses.																
<i>fc</i>	The number of remaining faulting accesses.																
<i>chk</i>	The value of the <i>acc_chk</i> field of the errdef.																
<i>ec</i>	The number of fault reports issued by the driver against this errdef (<i>mt</i> holds the time of the initial report).																
<i>s</i>	The severity level reported by the driver.																
<i>"message"</i>	Textual reason why the driver has reported a fault.																
<code>-h</code>	Display the command usage string.																
<code>-s collect_time</code>	If <i>acc_types</i> is given with the <i>-a</i> option and includes <i>log</i> , the errdef will log accesses for <i>collect_time</i> seconds (the default is to log until the log becomes full). Note that, if the errdef																

-p *policy*

specification matches multiple driver handles, multiple logging errdefs are registered with the bofi driver and logging terminates when all logs become full or when *collect_time* expires or when the associated errdefs are cleared. The current state of the log can be checked with the [th_manage\(1M\)](#) command, using the broadcast parameter. A log can be terminated by running [th_manage\(1M\)](#) with the `clear_errdefs` option or by sending a SIGALRM signal to the `th_define` process. See [alarm\(2\)](#) for the semantics of SIGALRM.

Applicable when the *acc_types* option includes `log`. The parameter modifies the policy used for converting from logged accesses to errdefs. All policies are inclusive:

- Use `rare` to bias error definitions toward rare accesses (default).
- Use `operator` to produce a separate error definition for each operator type (default).
- Use `common` to bias error definitions toward common accesses.
- Use `median` to bias error definitions toward median accesses.
- Use `maximal` to produce multiple error definitions for duplicate accesses.
- Use `unbiased` to create unbiased error definitions.
- Use `onebyte`, `twobyte`, `fourbyte`, or `eightbyte` to select errdefs corresponding to 1, 2, 4 or 8 byte accesses (if chosen, the `-xr` option is enforced in order to ensure that `ddi_rep_*()` calls are decomposed into *multiple single accesses*).
- Use `multibyte` to create error definitions for multibyte accesses performed using `ddi_rep_get*()` and `ddi_rep_put*()`.

Policies can be combined by adding together these options. See the NOTES section for further information.

`-x flags`

Applicable when the `acc_types` option includes `log`. The `flags` parameter modifies the way in which the `bofi` driver logs accesses. It is specified as a string containing any combination of the following letters:

- w Continuous logging (that is, the log will wrap when full).
- t Timestamp each log entry (access times are in seconds).
- r Log repeated I/O as individual accesses (for example, a `ddi_rep_get16(9F)` call which has a `repcount` of `N` is logged `N` times with each transaction logged as size 2 bytes. Without this option, the default logging behavior is to log this access once only, with a transaction size of twice the `repcount`).

`-C comment_string`

Applicable when the `acc_types` option includes `log`. It provides a comment string to be placed in any generated test scripts. The string must be enclosed in double quotes.

`-e fixup_script [args]`

Applicable when the `acc_types` option includes `log`. The output of a logging `errdefs` is to generate a test script for each driver access handle. Use this option to embed a command in the resulting script before the errors are injected. The generated test scripts will take an instance offline and bring it back online before injecting errors in order to bring the instance into a known fault-free state. The executable `fixup_script` will be called twice with the set of optional `args`— once just before the instance is taken offline and again after the instance has been brought online. The following variables are passed into the environment of the called executable:

<code>DRIVER_PATH</code>	Identifies the device path of the instance.
<code>DRIVER_INSTANCE</code>	Identifies the instance number of the device.

DRIVER_UNCONFIGURE	Has the value 1 when the instance is about to be taken offline.
DRIVER_CONFIGURE	Has the value 1 when the instance has just been brought online.

Typically, the executable ensures that the device under test is in a suitable state to be taken offline (unconfigured) or in a suitable state for error injection (for example configured, error free and servicing a workload). A minimal script for a network driver could be:

```
#!/bin/ksh

driver=xyznetdriver
ifnum=$driver$DRIVER_INSTANCE

if [[ $DRIVER_CONFIGURE = 1 ]]; then
    ifconfig $ifnum plumb
    ifconfig $ifnum ...
    ifworkload start $ifnum
elif [[ $DRIVER_UNCONFIGURE = 1 ]]; then
    ifworkload stop $ifnum
    ifconfig $ifnum down
    ifconfig $ifnum unplumb
fi
exit $?
```

The -e option must be the last option on the command line.

If the -a log option is selected but the -e option is not given, a default script is used. This script repeatedly attempts to detach and then re-attach the device instance under test.

示例

Examples of Error Definitions

```
th_define -n foo -i 1 -a log
```

Logs all accesses to all handles used by instance 1 of the foo driver while running the default workload (attaching and detaching the instance). Then generates a set of test scripts to inject appropriate errdefs while running that default workload.

```
th_define -n foo -i 1 -a log pio
```

Logs PIO accesses to each PIO handle used by instance 1 of the foo driver while running the default workload (attaching and detaching the instance). Then generates a set of test scripts to inject appropriate errdefs while running that default workload.

```
th_define -n foo -i 1 -p onebyte median -e fixup arg -now
```

Logs all accesses to all handles used by instance 1 of the foo driver while running the workload defined in the fixup script fixup with arguments arg and -now. Then generates a set of test scripts to inject appropriate errdefs while running that workload. The resulting error definitions are requested to focus upon single byte accesses to locations that are accessed a *median* number of times with respect to frequency of access to I/O addresses.

```
th_define -n se -l 0x20 1 -a pio_r -o OR 0x4 -c 10 1000
```

Simulates a stuck serial chip command by forcing 1000 consecutive read accesses made by any instance of the se driver to its command status register, thereby returning status busy.

```
th_define -n foo -i 3 -r 1 -a pio_r -c 0 1 -f 1 -o OR 0x100
```

Causes 0x100 to be ORed into the next physical I/O read access from any register in register set 1 of instance 3 of the foo driver. Subsequent calls in the driver to ddi_check_acc_handle() return DDI_FAILURE.

```
th_define -n foo -i 3 -r 1 -a pio_r -c 0 1 -o OR 0x0
```

Causes 0x0 to be ORed into the next physical I/O read access from any register in register set 1 of instance 3 of the foo driver. This is of course a no-op.

```
th_define -n foo -i 3 -r 1 -l 0x8100 1 -a pio_r -c 0 10 -o EQ 0x70003
```

Causes the next ten next physical I/O reads from the register at offset 0x8100 in register set 1 of instance 3 of the foo driver to return 0x70003.

```
th_define -n foo -i 3 -r 1 -l 0x8100 1 -a pio_w -c 100 3 -o AND  
0xffffffffffffefff
```

The next 100 physical I/O writes to the register at offset 0x8100 in register set 1 of instance 3 of the foo driver take place as normal. However, on each of the three subsequent accesses, the 0x1000 bit will be cleared.

```
th_define -n foo -i 3 -r 1 -l 0x8100 0x10 -a pio_r -c 0 1 -f 1 -o XOR 7
```

Causes the bottom three bits to have their values toggled for the next physical I/O read access to registers with offsets in the range 0x8100 to 0x8110 in register set 1 of instance 3 of the foo driver. Subsequent calls in the driver to ddi_check_acc_handle() return DDI_FAILURE.

```
th_define -n foo -i 3 -a pio_w -c 0 1 -o NO 0
```

Prevents the next physical I/O write access to any register in any register set of instance 3 of the foo driver from going out on the bus.

```
th_define -n foo -i 3 -l 0 8192 -a dma_r -c 0 1 -o OR 7
```

Causes 0x7 to be ORed into each long long in the first 8192 bytes of the next DMA read, using any DMA handle for instance 3 of the foo driver.

```
th_define -n foo -i 3 -r 2 -l 0 8 -a dma_r -c 0 1 -o OR 0x7070707070707070
```

Causes 0x70 to be ORed into each byte of the first long long of the next DMA read, using the DMA handle with sequential allocation number 2 for instance 3 of the foo driver.

```
th_define -n foo -i 3 -l 256 256 -a dma_w -c 0 1 -f 2 -o OR 7
```

Causes 0x7 to be ORed into each long long in the range from offset 256 to offset 512 of the next DMA write, using any DMA handle for instance 3 of the foo driver. Subsequent calls in the driver to `ddi_check_dma_handle()` return `DDI_FAILURE`.

```
th_define -n foo -i 3 -r 0 -l 0 8 -a dma_w -c 100 3 -o AND 0xffffffffffffefff
```

The next 100 DMA writes using the DMA handle with sequential allocation number 0 for instance 3 of the foo driver take place as normal. However, on each of the three subsequent accesses, the 0x1000 bit will be cleared in the first long long of the transfer.

```
th_define -n foo -i 3 -a intr -c 0 6 -o LOSE 0
```

Causes the next six interrupts for instance 3 of the foo driver to be lost.

```
th_define -n foo -i 3 -a intr -c 30 1 -o EXTRA 10
```

When the thirty-first subsequent interrupt for instance 3 of the foo driver occurs, a further ten interrupts are also generated.

```
th_define -n foo -i 3 -a intr -c 0 1 -o DELAY 1024
```

Causes the next interrupt for instance 3 of the foo driver to be delayed by 1024 microseconds.

附注

The policy option in the `th_define -p` syntax determines how a set of logged accesses will be converted into the set of error definitions. Each logged access will be matched against the chosen policies to determine whether an error definition should be created based on the access.

Any number of policy options can be combined to modify the generated error definitions.

Bytewise Policies

These select particular I/O transfer sizes. Specifying a byte policy will exclude other byte policies that have not been chosen. If none of the byte type policies is selected, all transfer sizes are treated equally. Otherwise, only those specified transfer sizes will be selected.

	onebyte	Create errdefs for one byte accesses (ddi_get8())
	twobyte	Create errdefs for two byte accesses (ddi_get16())
	fourbyte	Create errdefs for four byte accesses (ddi_get32())
	eightbyte	Create errdefs for eight byte accesses (ddi_get64())
	multibyte	Create errdefs for repeated byte accesses (ddi_rep_get*())
Frequency of Access Policies	The frequency of access to a location is determined according to the access type, location and transfer size (for example, a two-byte read access to address A is considered distinct from a four-byte read access to address A). The algorithm is to count the number of accesses (of a given type and size) to a given location, and find the locations that were most and least accessed (let <i>maxa</i> and <i>mina</i> be the number of times these locations were accessed, and <i>mean</i> the total number of accesses divided by total number of locations that were accessed). Then a rare access is a location that was accessed less than $(mean - mina) / 3 + mina$ times. Similarly for the definition of common accesses: $maxa - (maxa - mean) / 3$ A location whose access patterns lies within these cutoffs is regarded as a location that is accessed with median frequency. rare Create errdefs for locations that are rarely accessed. common Create errdefs for locations that are commonly accessed. median Create errdefs for locations that are accessed a median frequency.	
Policies for Minimizing errdefs	If a transaction is duplicated, either a single or multiple errdefs will be written to the test scripts, depending upon the following two policies: maximal Create multiple errdefs for locations that are repeatedly accessed. unbiased Create a single errdef for locations that are repeatedly accessed. operators For each location, a default operator and operand is typically applied. For maximal test coverage, this default may be modified using the operators policy so that a separate errdef is created for each of the possible corruption operators.	
另请参见	kill(1) , th_manage(1M) , alarm(2) , ddi_check_acc_handle(9F) , ddi_check_dma_handle(9F)	

引用名	th_manage – manage the fault injection test harness														
用法概要	<pre>th_manage name instance command th_manage path command</pre>														
描述	th_manage applies the action specified by <i>command</i> to the instance specified by <i>instance</i> of the driver specified by <i>name</i> (or the driver instance specified by <i>path</i>). The driver instance must be running fault injection specifications (errdefs) defined by th_define(1M). th_manage supports several commands that operate on the driver instance specified by <i>name</i> and <i>instance</i> (or <i>path</i>). The commands are: <table><tr><td>broadcast</td><td>Awaken all th_define processes, causing them to display their current status and exit if the errdef is now defunct (that is, if <i>count</i>, <i>failcount</i>, and <i>acc_chk</i> are all zero).</td></tr><tr><td>clear_acc_chk</td><td>Awaken all th_define processes. If <i>count</i> and <i>failcount</i> are already zero, then set <i>acc_chk</i> to zero, so that th_define exits once it has displayed its status.</td></tr><tr><td>clear_errdefs</td><td>Awaken all th_define processes. <i>count</i>, <i>failcount</i> and <i>acc_chk</i> are all set to zero so that all th_define commands exit once they have displayed their status.</td></tr><tr><td>clear_errors</td><td>Awaken all th_define processes. If <i>count</i> is already zero, set <i>failcount</i> and <i>acc_chk</i> to zero, so that th_define exits once it has displayed its status.</td></tr><tr><td>get_handles</td><td>List all the access handles.</td></tr><tr><td>start</td><td>Begin or resume execution of all errdefs.</td></tr><tr><td>stop</td><td>Suspend all errdefs for this <i>name</i> and <i>instance</i> (or <i>path</i>).</td></tr></table>	broadcast	Awaken all th_define processes, causing them to display their current status and exit if the errdef is now defunct (that is, if <i>count</i> , <i>failcount</i> , and <i>acc_chk</i> are all zero).	clear_acc_chk	Awaken all th_define processes. If <i>count</i> and <i>failcount</i> are already zero, then set <i>acc_chk</i> to zero, so that th_define exits once it has displayed its status.	clear_errdefs	Awaken all th_define processes. <i>count</i> , <i>failcount</i> and <i>acc_chk</i> are all set to zero so that all th_define commands exit once they have displayed their status.	clear_errors	Awaken all th_define processes. If <i>count</i> is already zero, set <i>failcount</i> and <i>acc_chk</i> to zero, so that th_define exits once it has displayed its status.	get_handles	List all the access handles.	start	Begin or resume execution of all errdefs.	stop	Suspend all errdefs for this <i>name</i> and <i>instance</i> (or <i>path</i>).
broadcast	Awaken all th_define processes, causing them to display their current status and exit if the errdef is now defunct (that is, if <i>count</i> , <i>failcount</i> , and <i>acc_chk</i> are all zero).														
clear_acc_chk	Awaken all th_define processes. If <i>count</i> and <i>failcount</i> are already zero, then set <i>acc_chk</i> to zero, so that th_define exits once it has displayed its status.														
clear_errdefs	Awaken all th_define processes. <i>count</i> , <i>failcount</i> and <i>acc_chk</i> are all set to zero so that all th_define commands exit once they have displayed their status.														
clear_errors	Awaken all th_define processes. If <i>count</i> is already zero, set <i>failcount</i> and <i>acc_chk</i> to zero, so that th_define exits once it has displayed its status.														
get_handles	List all the access handles.														
start	Begin or resume execution of all errdefs.														
stop	Suspend all errdefs for this <i>name</i> and <i>instance</i> (or <i>path</i>).														

示例

示例1 Useful Commands

To begin the tests, enter:

```
# th_manage foo 0 start
```

To check the status of the errdefs, enter:

```
# th_manage foo 0 broadcast
```

This causes each th_define process to print out its current status.

If the driver has reported a fatal error, you can take the driver offline using libdevice, clear the error condition by entering:

```
# th_manage foo 0 clear_acc_chk
```

示例1 Useful Commands (续)

or

```
# th_manage foo 0 clear_errors
```

and bring the driver online again using libdevice.

To terminate testing, enter:

```
# th_manage foo 0 clear_errdefs
```

另请参见

[th_define\(1M\)](#)

引用名	tic – terminfo compiler
用法概要	tic [-v [n]] [-c] file
描述	The command <code>tic</code> translates a terminfo file from the source format into the compiled format. The results are placed in the directory <code>/usr/share/lib/terminfo</code>. The compiled format is necessary for use with the library routines in curses(3CURSES). If the environment variable <code>TERMINFO</code> is set, the compiled results are placed there instead of <code>/usr/share/lib/terminfo</code>. Total compiled entries cannot exceed 4096 bytes. The name field cannot exceed 128 bytes. Terminal names exceeding 14 characters will be truncated to 14 characters and a warning message will be printed.
选项	The following options are supported: -c Specifies to check only file for errors. Errors in use= links are not detected. -v[n] Specify that (verbose) output be written to standard error trace information showing tic's progress. The optional integer <i>n</i> is a number from 1 to 10, indicating the desired level of detail of information. If <i>n</i> is omitted, the default level is 1. If <i>n</i> is specified and greater than 1, the level of detail is increased.
操作数	file Contains one or more terminfo terminal descriptions in source format [see terminfo(4)]. Each description in the file describes the capabilities of a particular terminal. When a use=entry-name field is discovered in a terminal entry currently being compiled, tic reads in the binary from <code>/usr/share/lib/terminfo</code> to complete the entry. (Entries created from file will be used first. If the environment variable <code>TERMINFO</code> is set, that directory is searched instead of <code>/usr/share/lib/terminfo</code>.) tic duplicates the capabilities in entry-name for the current entry, with the exception of those capabilities that are explicitly defined in the current entry.
文件	<code>/usr/share/lib/terminfo/??/*</code> Compiled terminal description database
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [captainfo\(1M\)](#), [infocmp\(1M\)](#), [curses\(3CURSES\)](#), [terminfo\(4\)](#), [attributes\(5\)](#)

附注

When an entry, for example, `entry_name_1`, contains a `use=entry_name_2` field, any canceled capabilities in `entry_name_2` must also appear in `entry_name_1` before `use=` for these capabilities to be canceled in `entry_name_1`.

引用名 tncfg – configure trusted networking properties

用法概要

```
tncfg [-t template] [-e|-S [files|ldap]] [subcommand]  
tncfg [-t template] [-e|-S [files|ldap]] -f command_file  
tncfg -z zone [-e] [subcommand]  
tncfg -z zone [-e] -f command_file  
tncfg help
```

描述 The tncfg utility creates, modifies, and displays the configuration of various networking properties related to Trusted Extensions. The command requires that the SMF service, svc:/system/labeld is enabled. It can be executed only in the global zone.

Templates A template is a collection of network security properties that define the rules for labeling packets received from remote hosts. Two host types are supported: cipso and unlabeled. Each template must specify one of these two host types. Hosts that are trusted to specify their own labels are assigned to cipso templates. Otherwise, hosts may be assigned to unlabeled templates, in which a single default label must be specified.

Hosts can be specified using hostnames, IP addresses, or masks. When masks are used, a prefix length that specifies how many bits are required for a match must be appended. Hosts cannot be assigned to more than one template. When masks are used, the entry with the longest matching prefix is used to associate a host with a template. Packets from hosts without a matching template are dropped.

Each template must include an upper and lower bound specifying the accreditation range of accepted labels. Additionally, up to four auxiliary labels can be specified to enumerate labels outside of this range. Services bound to multilevel ports can accept packets from hosts whose labels are within the accreditation range, or match one of the auxiliary labels.

Normally, the template settings and their corresponding hosts are persistently maintained in local files or by means of an LDAP directory, depending on the -S option. These settings are automatically loaded into kernel memory when the user commits the updates. If the -e (ephemeral) option is specified, only the current in-memory properties are displayed and updated. However, the list of hosts associated with an in-memory template is generally incomplete. To view the matching template for a specific host, use the get subcommand.

By default, an unlabeled template, admin_low, is installed with a default label of ADMIN_LOW, and two mask entries matching any IPv4 or IPv6 address, so that the global zone is initially able to contact any unlabeled hosts. It is recommended to remove these two mask entries once your network security policy is established. An additional template, cipso, is installed with no matching hosts. By default all local IP addresses are implicitly associated with this template, but it is recommended that they should be explicitly added to this or a new cipso-type template.

Searches for template and host entries are resolved in the order specified by means of the name service configuration file, `/etc/nsswitch.conf`. The keywords, `tnrhdb` and `tnrhttp`, are used to specify the search order for hosts and templates, respectively. Both the `files` and `ldap` repositories are supported, but it is recommended to specify `files` first.

Creating or modifying a template requires the authorization `solaris.label.network.manage`, which is included in the Object Label Management rights profile.

Zones

Zones are isolated execution environments described in [zones\(5\)](#). Trusted Extensions requires a zone brand called `labeled` to which special properties apply. Each labeled zone must have a unique label property at which it executes processes. This is also the label at which it will accept packets from remote hosts for services bound to single-level ports. Explicit multilevel port can also be specified. Services with the privilege `net_bindmlp` can bind to these ports, and accept packets within the accreditation range or the auxiliary label set associated with the remote host.

Non-global zones must be configured using [zonecfg\(1M\)](#) prior to configuring these properties. In general, updates to each zone's properties, including the global zone, are applied when it is booted. However, the multilevel port properties of running zones are reloaded into kernel memory when updates are committed. If the `-e` (ephemeral) option is specified, the zone must be in the ready or running state, and only its multilevel port properties can be updated.

Creating or modifying a zone's trusted networking properties requires the authorization `solaris.label.zone.manage`, which is included in the Object Label Management rights profile.

Properties

The set of valid properties depends on whether the `-t` or `-z` option was used. The two sets are referred to as the template context and the zone context.

Only a single property value can be specified at a time. Values containing white space must be quoted. An equal sign is required between the property and its value.

The values that can be specified in the template context properties are described below.

`name=template_name`

The initial value for the name is specified using `-t` option using the command line. If the name is changed, the current template properties are applied to the newly named template. In this way an existing template can be cloned for subsequent editing. However, to avoid conflicts, the host entries from the initial template are not copied to the new template. The specified name must not match an existing template.

`host_type=cipso|unlabeled`

When the `unlabeled` host type is used, the value specified using the `def_label` property is implicitly applied to the received packets. The `cipso` host type is used for hosts that are trusted to explicitly label their packets. The default is `unlabeled`.

`def_label=sensitivity_label`

The default label assigned to IP packets that are not explicitly labeled by means of `cipso` or `IPsec`.

`doi=integer`

A positive integer specifying the Domain of Interpretation for the binary representation of the labels. The default is 1.

`min_label=sensitivity_label`

The minimum label in the accreditation range for IP packets that are accepted by multilevel services.

`max_label=sensitivity_label`

The maximum label in the accreditation range for IP packets that are accepted by multilevel services.

`aux_label=sensitivity_label`

Additional labels, outside of the accreditation range, for IP packets that are accepted by multilevel services. Up to four labels may be specified, using the `add` subcommand repetitively.

`host=hostname|IP address[/prefix]`

A hostname or an IP address to which the template properties apply. For IP addresses, both IPv4 and IPv6 formats can be used, followed by an optional slash and prefix length specifying the number of bits to match again IP addresses. The IPv4 address `0.0.0.0` has an implied prefix length of zero, and matches any IPv4 address. Multiple host values can be specified, using the `add` subcommand repetitively. There is no specific limit.

The values that may be specified in the zone context properties are described below.

`name=zone_name`

The name of the zone, which must have previously been configured using [zoncfg\(1M\)](#).

The initial value for the name is specified using `-z` option on the command line. If the name is changed, the current zone properties are applied to the newly named zone. In this way an existing template can be cloned for subsequent editing. However, to avoid conflicts, the initial label value is not copied to the new zone configuration. The specified name must correspond to an existing zone without a trusted networking configuration.

`label=sensitivity_label`

The sensitivity label of the zone. It must be unique for each zone. The global zone value must be `admin_low`.

`visible=yes|no`

Specifies whether the zone responds to ping requests from hosts whose labels don't match the zone's label. The default is `no`.

`mlp_private=port[[-port2]]|tcp|udp`

A single port number, or a range of ports that privileged services can bind to and then accept requests from clients whose labels are with the accreditation range or the set of

auxiliary labels specified in their matching templates. The port specification must be followed by a protocol, either `tcp` or `udp`. This value applies to all interfaces that are private to the zone. Multiple `m1p_private` values can be specified, using the `add` subcommand repetitively. This is only limited by the number of available ports.

`m1p_shared=port[[-port2]/tcp|udp`

A single port number, or a range of ports that privileged services can bind to and then accept requests from clients whose labels are with the accreditation range or the set of auxiliary labels specified in their matching templates. The port specification must be followed by a protocol, either `tcp` or `udp`. This value applies to any `all-zones` interfaces, and must not overlap with the `m1p_shared` ports specifications for other zones. Multiple `m1p_shared` values can be specified, using the `add` subcommand repetitively. This is only limited by the number of available ports.

选项

The following options are supported:

`-e`

Specifies that the data is ephemeral, affecting only what is currently loaded into kernel memory.

`-f command_file`

Specifies the name of `tncfg` command file. *command_file* is a text file of `tncfg` subcommands, one per line.

`-t template`

Specifies the template name. If the named template does not exist, a new template is created. If neither `-t` nor `-z` is specified, the template context is assumed using `cipso` as the default template name.

`-S repository`

The valid repositories are `files` and `ldap`. The repository specifies which name service will be updated. The default repository is `files`.

`-z zone`

Specifies the zone name. The zone must have previously been configured by means of [zonecfg\(1M\)](#).

子命令

Subcommands can be provided on the command line or interactively. Multiple subcommands, separated by semicolons, can be specified on the command line by enclosing the entire set in quotation marks. The lack of subcommands implies an interactive session, during which auto-completion of subcommands can be invoked using the tab key.

The `add`, `clear`, and `remove` subcommands are used for properties that can accept multiple values. However, only one value can be specified at a time.

Subcommands which can result in destructive actions or loss of work have an `-F` option to force the action. If input is from a terminal device, the user is prompted when appropriate if such a subcommand is given without the `-F` option. Otherwise, the action is disallowed, with a diagnostic message written to standard error.

The following subcommands are supported:

add property-name=property-value

Adds the specified value to the current property values. This subcommand can only be applied to properties that accept multiple values. Use the *set* subcommand for single-value properties.

clear property-name

Clears all of the values for the property. Only those properties that accept multiple assignments, using the *add* subcommand, can be cleared.

commit

Commits the current configuration from memory to stable storage and into the kernel. The configuration must be committed for the changes to take effect. Until the in-memory configuration is committed, you can remove changes with the *revert* subcommand. The *commit* operation is attempted automatically upon completion of a *tncfg* session. Since a configuration must be correct to be committed, this operation automatically does a *verify*.

delete [-F]

Deletes the specified template or zone configuration from the current name service.

Specify the *-F* option to force the action. If the deletion is allowed, its action is instantaneous and the session is terminated.

export [-f output-file]

Displays configuration to standard output. Use the *-f* option to display the configuration to *output-file*. This option produces output in a form suitable for use in a command file.

get host=hostname | IP address[/prefix]

Displays the template name corresponding to the specified host using the kernel's in-memory mapping.

help [usage] [subcommands] [properties] [subcommand] [property]

Displays general help or help about a given topic.

info property-name

Displays information about the current template or zone, or the specified property in a parseable format.

list

Lists the names of the templates or zones that have been configured.

remove property-name=property-value

Removes the specified value from the property. Only those properties that accept multiple assignments, using the *add* subcommand, can be removed.

set property-name=property-value

Sets a given property name to the given value. Properties that can take multiple values are assigned using the *add* subcommand, instead of *set*.

verify

Verifies the current configuration for correctness:

- The required properties are specified;
- the values are valid for each key word;
- the user is authorized to specify the values.

revert [-F]

Causes the configuration to revert to the last committed state. The -F option can be used to force the action.

exit [-F]

Exits the tncfg session. A commit is automatically attempted if needed. You can also use an EOF character to exit tncfg. The -F option can be used to force the action.

示例**示例 1 Using the info Subcommand**

The command below displays the properties of a cipso template are displayed. The subcommand is specified on the command line.

```
example% tncfg -t cipso info
      name=cipso
      host_type=cipso
      doi=1
      min_label=ADMIN_LOW
      max_label=ADMIN_LOW
      host=10.5.233.74
```

示例 2 Using the export Subcommand

The following example shows an interactive session that exports the configuration of a zone in a format the could be imported to another machine with an equivalent zone.

```
example% tncfg -t public
tncfg:public> export
set name=public
set host_type=unlabeled
set doi=1
set def_label="PUBLIC"
set min_label="PUBLIC"
set max_label="CONFIDENTIAL : NEED TO KNOW"
add aux_label="SANDBOX PLAYGROUND"
add host=myserver.oracle.com
add host=10.5.0.0/16
tncfg:public> exit
```

示例 3 Assigning Properties to a Zone

In the following example, the public zone is configured to be a multi-level NFS server.

示例 3 Assigning Properties to a Zone (续)

```
example% tncfg -z public
tncfg:public> info
    name=public
    label=PUBLIC
    visible=no
tncfg:public> add mlp_private=111/tcp
tncfg:public> add mlp_private=111/ucp
tncfg:public> add mlp_private=2049/tcp
tncfg:public> commit
tncfg:public> exit
```

退出状态

- 0 Successful completion.
- 1 An error occurred.

文件

- /etc/security/tsol/tnrhtp
- /etc/security/tsol/tnrhdb
- /etc/security/tsol/tnzonecfg

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/trusted
Interface Stability	See below.

The invocation and subcommands are committed. Output, except for the export and info subcommands, is Not-an-Interface.

另请参见

[tnctl\(1M\)](#), [tnd\(1M\)](#), [tninfo\(1M\)](#), [txzonemgr\(1M\)](#), [zonecfg\(1M\)](#), [nsswitch.conf\(4\)](#), [attributes\(5\)](#), [labels\(5\)](#), [zones\(5\)](#)

附注

The Labeled Zone Manager, [txzonemgr\(1M\)](#), is an alternative application for configuring Trusted Extensions. It invokes the tncfg command internally, and provides an interactive GUI-based user interface.

引用名	tnchkdb – 检查可信网络数据库的文件语法
用法概要	<code>/usr/sbin/tnchkdb [-h path] [-t path] [-z path]</code>
描述	tnchkdb 检查 tnrhttp、tnrhdb 和 tnzonecfg 数据库的语法。缺省情况下，各个文件的 <i>path</i> 为： ▪ <code>/etc/security/tsol/tnrhttp</code> ▪ <code>/etc/security/tsol/tnrhdb</code> ▪ <code>/etc/security/tsol/tnzonecfg</code> 通过在命令行上使用 <code>-h</code> (tnrhdb)、<code>-t</code> (tnrhttp) 和 <code>-z</code> (tnzonecfg) 选项指定路径，可以为任一文件或所有文件指定其他路径。在将文件安装为新系统数据库之前测试修改的文件集时，这些选项非常有用。 将检查所有三个数据库文件的完整性。如果所有文件句法正确并且语义尽可能正确，则 <code>tnchkdb</code> 返回退出状态 0。如果一个或多个文件有错误，则返回退出状态 1。如果命令行有问题（例如不可读取的文件），则返回退出状态 2。错误将写入到标准错误。 为了避免级联错误，当 tnrhttp 中存在错误时，tnrhdb 中的模板名称将无法通过验证。 tnchkdb 可以在任何标签上运行，但是标准 <code>/etc/security/tsol</code> 文件仅在全局区域中才可见。
选项	<code>-h [path]</code> 检查 <i>path</i> 中的 tnrhdb 的语法是否正确。如果未指定 <i>path</i>，则检查 <code>/etc/security/tsol/tnrhdb</code>。 <code>-t [path]</code> 检查 <i>path</i> 中的 tnrhttp 的语法是否正确。如果未指定 <i>path</i>，则检查 <code>/etc/security/tsol/tnrhttp</code>。 <code>-z [path]</code> 检查 <i>path</i> 中的 tnzonecfg 的语法是否正确。如果未指定 <i>path</i>，则检查 <code>/etc/security/tsol/tnzonecfg</code>。
示例	示例1 样例错误消息 tnchkdb 命令检查 CIPSO 错误。在本例中，<code>admin_low</code> 模板的缺省标签有一个错误的值 <code>ADMIN_HIGH</code>。 <pre># tnchkdb checking /etc/security/tsol/tnrhttp ... tnchkdb: def_label classification 7fff is invalid for cipso labels: line 14 entry admin_low tnchkdb: def_label compartments 241-256 must be zero for cipso labels: line 14 entry admin_low checking /etc/security/tsol/tnrhdb ... checking /etc/security/tsol/tnzonecfg ...</pre>

文件

/etc/security/tsol/tnrhdb

可信网络远程主机数据库

/etc/security/tsol/tnrhtp

可信网络远程主机模板

/etc/security/tsol/tnzonecfg

可信区域配置数据库

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/trusted
接口稳定性	请参见下文。

命令行为 Committed（已确定）。输出为 Uncommitted（未确定）。

另请参见

[tnd\(1M\)](#)、[tnctl\(1M\)](#)、[attributes\(5\)](#)

《[Trusted Extensions 配置和管理](#)》

附注

仅当系统配置有 Trusted Extensions 时，本手册页中介绍的功能才可用。

当使用 LDAP 来提供缺少的模板时，可能具有有效但不一致的 tnrhtp 和 tnrhdb 配置。

引用名	tnctl – 配置 Trusted Extensions 网络参数
用法概要	<pre>/usr/sbin/tnctl [-dfv] [-h host [/prefix] [:template]] [-m zone:mlp:shared-mlp] [-t template [:key=val [:key=val]]] [-HTz] file]</pre>
描述	tnctl 提供了一个接口来处理 Solaris 内核中的可信网络参数。 作为 Solaris Trusted Extensions 初始化的一部分，在系统引导期间，tnctl 通过 smf(5) 脚本在全局区域中运行。tnctl 命令不是计划在常规系统管理过程中使用的。相反，如果本地可信网络数据库文件已修改，则管理员首先发出 tnchkdb(1M) 来检查语法，然后使用以下命令刷新内核副本： <pre># svcadm restart svc:/network/tnctl</pre> 有关在正在运行的系统上更改远程主机和模板信息的风险，请参见“警告”部分。
选项	-d 从内核中删除匹配的条目。缺省行为是添加新条目。 在删除 MLP 时，MLP 范围必须完全匹配。MLP 指定为以下格式： <pre>port[-port]/protocol</pre> 其中 <i>port</i> 可以是 1 到 65535 之间的数字。或任何已知服务（请参见 services(4)），<i>protocol</i> 可以是 1 到 255 之间的数字，或任何已知协议（请参见 protocols(4)）。 -f 在加载命令行上指定的条目之前刷新所有内核条目。除非至少一个条目解析成功，否则不进行刷新。 -v 打开详细模式。 -h host[/prefix][:template] 为指定的 <i>host</i> 更新本地主机上的内核远程主机高速缓存，或者，如果给定了模板名称，则更改内核的高速缓存以使用指定的 <i>template</i>。如果没有指定 <i>prefix</i>，则根据用于解释 <i>tnrhdb</i> 的规则确定隐式前缀长度。如果指定了 -d，则不能指定模板名称。 -m zone:mlp:shared-mlp 为指定的 <i>zone</i> 修改内核的多级别端口 (multilevel port, MLP) 配置高速缓存。<i>zone</i> 指定要更新的区域。<i>mlp</i> 和 <i>shared-mlp</i> 为特定于区域的 IP 地址和共享的 IP 地址指定 MLP。<i>shared-mlp</i> 字段仅在全局区域中有效。 -t template[key=val[:key=val]] 为 <i>template</i> 更新内核模板高速缓存，或者，如果给定了 <i>key=val</i> 对的列表，则更改内核的高速缓存以使用指定的条目。如果指定了 -d，则不能指定 <i>key=val</i> 对。 -T file 将 <i>file</i> 中的所有模板条目加载到内核高速缓存中。

- H *file*
将 *file* 中的所有远程主机条目加载到内核高速缓存中。
- z *file*
仅将 *file* 中的全局区域的 MLP 加载到内核高速缓存中。要为非全局区域重新加载 MLP，请重新引导区域：

```
# zoneadm -z non-global zone reboot
```

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/trusted
接口稳定性	Uncommitted（未确定）

- 文件
- /etc/security/tsol/tnrhdb 可信网络远程主机数据库
 - /etc/security/tsol/tnrhttp 可信网络远程主机模板
 - /etc/security/tsol/tnzonecfg 可信区域配置数据库
 - /etc/nsswitch.conf 名称服务转换器的配置文件

另请参见 [svcs\(1\)](#)、[svcadm\(1M\)](#)、[tninfo\(1M\)](#)、[tnd\(1M\)](#)、[tnchkdb\(1M\)](#)、[zoneadm\(1M\)](#)、[nsswitch.conf\(4\)](#)、[pr](#)
《[Trusted Extensions 配置和管理](#)》

警告 当网络处于运行状态时更改模板可能会改变某些主机（数量不定）的安全视图。

附注 仅当系统配置有 Trusted Extensions 时，本手册页中介绍的功能才可用。

tnctl 服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

```
svc:/network/tnctl
```

可以使用 [svcs\(1\)](#) 查询该服务的状态。可以使用 [svcadm\(1M\)](#) 对该服务执行管理操作（如刷新内核高速缓存），如下所示：

```
svcadm restart svc:/network/tnctl
```

引用名	tnd – 可信网络守护进程						
用法概要	<code>/usr/sbin/tnd [-p poll-interval]</code>						
描述	tnd（可信网络守护进程）使用可信的网络数据库初始化内核，还根据需要从 LDAP 服务器和本地文件重新装入数据库。tnd 在装入配置数据库时遵循 <code>nsswitch.conf(4)</code> 文件中指定的顺序。 对 tnd 的设计安排是，仅当系统已配置为 LDAP 客户机时，才在引导进程中通过 <code>svc:/network/tnd:smf(5)</code> 服务在全局区域中启动。只以本地文件作为可信网络数据库的系统使用 <code>tnctl(1M)</code> 而不是 tnd。 tnd 将以下数据库装入到内核中：远程主机数据库 <code>tnrhdb</code> 和远程主机模板数据库 <code>tnrhtp</code>。tnd 还定期扫描关联的 LDAP 数据库或本地数据库中的更改并相应地更新内核高速缓存。 如果修改了本地可信网络数据库文件，则管理员应运行 <code>tnchkdb(1M)</code> 来检查语法，还应运行 <code>svcadm refresh svc:/network/tnd</code> 来启动由 tnd 执行的即时数据库扫描。 tnd 的设计安排是通过 <code>smf(5)</code> 脚本启动并在全局区域中运行。以下 <code>svcadm</code> 命令指示 tnd 执行特定操作： <table><tr><td><code>svcadm refresh svc:/network/tnd</code></td><td>启动对本地和 LDAP <code>tnrhdb</code> 和 <code>tnrhtp</code> 数据库的重新扫描。tnd 使用找到的任何更改来更新内核数据库。</td></tr><tr><td><code>svcadm disable svc:/network/tnd</code></td><td>终止 tnd 守护进程。不对内核数据库进行任何更改。</td></tr></table> 在由以下服务管理工具 (Service Management Facility, SMF) 属性的值确定的调试模式下运行 tnd： <pre>tnd/debug_level = 0</pre> 如上所示，值 0 阻止收集调试信息；1 打开调试。缺省值为 "0"。调试输出将发送到 <code>/var/tsxol/tndlog</code> 日志文件。 选项 <table><tr><td><code>-p poll-interval</code></td><td>将轮询间隔设置为 <code>poll-interval</code> 秒。<code>poll-interval</code> 缺省值是 1800 秒（30 分钟）。</td></tr></table> 示例 示例1 更改轮询间隔 以下命令将轮询间隔更改为 1 小时，并将此间隔放入 SMF 系统信息库中。下次引导时，tnd 轮询间隔将变为 1 小时。 <pre># svccfg -s network/tnd setprop tnd/poll_interval=3600</pre> 以下命令更改轮询间隔，但是不更新系统信息库。下次引导时，tnd 轮询间隔保持为缺省值 30 分钟。	<code>svcadm refresh svc:/network/tnd</code>	启动对本地和 LDAP <code>tnrhdb</code> 和 <code>tnrhtp</code> 数据库的重新扫描。tnd 使用找到的任何更改来更新内核数据库。	<code>svcadm disable svc:/network/tnd</code>	终止 tnd 守护进程。不对内核数据库进行任何更改。	<code>-p poll-interval</code>	将轮询间隔设置为 <code>poll-interval</code> 秒。 <code>poll-interval</code> 缺省值是 1800 秒（30 分钟）。
<code>svcadm refresh svc:/network/tnd</code>	启动对本地和 LDAP <code>tnrhdb</code> 和 <code>tnrhtp</code> 数据库的重新扫描。tnd 使用找到的任何更改来更新内核数据库。						
<code>svcadm disable svc:/network/tnd</code>	终止 tnd 守护进程。不对内核数据库进行任何更改。						
<code>-p poll-interval</code>	将轮询间隔设置为 <code>poll-interval</code> 秒。 <code>poll-interval</code> 缺省值是 1800 秒（30 分钟）。						

示例1 更改轮询间隔 (续)

文件

tnd -p 3600

/etc/security/tsol/tnrhdb可信网络远程主机数据库

/etc/security/tsol/tnrhttp可信网络远程主机模板

/etc/security/tsol/tnzonecfg可信区域配置数据库

/etc/nsswitch.conf名称服务转换器的配置文件

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/trusted
接口稳定性	请参见下文。

命令调用为 Committed（已确定）。服务是 Private（专用）。

另请参见

[svcs\(1\)](#)、[svcadm\(1M\)](#)、[tninfo\(1M\)](#)、[tnctl\(1M\)](#)、[tnchkdb\(1M\)](#)、[nsswitch.conf\(4\)](#)、[attributes\(5\)](#)、
[《Trusted Extensions 配置和管理》](#)

附注

仅当系统配置有 Trusted Extensions 时，本手册页中介绍的功能才可用。

tnd 服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

svc:/network/tnd

可以使用 [svcs\(1\)](#) 查询该服务的状态。可以使用 [svcadm\(1M\)](#) 对该服务执行管理操作（如请求重新启动守护进程），如下所示：

svcadm restart svc:/network/tnd

引用名	tninfo – 显示内核级网络信息和统计信息
用法概要	<code>/usr/sbin/tninfo [-h <i>hostname</i>] [-m <i>zone-name</i>] [-t <i>template</i>]</code>
描述	tninfo 提供了一个接口来检索并显示内核级网络信息和统计信息。
选项	<code>-h <i>hostname</i></code> 显示远程主机高速缓存中指定的主机的安全结构。输出应当反映 <code>tnrhdb</code> 数据库中指定的内容。 <code>-m <i>zone-name</i></code> 显示与指定区域关联的 MLP 配置。输出应当反映 <code>tnzonecfg</code> 数据库中指定的内容。 <code>-t <i>template</i></code> 显示与指定的 <i>template</i> 关联的结构。输出应当反映 <code>tnrhtp</code> 数据库中指定的内容。

示例

示例1 显示内核中缓存的远程主机结构

此示例显示内核中缓存的远程主机结构。输出反映 `tnrhdb` 数据库中的定义。

```
# tninfo -h machine1
IP address= 192.168.8.61
Template = cipso
```

示例2 显示全局区域的多级别端口

此示例为全局区域显示内核中缓存的 MLP。输出反映 `tnzonecfg` 数据库中的定义，以及动态分配的 MLP。`private` 表示特定于区域的 MLP。

```
# tninfo -m global
private:23/tcp;111/tcp;111/udp;515/tcp;2049/tcp;6000-6003/tcp;
        32812/tcp;36698/ip;38634/tcp;64365/ip
shared: 6000-6003/tcp
```

示例3 显示 `cipso` 模板定义

此示例显示内核中缓存的 `cipso` 模板定义。输出反映 `tnrhtp` 数据库中的定义。

```
# tninfo -t cipso
=====
Remote Host Template Table Entries:
-----
template: cipso
host_type: CIPSO
doi: 1
min_sl: ADMIN_LOW
hex: ADMIN_LOW
max_sl: ADMIN_HIGH
hex: ADMIN_HIGH
```

文件 /etc/security/tsol/tnrhdb 可信网络远程主机数据库
 /etc/security/tsol/tnrhtp 可信网络远程主机模板
 /etc/security/tsol/tnzonecfg 可信区域配置数据库

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/trusted
接口稳定性	请参见下文。

命令行为 Committed（已确定）。输出为 Uncommitted（未确定）。

另请参见 [tnd\(1M\)](#)、[tnctl\(1M\)](#)、[attributes\(5\)](#)

 《[Trusted Extensions 配置和管理](#)》

附注 仅当系统配置有 Trusted Extensions 时，本手册页中介绍的功能才可用。

引用名	tpmadm – 管理可信平台模块
用法概要	<pre>tpmadm status tpmadm init tpmadm clear [owner lock] tpmadm auth tpmadm keyinfo [uuid] tpmadm deletekey uuid tpmadm pcrextend pcr [filename] tpmadm pcrreset pcr</pre>
描述	可信平台模块 (Trusted Platform Module, TPM) 是为受保护的密钥存储和用于引导操作系统的软件的可靠度量提供的一个硬件组件。tpmadm 实用程序用于初始化并管理 TPM，以使其可供操作系统和其他程序使用。 TPM 子系统可以存储和管理数目不受限的密钥供操作系统和用户使用。每个密钥由一个全局唯一标识符（或称 UUID）来标识。 虽然 TPM 在任一给定时间只能保留有限数量的密钥，但支持软件可以根据需要自动加载和卸载密钥。当密钥存储在 TPM 外部时，其父密钥始终会对其进行加密或“包装”，以使得密钥不会在 TPM 外部以可读形式公开。 TPM 必须由平台所有者进行初始化，然后才可供使用。此过程包括设置用于对特权操作进行授权的所有者口令。 虽然 TPM 所有者类似于传统的超级用户，但两者仍有重要的区别。首先，对 TPM 功能的访问与进程特权无关。不论调用方进程的特权等级如何，所有特权操作都需要知道所有者口令。其次，TPM 所有者无法覆盖由 TPM 密钥保护的数据的访问控制。所有者可以通过重新初始化 TPM 有效地销毁数据，但是无法访问已使用其他用户所拥有的 TPM 密钥加密的数据。
子命令	以如下格式使用下列子命令： <pre># tpmadm <subcommand> [operand]</pre> status 报告 TPM 的状态信息。输出包括是否已建立了 TPM 的所有权、当前的 PCR 内容，以及 TPM 资源（如通信会话和已加载的密钥）的使用情况等等基本信息。 init 初始化 TPM 以供使用。这包括通过设置所有者授权口令来获取 TPM 的所有权。获取 TPM 的所有权将创建新的存储根密钥，这是该 TPM 创建的所有密钥的祖先。在发出此命令后，要想重新初始化 TPM，必须使用 BIOS 操作重置 TPM。 auth 更改 TPM 的所有者授权口令。

clear lock
清除验证尝试失败的次数。验证尝试失败多次后，TPM 对后续尝试的响应将变得更慢，以挫败那些试图通过穷举搜索来获取所有者口令的人。此命令（需要正确的所有者口令）将重置失败尝试的次数。

clear owner
取消激活 TPM 并将其返回无所有者状态。此操作（需要当前的 TPM 所有者口令）使绑定到 TPM 的所有密钥和数据失效。要想重新使用 TPM，必须重新启动系统，从 BIOS 或 ILOM 预引导环境重新激活 TPM，且必须使用 `tpmadm init` 命令重新初始化 TPM。

keyinfo *[uuid]*
报告存储在 TPM 子系统中的密钥的相关信息。未指定其他参数时，此子命令将生成所有密钥的简短列表。如果指定了单个密钥的 UUID，则将显示该密钥的详细信息。

deletekey *uuid*
从 TPM 子系统的持久性存储中删除具有指定 UUID 的密钥。

pcrextend *pcr [filename]*
创建 *filename* 内容的 SHA-1 散列，并使用散列值作为将扩展的数据对指示的 PCR 执行 PCR 扩展操作。如果未指定文件名，则从 `stdin` 读取数据。

pcrreset *pcr*
将指示的 PCR 重置为其最初状态（全零）。

退出状态 在完成所请求的操作之后，`tpmadm` 以下列状态值之一退出。

- 0
成功终止。
- 1
失败。无法完成所请求的操作。
- 2
用法错误。调用 `tpmadm` 命令时使用了无效的参数。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Committed（已确定）

另请参见 [prof_attr\(4\)](#)、[attributes\(5\)](#)

另请参见 SUNWtss 软件包中的 `tcsd(8)` 手册页。

TCG 软件栈 (TCG Software Stack, TSS) 规范：

<https://www.trustedcomputinggroup.org/specs/TSS>（截止到出版日期）

附注

tpmadm 通过 tcscd 服务与 TPM 设备通信。在使用 tpmadm 命令之前，tcscd 必须正在运行。如果 tcscd 没有正在运行，tpmadm 将生成以下错误：

```
Connect context: Communication failure (0x3011)
```

有关更多详细信息，请参见 tcscd(8)。

引用名

tracert – print the route packets take to network host

用法概要

```
tracert [-adFIlnSvx] [-A addr_family] [-c traffic_class]
        [-f first_hop] [-g gateway [-g gateway...] | -r]
        [-i iface] [-L flow_label] [-m max_hop]
        [-P pause_sec] [-p port] [-Q max_timeout]
        [-q nqueries] [-s src_addr] [-t tos] [-w wait_time] host
        [packetlen]
```

描述

The Internet is a large and complex aggregation of network hardware, connected by gateways. Tracking the route a packet follows can be difficult. The utility `tracert` traces the route that an IP packet follows to another internet host.

The `tracert` utility utilizes the both the IPv4 and IPv6 protocols. Use the `-A` option to override the default behavior. `tracert` uses the IPv4 protocol *ttl* (time to live) field or the IPv6 field *hop limit*. It attempts to elicit an ICMP or ICMP6 `TIME_EXCEEDED` response from each *gateway* along the path, and a `PORT_UNREACHABLE` (or `ECHO_REPLY` if `-I` is used) response from the destination host. It starts by sending probes with a *ttl* or *hop limit* of 1 and increases by 1 until it either gets to the host, or it hits the maximum *max_hop*. The default maximum *max_hop* is 30 hops, but this can be set by the `-m` option.

Three probes are sent at each *ttl* (*hop limit*) setting, and a line is printed showing the *ttl* (*hop limit*), the hostname and the address of the gateway, and the *rtt* (round trip time) of each probe. The number of probes may be specifically set using the `-q` option. If the probe answers come from different gateways, the hostname and the address of each responding system will be printed. If there is no response within a 5 second timeout interval, an asterisk (*) is displayed for that probe. The `-w` option may be used to set the timeout interval. Other possible annotations that may appear after the time are:

! the *ttl* (*hop limit*) value in the received packet is <= 1.

!H host unreachable.

!X communication administratively prohibited.

<!N> ICMP (ICMP6) unreachable code N.

The following annotations appear only for IPv4:

!F fragmentation needed. This should never occur. If this is seen, the associated gateway is broken.

!N network unreachable.

!P
protocol unreachable.

!S
source route failed. It is likely that the gateway does not support source routing.

!T
unreachable for the specified tos (type-of-service).

!U
source host isolated or precedence problem.

The following annotations appear only for IPv6:

!A
host unreachable for a reason other than lack of an entry in the routing table.

!B
packet too big.

!E
destination is not a neighbor.

!R
unrecognized next header.

If almost all the probes result in some kind of unreachable code, then `tracert` gives up and exits.

The destination *host* is not supposed to process the UDP probe packets, so the destination *port* default is set to an unlikely value. However, if some application on the destination is using that value, the value of *port* can be changed with the `-p` option.

The only mandatory parameter is the destination *host* name or IP number. The default probe datagram length is 40 bytes (60 bytes for IPv6), but this may be increased by specifying a packet length (in bytes) after the destination *host* name.

All integer arguments to `tracert` can be specified in either decimal or hexadecimal notation. For example, *packetlen* can be specified either as 256 or 0x100.

选项

`-A addr_family`

Specify the address family of the target host. *addr_family* can be either `inet` or `inet6`. Address family determines which protocol to use. For an argument of `inet`, IPv4 is used. For `inet6`, IPv6 is used.

By default, if the name of a host is provided, not the literal IP address, and a valid IPv6 address exists in the name service database, `tracert` will use this address. Otherwise, if the name service database contains an IPv4 address, it will try the IPv4 address.

Specify the address family `inet` or `inet6` to override the default behavior. If the argument specified is `inet`, `tracert` will use the IPv4 address associated with the hostname. If none exists, `tracert` will state that the host is unknown and exit. It will not try to determine if an IPv6 address exists in the name service database.

If the specified argument is `inet6`, `tracert` will use the IPv6 address that is associated with the hostname. If none exists, `tracert` will state that the host is unknown and exit.

-a

Probe all of the addresses of a multi-homed destination. The output looks like `tracert` has been run once for each IP address of the destination. If this option is used together with `-A`, `tracert` probes only the addresses that are of the specified address family. While probing one of the addresses of the destination, user can skip to the next address by sending a `SIGINT`, or exit `tracert` by sending a `SIGQUIT` signal. See [signal\(3C\)](#)

-c *traffic_class*

Specify the traffic class of probe packets. The value must be an integer in the range from 0 to 255. Gateways along the path may route the probe packet differently depending upon the value of *traffic_class* set in the probe packet. This option is valid only on IPv6.

-d

Set the `SO_DEBUG` socket option.

-F

Turn off fragmentation. For IPv4, this means setting the Don't Fragment bit. For IPv4 and IPv6, this means do not allow fragmentation as the datagrams are sent. If the packetlen exceeds the MTU, then `tracert` may report that sending failed due to Message too long.

-f *first_hop*

Set the starting *ttl* (*hop limit*) value to *first_hop*, to override the default value 1. `tracert` skips processing for those intermediate gateways which are less than *first_hop* hops away.

-g *gateway*

Specify a loose source route *gateway*. The user can specify more than one *gateway* by using `-g` for each gateway. The maximum number of gateways is 8 for IPv4 and 127 for IPv6. Note that some factors such as the link MTU can further limit the number of gateways for IPv6. This option cannot be used with the `-r` option.

Only users with the `{PRIV_NET_RAWACCESS}` privilege can specify a loose source route with this option.

-I

Use ICMP (ICMP6) ECHO instead of UDP datagrams.

-i *iface*

For IPv4, this option specifies a network interface to obtain the source IP address. This is normally only useful on a multi-homed host. The `-s` option is also another way to do this.

For IPv6, it specifies the network interface on which probe packets are transmitted. The argument can be either an interface index, for example, 1, 2, or an interface name, for example, `eri0`, `hme0`.

-L *flow_label*

Specify the flow label of probe packets. The value must be an integer in the range from 0 to 1048575. This option is valid only on IPv6.

-l

Print the value of the *ttl (hop limit)* field in each packet received.

-m *max_hop*

Set the maximum *ttl (hop limit)* used in outgoing probe packets. The default is 30 hops, which is the same default used for TCP connections.

-n

Print hop addresses numerically rather than symbolically and numerically. This saves a nameserver address-to-name lookup for each gateway found on the path.

-P *pause_sec*

Specify a delay, in seconds, to pause between probe packets. This may be necessary if the final destination does not accept undeliverable packets in bursts. By default, `tracert` sends the next probe as soon as it has received a reply. Note that *pause_sec* is a real number.

-p *port*

Set the base UDP *port* number used in probes. The default is 33434. `tracert` hopes that nothing is listening on UDP *ports* ($\text{base} + (\text{nhops} - 1) * \text{nqueries}$) to ($\text{base} + (\text{nhops} * \text{nqueries}) - 1$) at the destination host, so that an ICMP (ICMP6) `PORT_UNREACHABLE` message will be returned to terminate the route tracing. If something is listening on a *port* in the default range, this option can be used to select an unused *port* range. *nhops* is defined as the number of hops between the source and the destination.

-Q *max_timeout*

Stop probing this hop after *max_timeout* consecutive timeouts are detected. The default value is 5. Useful in combination with the `-q` option if you have specified a large *nqueries* probe count.

-q *nqueries*

Set the desired number of probe queries. The default is 3.

-r

Bypass the normal routing tables and send directly to a host on an attached network. If the host is not on a directly-attached network, an error is returned. This option can be used to send probes to a local host through an interface that has been dropped by the router daemon. See [in .routed\(1M\)](#). You cannot use this option if the `-g` option is used.

-S

Display a summary of how many probes were not answered for each hop.

-s *src_addr*

Use the following address, which usually is given as a literal IP address, not a hostname, as the source address in outgoing probe packets. On multi-homed hosts, those with more than one IP address, this option can be used to force the source address to be something other than the IP address `tracert` picks by default. If the IP address is not one of this machine's interface addresses, an error is returned and nothing is sent. For IPv4, when used together with the `-i` option, the given IP address should be configured on the specified interface. Otherwise, an error will be returned. In the case of IPv6, the interface name and the source address do not have to match.

-t *tos*

Set the *tos*(type-of-service) in probe packets to the specified value. The default is zero. The value must be an integer in the range from 0 to 255. Gateways along the path may route the probe packet differently depending upon the *tos* value set in the probe packet. This option is valid only on IPv4.

-v

Verbose output. For each hop, the size and the destination of the response packets is displayed. Also ICMP (ICMP6) packets received other than `TIME_EXCEEDED` and `UNREACHABLE` are listed as well.

-w *waittime*

Set the time, in seconds, to wait for a response to a probe. The default is 5 seconds.

-x

Prevent `tracert` from calculating checksums. Checksums are usually required for the last hop when using ICMP ECHO probes. This option is valid only on IPv4. See the `-I` option.

When specified from within a shared-IP zone, this option has no effect as the checksum is always calculated by the operating system in this case.

操作数

The following operands are supported:

host

The network host.

示例**示例 1 Sample Output From the `tracert` Utility**

Some sample output from the `tracert` utility might be:

```
istanbul% tracert london
tracert: Warning: london has multiple addresses; \
    using 4::114:a00:20ff:ab3d:83ed
tracert: Warning: Multiple interfaces found; \
    using 4::56:a00:20ff:fe93:8dde @ eri0:2
tracert to london (4::114:a00:20ff:ab3d:83ed), 30 hops max, \
    60 byte packets
1  frbldg7c-86 (4::56:a00:20ff:felf:65a1)  1.786 ms  1.544 ms  1.719 ms
```

示例 1 Sample Output From the traceroute Utility (续)

```

2  frbldg7b-77 (4::255:0:0:c0a8:517)  2.587 ms 3.001 ms  2.988 ms
3  london (4::114:a00:20ff:ab3d:83ed)  3.122 ms 2.744 ms  3.356 ms

```

The target host, london, has both IPv4 and IPv6 addresses in the name service database. According to the default behavior, traceroute uses IPv6 address of the destination host.

示例 2 Using the traceroute Utility For a Host Which has Only IPv4 Addresses

In the following examples, traceroute is tracking the route to host sanfrancisco, which has only IPv4 addresses in the name service database. Therefore traceroute uses only IPv4 addresses. The following shows the 7-hop path that a packet would follow from the host istanbul to the host sanfrancisco.

```

istanbul% traceroute sanfrancisco
traceroute: Warning: Multiple interfaces found; using 172.31.86.247 @eri0
traceroute to sanfrancisco (172.29.64.39), 30 hops max, 40 byte packets
1  frbldg7c-86 (172.31.86.1)  1.516 ms  1.283 ms  1.362 ms
2  bldg1a-001 (172.31.1.211)  2.277 ms  1.773 ms  2.186 ms
3  bldg4-bldg1 (172.30.4.42)  1.978 ms  1.986 ms  13.996 ms
4  bldg6-bldg4 (172.30.4.49)  2.655 ms  3.042 ms  2.344 ms
5  ferbldg11a-001 (172.29.1.236)  2.636 ms  3.432 ms  3.830 ms
6  frbldg12b-153 (172.29.153.72)  3.452 ms  3.146 ms  2.962 ms
7  sanfrancisco (172.29.64.39)  3.430 ms  3.312 ms  3.451 ms

```

示例 3 Using the traceroute Utility With Source Routing

The following example shows the path of a packet that goes from istanbul to sanfrancisco through the hosts cairo and paris, as specified by the -g option. The -I option makes traceroute send ICMP ECHO probes to the host sanfrancisco. The -i options sets the source address to the IP address configured on the interface qe0.

```

istanbul% traceroute -g cairo -g paris -i qe0 -q 1 -I sanfrancisco
traceroute to sanfrancisco (172.29.64.39), 30 hops max, 56 byte packets
1  frbldg7c-86 (172.31.86.1)  2.012 ms
2  flrbldg7u (172.31.17.131)  4.960 ms
3  cairo (192.168.163.175)  4.894 ms
4  flrbldg7u (172.31.17.131)  3.475 ms
5  frbldg7c-017 (172.31.17.83)  4.126 ms
6  paris (172.31.86.31)  4.086 ms
7  frbldg7b-82 (172.31.82.1)  6.454 ms
8  bldg1a-001 (172.31.1.211)  6.541 ms
9  bldg6-bldg4 (172.30.4.49)  6.518 ms
10 ferbldg11a-001 (172.29.1.236)  9.108 ms
11 frbldg12b-153 (172.29.153.72)  9.634 ms
12 sanfrancisco (172.29.64.39)  14.631 ms

```

退出状态

The following exit values are returned:

0

Successful operation.

 ≥ 0

An error occurred.

属性

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

```
netstat(1M), signal(3C), ping(1M), attributes(5), privileges(5), zones(5)
```

警告

This utility is intended for use in network testing, measurement and management. It should be used primarily for manual fault isolation. Because of the load it could impose on the network, it is unwise to use `traceroute(1M)` during normal operations or from automated scripts.

引用名 trapstat – report trap statistics

用法概要

```
/usr/sbin/trapstat [-t | -T | -e entry]
                  [-C processor_set_id | -c cpulist] [-P] [-a]
                  [-r rate] [ [interval [count]] | command | [args]]
```

```
/usr/sbin/trapstat -l
```

描述

The trapstat utility gathers and displays run-time trap statistics on UltraSPARC-based systems. The default output is a table of trap types and CPU IDs, with each row of the table denoting a trap type and each column of the table denoting a CPU. If standard output is a terminal, the table contains as many columns of data as can fit within the terminal width; if standard output is not a terminal, the table contains at most six columns of data. By default, data is gathered and displayed for all CPUs; if the data cannot fit in a single table, it is printed across multiple tables. The set of CPUs for which data is gathered and displayed can be optionally specified with the *-c* or *-C* option.

Unless the *-r* option or the *-a* option is specified, the value displayed in each entry of the table corresponds to the number of traps per second. If the *-r* option is specified, the value corresponds to the number of traps over the interval implied by the specified sampling rate; if the *-a* option is specified, the value corresponds to the accumulated number of traps since the invocation of trapstat.

By default, trapstat displays data once per second, and runs indefinitely; both of these behaviors can be optionally controlled with the *interval* and *count* parameters, respectively. The *interval* is specified in seconds; the *count* indicates the number of intervals to be executed before exiting. Alternatively, *command* can be specified, in which case trapstat executes the provided command and continues to run until the command exits. A positive integer is assumed to be an *interval*; if the desired *command* cannot be distinguished from an integer, the full path of *command* must be specified.

UltraSPARC I (obsolete), II, and III handle translation lookaside buffer (TLB) misses by trapping to the operating system. TLB miss traps can be a significant component of overall system performance for some workloads; the *-t* option provides in-depth information on these traps. When run with this option, trapstat displays both the rate of TLB miss traps and the percentage of time spent processing those traps. Additionally, TLB misses that hit in the translation storage buffer (TSB) are differentiated from TLB misses that further miss in the TSB. (The TSB is a software structure used as a translation entry cache to allow the TLB to be quickly filled; it is discussed in detail in the *UltraSPARC II User's Manual*.) The TLB and TSB miss information is further broken down into user- and kernel-mode misses.

Workloads with working sets that exceed the TLB reach may spend a significant amount of time missing in the TLB. To accommodate such workloads, the operating system supports multiple page sizes: larger page sizes increase the effective TLB reach and thereby reduce the number of TLB misses. To provide insight into the relationship between page size and TLB miss rate, trapstat optionally provides in-depth TLB miss information broken down by page size using the *-T* option. The information provided by the *-T* option is a superset of that provided by the *-t* option; only one of *-t* and *-T* can be specified.

选项

The following options are supported:

- a Displays the number of traps as accumulating, monotonically increasing values instead of per-second or per-interval rates.
- c cpulist Enables trapstat only on the CPUs specified by *cpulist*.

cpulist can be a single processor ID (for example, 4), a range of processor IDs (for example, 4-6), or a comma separated list of processor IDs or processor ID ranges (for example, 4,5,6 or 4,6-8).
- C processor_set_id Enables trapstat only on the CPUs in the processor set specified by *processor_set_id*.

trapstat modifies its output to always reflect the CPUs in the specified processor set. If a CPU is added to the set, trapstat modifies its output to include the added CPU; if a CPU is removed from the set, trapstat modifies its output to exclude the removed CPU. At most one processor set can be specified.
- e entrylist Enables trapstat only for the trap table entry or entries specified by *entrylist*. A trap table entry can be specified by trap number or by trap name (for example, the level-10 trap can be specified as 74, 0x4A, 0x4a, or level-10).

entrylist can be a single trap table entry or a comma separated list of trap table entries. If the specified trap table entry is not valid, trapstat prints a table of all valid trap table entries and values. A list of valid trap table entries is also found in *The SPARC Architecture Manual, Version 9* and the *Sun Microelectronics UltraSPARC II User's Manual*. If the parsable option (-P) is specified in addition to the -e option, the format of the data is as follows:

Field	Contents
1	Timestamp (nanoseconds since start)
2	CPU ID
3	Trap number (in hexadecimal)
4	Trap name
5	Trap rate per interval

Each field is separated with whitespace. If the format is modified, it will be modified by adding potentially new fields beginning with field 6; exant fields will remain unchanged.

-l Lists trap table entries. By default, a table is displayed containing all valid trap numbers, their names and a brief description. The trap name is used in both the default output and in the *entrylist* parameter for the -e argument. If the parsable option (-P) is specified in addition to the -l option, the format of the data is as follows:

Field	Contents
1	Trap number in hexadecimal
2	Trap number in decimal
3	Trap name
Remaining	Trap description

-P Generates parsable output. When run without other data gathering modifying options (that is, -e, -t or -T), trapstat's the parsable output has the following format:

Field	Contents
1	Timestamp (nanoseconds since start)
2	CPU ID
3	Trap number (in hexadecimal)
4	Trap name
5	Trap rate per interval

Each field is separated with whitespace. If the format is modified, it will be modified by adding potentially new fields beginning with field 6; extant fields will remain unchanged.

-r *rate* Explicitly sets the sampling rate to be *rate* samples per second. If this option is specified, trapstat's output changes from a traps-per-second to traps-per-sampling-interval.

-t Enables TLB statistics.

A table is displayed with four principal columns of data: *itlb-miss*, *itsb-miss*, *dtlb-miss*, and *dtsb-miss*. The columns contain both the rate of the corresponding event and the percentage of CPU time spent processing the event. The percentage of CPU time is given only in terms of a single CPU. The rows of the table correspond to CPUs,

with each CPU consuming two rows: one row for user-mode events (denoted with u) and one row for kernel-mode events (denoted with k). For each row, the percentage of CPU time is totalled and displayed in the rightmost column. The CPUs are delineated with a solid line. If the parsable option (-P) is specified in addition to the -t option, the format of the data is as follows:

Field	Contents
1	Timestamp (nanoseconds since start)
2	CPU ID
3	Mode (k denotes kernel, u denotes user)
4	I-TLB misses
5	Percentage of time in I-TLB miss handler
6	I-TSB misses
7	Percentage of time in I-TSB miss handler
8	D-TLB misses
9	Percentage of time in D-TLB miss handler
10	D-TSB misses
11	Percentage of time in D-TSB miss handler

Each field is separated with whitespace. If the format is modified, it will be modified by adding potentially new fields beginning with field 12; extant fields will remain unchanged.

-T

Enables TLB statistics, with page size information. As with the -t option, a table is displayed with four principal columns of data: *itlb-miss*, *itsb-miss*, *dtlb-miss*, and *dtsb-miss*. The columns contain both the absolute number of the corresponding event, and the percentage of CPU time spent processing the event. The percentage of CPU time is given only in terms of a single CPU. The rows of the table correspond to CPUs, with each CPU consuming two sets of rows: one set for user-level events (denoted with u) and one set for kernel-level events (denoted with k). Each set, in turn, contains as many rows as

there are page sizes supported (see [getpagesizes\(3C\)](#)). For each row, the percentage of CPU time is totalled and displayed in the right-most column. The two sets are delineated with a dashed line; CPUs are delineated with a solid line. If the parsable option (-P) is specified in addition to the -T option, the format of the data is as follows:

Field	Contents
1	Timestamp (nanoseconds since start)
2	CPU ID
3	Mode k denotes kernel, u denotes user)
4	Page size, in decimal
5	I-TLB misses
6	Percentage of time in I-TLB miss handler
7	I-TSB misses
8	Percentage of time in I-TSB miss handler
9	D-TLB misses
10	Percentage of time in D-TLB miss handler
11	D-TSB misses
12	Percentage of time in D-TSB miss handler

Each field is separated with whitespace. If the format is modified, it will be modified by adding potentially new fields beginning with field 13; extant fields will remain unchanged.

示例

示例 1 Using trapstat Without Options

When run without options, trapstat displays a table of trap types and CPUs. At most six columns can fit in the default terminal width; if (as in this example) there are more than six CPUs, multiple tables are displayed:

```
example# trapstat
vct  name          |      cpu0      cpu1      cpu4      cpu5      cpu8      cpu9
-----+-----
 24 cleanwin      |      6446      4837      6368      2153      2623      1321
 41 level-1       |       100         0         0         0         1         0
```

示例 1 Using trapstat Without Options (续)

44 level-4		0	1	1	1	0	0
45 level-5		0	0	0	0	0	0
47 level-7		0	0	0	0	9	0
49 level-9		100	100	100	100	100	100
4a level-10		100	0	0	0	0	0
4d level-13		6	10	7	16	13	11
4e level-14		100	0	0	0	1	0
60 int-vec		2607	2740	2642	2922	2920	3033
64 itlb-miss		3129	2475	3167	1037	1200	569
68 dtlb-miss		121061	86162	109838	37386	45639	20269
6c dtlb-prot		997	847	1061	379	406	184
84 spill-user-32		2809	2133	2739	200806	332776	454504
88 spill-user-64		45819	207856	93487	228529	68373	77590
8c spill-user-32-cln		784	561	767	274	353	215
90 spill-user-64-cln		9	37	17	39	12	13
98 spill-kern-64		62913	50145	63869	21916	28431	11738
a4 spill-asuser-32		1327	947	1288	460	572	335
a8 spill-asuser-64		26	48	18	54	10	14
ac spill-asuser-32-cln		4580	3599	4555	1538	1978	857
b0 spill-asuser-64-cln		26	0	0	2	0	0
c4 fill-user-32		2862	2161	2798	191746	318115	435850
c8 fill-user-64		45813	197781	89179	217668	63905	74281
cc fill-user-32-cln		3802	2833	3733	10153	16419	19475
d0 fill-user-64-cln		329	10105	4873	10603	4235	3649
d8 fill-kern-64		62519	49943	63611	21824	28328	11693
108 syscall-32		2285	1634	2278	737	957	383
126 self-xcall		100	0	0	0	0	0
vct name		cpu12	cpu13	cpu14	cpu15		
24 cleanwin		5435	4232	6302	6104		
41 level-1		0	0	0	0		
44 level-4		2	0	0	1		
45 level-5		0	0	0	0		
47 level-7		0	0	0	0		
49 level-9		100	100	100	100		
4a level-10		0	0	0	0		
4d level-13		15	11	22	11		
4e level-14		0	0	0	0		
60 int-vec		2813	2833	2738	2714		
64 itlb-miss		2636	1925	3133	3029		
68 dtlb-miss		90528	70639	107786	103425		
6c dtlb-prot		819	675	988	954		
84 spill-user-32		175768	39933	2811	2742		
88 spill-user-64		0	241348	96907	118298		
8c spill-user-32-cln		681	513	753	730		

示例1 Using trapstat Without Options (续)

90	spill-user-64-cln		0	42	16	20
98	spill-kern-64		52158	40914	62305	60141
a4	spill-asuser-32		1113	856	1251	1208
a8	spill-asuser-64		0	64	16	24
ac	spill-asuser-32-cln		3816	2942	4515	4381
b0	spill-asuser-64-cln		0	0	0	0
c4	fill-user-32		170744	38444	2876	2784
c8	fill-user-64		0	230381	92941	111694
cc	fill-user-32-cln		8550	3790	3612	3553
d0	fill-user-64-cln		0	10726	4495	5845
d8	fill-kern-64		51968	40760	62053	59922
108	syscall-32		1839	1495	2144	2083
126	self-xcall		0	0	0	0

示例2 Using trapset with CPU Filtering

The -c option can be used to limit the CPUs on which trapstat is enabled. This example limits CPU 1 and CPUs 12 through 15.

example# **trapstat -c 1,12-15**

vct	name		cpu1	cpu12	cpu13	cpu14	cpu15
24	cleanwin		6923	3072	2500	3518	2261
44	level-4		3	0	0	1	1
49	level-9		100	100	100	100	100
4d	level-13		23	8	14	19	14
60	int-vec		2559	2699	2752	2688	2792
64	itlb-miss		3296	1548	1174	1698	1087
68	dtlb-miss		114788	54313	43040	58336	38057
6c	dtlb-prot		1046	549	417	545	370
84	spill-user-32		66551	29480	301588	26522	213032
88	spill-user-64		0	318652	111239	299829	221716
8c	spill-user-32-cln		856	347	331	416	293
90	spill-user-64-cln		0	55	21	59	39
98	spill-kern-64		66464	31803	24758	34004	22277
a4	spill-asuser-32		1423	569	560	698	483
a8	spill-asuser-64		0	74	32	98	46
ac	spill-asuser-32-cln		4875	2250	1728	2384	1584
b0	spill-asuser-64-cln		0	2	0	1	0
c4	fill-user-32		64193	28418	287516	27055	202093
c8	fill-user-64		0	305016	106692	288542	210654
cc	fill-user-32-cln		6733	3520	15185	2396	12035
d0	fill-user-64-cln		0	13226	3506	12933	11032
d8	fill-kern-64		66220	31680	24674	33892	22196
108	syscall-32		2446	967	817	1196	755

示例 3 Using trapstat with TLB Statistics

The -t option displays in-depth TLB statistics, including the amount of time spent performing TLB miss processing. The following example shows that the machine is spending 14.1 percent of its time just handling D-TLB misses:

```
example# trapstat -t
cpu m| itlb-miss %tim itsb-miss %tim | dtlb-miss %tim dtsb-miss %tim |%tim
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
  0 u|      2571  0.3          0  0.0 |    10802  1.3          0  0.0 |  1.6
  0 k|          0  0.0          0  0.0 |   106420 13.4        184  0.1 |13.6
-----+-----+-----+-----+-----+-----+-----+
  1 u|      3069  0.3          0  0.0 |    10983  1.2        100  0.0 |  1.6
  1 k|         27  0.0          0  0.0 |   106974 12.6         19  0.0 |12.7
-----+-----+-----+-----+-----+-----+-----+
  2 u|      3033  0.3          0  0.0 |    11045  1.2        105  0.0 |  1.6
  2 k|         43  0.0          0  0.0 |   107842 12.7        108  0.0 |12.8
-----+-----+-----+-----+-----+-----+-----+
  3 u|      2924  0.3          0  0.0 |    10380  1.2        121  0.0 |  1.6
  3 k|         54  0.0          0  0.0 |   102682 12.2         16  0.0 |12.2
-----+-----+-----+-----+-----+-----+-----+
  4 u|      3064  0.3          0  0.0 |    10832  1.2        120  0.0 |  1.6
  4 k|         31  0.0          0  0.0 |   107977 13.0        236  0.1 |13.1
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
ttl |      14816  0.3          0  0.0 |   585937 14.1       1009  0.0 |14.5
```

示例 4 Using trapstat with TLB Statistics and Page Size Information

By specifying the -T option, trapstat shows TLB misses broken down by page size. In this example, CPU 0 is spending 7.9 percent of its time handling user-mode TLB misses on 8K pages, and another 2.3 percent of its time handling user-mode TLB misses on 64K pages.

```
example# trapstat -T -c 0
cpu m size| itlb-miss %tim itsb-miss %tim | dtlb-miss %tim dtsb-miss %tim |%tim
-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
  0 u  8k|      1300  0.1        15  0.0 |   104897  7.9         90  0.0 |  8.0
  0 u 64k|          0  0.0          0  0.0 |    29935  2.3          7  0.0 |  2.3
  0 u512k|          0  0.0          0  0.0 |     3569  0.2          2  0.0 |  0.2
  0 u  4m|          0  0.0          0  0.0 |        233  0.0          2  0.0 |  0.0
-----+-----+-----+-----+-----+-----+-----+-----+
  0 k  8k|         13  0.0          0  0.0 |    71733  6.5        110  0.0 |  6.5
  0 k 64k|          0  0.0          0  0.0 |          0  0.0          0  0.0 |  0.0
  0 k512k|          0  0.0          0  0.0 |          0  0.0        206  0.1 |  0.1
  0 k  4m|          0  0.0          0  0.0 |          0  0.0          0  0.0 |  0.0
=====+=====+=====+=====+=====+=====+=====+=====+=====+=====
          ttl |      1313  0.1         15  0.0 |   210367 17.1        417  0.2 |17.5
```

示例 5 Using trapstat with Entry Filtering

By specifying the `-e` option, `trapstat` displays statistics for only specific trap types. Using this option minimizes the probe effect when seeking specific data. This example yields statistics for only the `dtlb-prot` and `syscall-32` traps on CPUs 12 through 15:

```
example# trapstat -e dtlb-prot,syscall-32 -c 12-15
```

vct	name		cpu12	cpu13	cpu14	cpu15
6c	dtlb-prot		817	754	1018	560
108	syscall-32		1426	1647	2186	1142

vct	name		cpu12	cpu13	cpu14	cpu15
6c	dtlb-prot		1085	996	800	707
108	syscall-32		2578	2167	1638	1452

示例 6 Using trapstat with a Higher Sampling Rate

The following example uses the `-r` option to specify a sampling rate of 1000 samples per second, and filter only for the level-10 trap. Additionally, specifying the `-P` option yields parsable output.

Notice the timestamp difference between the level-10 events: 9,998,000 nanoseconds and 10,007,000 nanoseconds. These level-10 events correspond to the system clock, which by default ticks at 100 hertz (that is, every 10,000,000 nanoseconds).

```
example# trapstat -e level-10 -P -r 1000
1070400 0 4a level-10 0
2048600 0 4a level-10 0
3030400 0 4a level-10 1
4035800 0 4a level-10 0
5027200 0 4a level-10 0
6027200 0 4a level-10 0
7027400 0 4a level-10 0
8028200 0 4a level-10 0
9026400 0 4a level-10 0
10029600 0 4a level-10 0
11028600 0 4a level-10 0
12024000 0 4a level-10 0
13028400 0 4a level-10 1
14031200 0 4a level-10 0
15027200 0 4a level-10 0
16027600 0 4a level-10 0
17025000 0 4a level-10 0
18026000 0 4a level-10 0
19027800 0 4a level-10 0
20025600 0 4a level-10 0
21025200 0 4a level-10 0
22025000 0 4a level-10 0
```

示例 6 Using trapstat with a Higher Sampling Rate (续)

```
23035400 0 4a level-10 1
24027400 0 4a level-10 0
25026000 0 4a level-10 0
26027000 0 4a level-10 0
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	
Human Readable Output	Uncommitted
Parsable Output	Committed

另请参见

[lockstat\(1M\)](#), [pmap\(1\)](#), [psrset\(1M\)](#), [psrinfo\(1M\)](#), [pbind\(1M\)](#), [ppgsz\(1\)](#), [getpagesizes\(3C\)](#)

Sun Microelectronics UltraSPARC II User's Manual, January 1997, STP1031,
The SPARC Architecture Manual, Version 9, 1994, Prentice-Hall.

附注

When enabled, `trapstat` induces a varying probe effect, depending on the type of information collected. While the precise probe effect depends upon the specifics of the hardware, the following table can be used as a rough guide:

Option	Approximate probe effect
default	3-5% per trap
-e	3-5% per specified trap
-t, -T	40-45% per TLB miss trap hitting in the TSB, 25-30% per TLB miss trap missing in the TSB

These probe effects are *per trap* not for the system as a whole. For example, running `trapstat` with the default options on a system that spends 7% of total time handling traps induces a performance degradation of less than one half of one percent; running `trapstat` with the `-t` or `-T` option on a system spending 5% of total time processing TLB misses induce a performance degradation of no more than 2.5%.

When run with the `-t` or `-T` option, `trapstat` accounts for its probe effect when calculating the `%tim` fields. This assures that the `%tim` fields are a reasonably accurate indicator of the time a given workload is spending handling TLB misses — regardless of the perturbing presence of `trapstat`.

While the *%tim* fields include the explicit cost of executing the TLB miss handler, they do *not* include the implicit costs of TLB miss traps (for example, pipeline effects, cache pollution, etc). These implicit costs become more significant as the trap rate grows; if high *%tim* values are reported (greater than 50%), you can accurately infer that much of the balance of time is being spent on the implicit costs of the TLB miss traps.

Due to the potential system wide degradation induced, only the super-user can run `trapstat`.

Due to the limitation of the underlying statistics gathering methodology, only one instance of `trapstat` can run at a time.

引用名

ttymon – port monitor for terminal ports

用法概要

```
/usr/lib/ttymon -g [-d device] [-h] [-t timeout]
                  [-l ttylabel] [-p prompt] [-m modules] [-T termtype]
```

描述

ttymon is a STREAMS-based TTY port monitor. Its function is to monitor ports, to set terminal modes, baud rates, and line disciplines for the ports, and to connect users or applications to services associated with the ports. Each instance of ttymon monitors one port, specified at startup. When an instance of ttymon is started, ttymon first initializes the line disciplines, if they are specified, and the speed and terminal settings. For ports with entries in /etc/logindevperm, device owner, group and permissions are set. (See [logindevperm\(4\)](#).) The values used for initialization are taken from the appropriate entry in the TTY settings file. This file is maintained by the [sttydefs\(1M\)](#) command. Default line disciplines on ports are usually set up by the [autopush\(1M\)](#) command of the Autopush Facility.

ttymon then writes the prompt and waits for user input. If the user indicates that the speed is inappropriate by pressing the BREAK key, ttymon tries the next speed and writes the prompt again. When valid input is received, ttymon creates a utmpx entry (see [utmpx\(4\)](#)), and execs the login service for the port. Valid input consists of a string of at least one non-newline character, terminated by a carriage return.

If *autobaud* is enabled for a port, ttymon will try to determine the baud rate on the port automatically. Users must enter a carriage return before ttymon can recognize the baud rate and print the prompt. Currently, the baud rates that can be determined by *autobaud* are 110, 1200, 2400, 4800, and 9600.

SMF Service Description

The primary [smf\(5\)](#) service which invokes ttymon is `svc:/system/console-login`, which may have multiple service instances. Instances are described in greater detail below. The service provides a number of properties within the property group ttymon to control the invocation, as follows:

NAME	TYPE	TTYMON OPTION
device	cstring	[-d device]
nohangup	boolean	[-h]
label	cstring	[-l label]
modules	cstring	[-m module1,module2]
prompt	cstring	[-p prompt]
timeout	count	[-t timeout]
terminal_type	cstring	[-T termtype]

If any value is the empty string or an integer set to zero, then the option is not passed to the ttymon invocation.

`svc:/system/console-login:default`
The default instance always represents the ttymon that offers login on the system hardware console.

See EXAMPLES for an example of how to modify settings for the system console.


```
svc:/system/console-login:{vt2, vt3, vt4, vt5, vt6}
```

Additional service instances are provided for the system's virtual consoles. If virtual consoles are not available, these services will automatically disable themselves. See [vtdaemon\(1M\)](#).

```
svc:/system/console-login:{terma, termb}
```

svc:/system/console-login:terma and svc:/system/console-login:termb are provided as a convenience and can assist the user in setting up login services for additional ports /dev/term/a and /dev/term/b. These services are disabled by default.

Creating Additional Instances

The user can configure additional service instances for additional devices. This can be accomplished in any of these ways:

- Manually creating the service instance using [svccfg\(1M\)](#).
- Creating the service in a service profile (see [smf\(5\)](#)).
- Creating a service manifest for additional service instance(s).

See EXAMPLES for an example of manually configuring the service using [svccfg](#).

SMF Service Errors

In most cases when an instance of the console-login service is misconfigured, it will transition itself to the maintenance state. Use `svcs -l` (see [svcs\(1\)](#)) to determine the location of the service's log file and consult the log for additional information.

In some error cases, the service may respawn indefinitely. Disable the service using [svcadm\(1M\)](#), then consult the service log for additional messages or information to help resolve the problem.

安全

ttymon uses [pam\(3PAM\)](#) for session management. The PAM configuration policy, listed through /etc/pam.conf, specifies the modules to be used for ttymon. Here is a partial pam.conf file with entries for ttymon using the UNIX session management module.

```
ttymon session required /usr/lib/security/pam_unix_session.so.1
```

If there are no entries for the ttymon service, then the entries for the “other” service will be used.

选项

The following options are supported:

- g The -g option is required for historical reasons.
- dddevice *device* is the full path name of the port to which ttymon is to attach. If this option is not specified, file descriptor 0 must be set up by the invoking process to a TTY port.
- h If the -h flag is not set, ttymon will force a hangup on the line by setting the speed to zero before setting the speed to the default or specified speed.
- l`ttylabel` *ttylabel* is a link to a speed and TTY definition in the `ttdefs` file. This definition tells ttymon at what speed to run initially, what the initial TTY

- settings are, and what speed to try next if the user indicates that the speed is inappropriate by pressing the BREAK key. The default speed is 9600 baud.
- mmodules* When initializing the port, `ttymon` will pop all modules on the port, and then push *modules* in the order specified. *modules* is a comma-separated list of pushable modules. Default modules on the ports are usually set up by the Autopush Facility.
 - pprompt* Allows the user to specify a prompt string. The default prompt is `Login:.`
 - ttimeout* Specifies that `ttymon` should exit if no one types anything in *timeout* seconds after the prompt is sent.
 - Ttermtype* Sets the TERM environment variable to *termtype*.
 - v* Enables verbose messaging.

示例

示例 1 Setting the Terminal Type for the System Console

The following example sets the value of the terminal type (`-T`) option for the system console `ttymon` invocation:

```
# svccfg -s svc:/system/console-login:default \
    "setprop ttymon/terminal_type = xterm"
# svcadm refresh svc:/system/console-login:default
```

示例 2 Creating a Service Instance for an Additional Serial Device

In this example, the user wishes to configure an additional instance of the `svc:/system/console-login` service in order to offer login services over a terminal connected by means of a USB serial adapter. Assume that the USB serial port is present as `/dev/term/1`, and the user plans to connect a `vt100` terminal to it. In this case, the service instance can be named `term1` (or any other name) and defined as follows:

```
# svccfg -s svc:/system/console-login "add term1"
# SVC=svc:/system/console-login:term1
# svccfg -s $SVC "addpg ttymon application"
# svccfg -s $SVC "setprop ttymon/device = /dev/term/1"
# svccfg -s $SVC "setprop ttymon/terminal_type = vt100"
# svcadm refresh $SVC
# svcadm enable $SVC
```

环境变量

If any of the `LC_*` variables (`LC_CTYPE`, `LC_MESSAGES`, `LC_TIME`, `LC_COLLATE`, `LC_NUMERIC`, and `LC_MONETARY`) (see [environ\(5\)](#)) are not set in the environment, the operational behavior of `ttymon` for each corresponding locale category is determined by the value of the `LANG` environment variable. If `LC_ALL` is set, its contents are used to override both the `LANG` and the other `LC_*` variables. If none of the above variables is set in the environment, the “C” (U.S. style) locale determines how `ttymon` behaves.

LC_CTYPE Determines how ttymon handles characters. When LC_CTYPE is set to a valid value, ttymon can display and handle text and filenames containing valid characters for that locale. ttymon can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. ttymon can also handle EUC characters of 1, 2, or more column widths. In the “C” locale, only characters from ISO 8859-1 are valid.

文件 /etc/logindevperm
Contains information that is used by login(1) and ttymon to change the owner, group, and permissions of devices upon logging into or out of a console device.

属性 See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [svcs\(1\)](#), [ct\(1C\)](#), [cu\(1C\)](#), [autopush\(1M\)](#), [sttydefs\(1M\)](#), [svcadm\(1M\)](#), [svccfg\(1M\)](#), [uucico\(1M\)](#), [vtdaemon\(1M\)](#), [pam\(3PAM\)](#), [logindevperm\(4\)](#), [pam.conf\(4\)](#), [utmpx\(4\)](#), [attributes\(5\)](#), [environ\(5\)](#), [pam_authtok_check\(5\)](#), [pam_authtok_get\(5\)](#), [pam_authtok_store\(5\)](#), [pam_dhkeys\(5\)](#), [pam_passwd_auth\(5\)](#), [pam_unix_account\(5\)](#), [pam_unix_auth\(5\)](#), [pam_unix_session\(5\)](#), [smf\(5\)](#)

《Oracle Solaris 管理：常见任务》

附注

Service Access Facility (SAF and SAC) ttymon was formerly a component of the Service Access Facility and was invoked by sac, the Service Access Controller. This facility has been removed in this release of Solaris, and a conversion to SMF of relevant portions was performed.

Competition for Ports If a port is monitored by more than one ttymon, it is possible for the ttymons to send out prompt messages in such a way that they compete for input.

It is possible that two svc:/system/console-login service instances could refer to the same underlying device. For example, if the system's hardware console is connected (due to settings or autodetection in firmware) to serial port A, then both the svc:/system/console-login:default and svc:/system/console-login:terma services will refer to same underlying hardware device. Care should be taken when defining or enabling additional service instances to avoid this situation, or the two ttymons will compete for input.

引用名	tunefs – tune an existing UFS file system
用法概要	<pre>tunefs [-a maxcontig] [-d rotdelay] [-e maxbpg] [-m minfree] [-o space time] special filesystem</pre>
描述	tunefs is designed to change the dynamic parameters of a file system that affect the layout policies. When using tunefs with <i>filesystem</i>, <i>filesystem</i> must be in <i>/etc/vfstab</i>. The parameters that can be changed are indicated by the options given below.
选项	The following options are supported: -a maxcontig The maximum number of logical blocks, belonging to one file, that is allocated contiguously. The default is calculated as follows: $\text{maxcontig} = \text{disk drive maximum transfer size} / \text{disk block size}$ If the disk drive's maximum transfer size cannot be determined, the default value for maxcontig is calculated from kernel parameters as follows: If maxphys is less than ufs_maxmaxphys, which is 1 Mbyte, then maxcontig is set to maxphys. Otherwise, maxcontig is set to ufs_maxmaxphys. You can set maxcontig to any positive integer value. The actual value will be the lesser of what has been specified and what the hardware supports. -d rotdelay This parameter is obsolete as of the Solaris 10 release. The value is always set to 0, regardless of the input value. -e maxbpg Indicates the maximum number of contiguous logical blocks any single file can allocate from a cylinder group before it is forced to begin allocating blocks from another cylinder group. Typically this value is set to approximately one quarter of the total contiguous logical blocks in a cylinder group. The intent is to prevent any single file from using up all the blocks in a single cylinder group, thus degrading access times for all files subsequently allocated in that cylinder group. The effect of this limit is to cause big files to do long seeks more frequently than if they were allowed to allocate all the blocks in a cylinder group before seeking elsewhere. For file systems with exclusively large files, this parameter should be set higher. -m minfree Specifies the minimum free space threshold, or the percentage of space held back from normal users. This value can be set to 0. However, up to a factor of three in throughput will be lost over the performance obtained at a 10% threshold. <i>Note:</i> If the value is raised above the current usage level, users will be unable to allocate files until enough files have been deleted to get under the higher threshold.

`-o space | time` The file system can either be instructed to try to minimize the *time* spent allocating blocks, or to try to minimize the *space* fragmentation on the disk. The default is *time*.

Generally, you should optimize for time unless the file system is over 90% full.

用法 See [largefile\(5\)](#) for the description of the behavior of tunefs when encountering files greater than or equal to 2 Gbyte (2³¹ bytes).

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [mkfs_ufs\(1M\)](#), [newfs\(1M\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

引用名	txzonemgr – Trusted Extensions 区域管理器配置实用程序
用法概要	<pre>/usr/sbin/txzonemgr [-c -d[f] -Z file]</pre>
描述	txzonemgr shell 脚本提供了简单的基于菜单的 GUI 向导，在启用了 Trusted Extensions 的系统上创建、安装、初始化和引导有标签区域。此脚本还提供用于网络选项、名称服务选项的菜单项，以及用于使全局区域成为现有 LDAP 服务器的客户机的菜单项。缺省情况下，所有区域都配置为使用与全局区域相同的名称服务和 IP 地址。 txzonemgr 还可以在命令行模式（而非 GUI 模式）下使用下述的合适选项执行一组有限的命令。 txzonemgr 由在 "Zone Management"（区域管理）权限配置文件中授予的角色或全局区域中的 root 运行。
选项	没有指定选项时，txzonemgr 作为基于菜单的 GUI 运行。使用任何选项都会导致命令行运行。除了下面指定的选项以外，其他选项不能组合使用。 所支持的选项如下： -c 创建缺省区域。要求预先不存在任何区域。创建的区域的编号、名称和属性可能会根据情况而变。但是，通常会创建一个 "public"（公共）区域和至少一个其他区域。 -d 销毁所有区域。将停止、卸载并删除所有区域。只有同时指定了 -f 选项时才提示您进行确认。 -f 强制。可以与 -d 一起使用来覆盖确认提示。 -Z <i>file</i> 使用文件中指定的列表创建一组区域。文件中的每行都必须包含区域名称和其相应的标签（以空格分隔）。即使标签包含空格，也不应当用引号括起。如果存在名为“快照”的现有区域，则从其中克隆所有指定的区域。否则将安装列表中的第一个区域，然后从中克隆快照区域。接下来从快照克隆剩下的区域。这些区域配置为需要手动引导。但是，当授权用户向 GNOME 工作区分配有效标签时，将自动引导相应区域。
退出状态	对于 GUI 模式，不返回任何值。对于命令行运行，返回以下退出值： 0 成功完成。 1 出现错误。 2 用法无效。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
命令接口稳定性	Committed（已确定）
交互对话	Not an Interface（不是接口）

另请参见 [zenity\(1\)](#)（非 SunOS 手册页）、[ipadm\(1M\)](#)、[zoneadm\(1M\)](#)、[zonecfg\(1M\)](#)、[attributes\(5\)](#)、[rbac\(5\)](#)、[zones\(5\)](#)
《Solaris Trusted Extensions 管理员规程》

附注 如果是从 JDS 管理区域，请使用 txzonemgr 而不是 CDE 操作。

txzonemgr 使用 zenity 命令。有关详细信息，请参见 [zenity\(1\)](#) 手册页（不是 SunOS 集合的一部分）。

引用名 tzreload – notify timezone update

用法概要 /usr/sbin/tzreload

描述 The tzreload command notifies active (running) processes to reread timezone information. The timezone information is cached in each process, absent a tzreload command, is never reread until a process is restarted. In response to a tzreload command, active processes reread the current timezone information at the next call to [ctime\(3C\)](#) and [mktime\(3C\)](#). The tzreload notification is sent to processes within the current zone.

tzreload causes processes which are using the system timezone (in /etc/default/init) to reread the contents of that file.

In addition to notifying active processes, the tzreload command also notifies [cron\(1M\)](#), to reinitialize the job scheduler with the new timezone information.

文件 /usr/share/lib/zoneinfo
Standard zone information directory.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [cron\(1M\)](#), [zdump\(1M\)](#), [zic\(1M\)](#), [ctime\(3C\)](#), [mktime\(3C\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 Although tzreload reinitializes [cron\(1M\)](#), applications that are affected by timezone changes still need to be restarted or reinitialized if they do not reread the new timezone information before timezone changes take place.

The timezone update service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

svc:/system/timezone:default

Administrative actions on this service, such as enabling, disabling, or requesting refresh, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

Refreshing this service causes the tzreload command to be run, notifying running processes to reload their timezone caches. Disabling this service is not recommended.

引用名	tzselect – select a time zone
用法概要	/usr/bin/tzselect
描述	The tzselect program asks you a series of questions about the current location and outputs the resulting time zone description to standard output. The output is suitable as a value for the TZ environment variable. All user interaction is through standard input and standard error.
选项	The tzselect command has no options.
退出状态	The following exit values are returned: 0 Timezone information was successfully obtained. >0 An error occurred.
文件	/usr/share/lib/zoneinfo directory containing timezone data files
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [zdump\(1M\)](#), [zic\(1M\)](#), [ctime\(3C\)](#), [attributes\(5\)](#)

引用名 uadmin – administrative control

用法概要 /usr/sbin/uadmin *cmd fcn* [*mdep*]

描述 The uadmin command provides control for basic administrative functions. This command is tightly coupled to the system administration procedures and is not intended for general use. It may be invoked only by the super-user.

Both the *cmd* (command) and *fcn* (function) arguments are converted to integers and passed to the uadmin system call. The optional *mdep* (machine dependent) argument is only available for the *cmd* values of 1 (A_REBOOT), 2 (A_SHUTDOWN), or 5 (A_DUMP). For any other *cmd* value, no *mdep* command-line argument is allowed.

When passing an *mdep* value that contains whitespaces, the string must be grouped together as a single argument enclosed within quotes, for example:

```
uadmin 1 1 "-s kernel/unix"
```

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [halt\(1M\)](#), [reboot\(1M\)](#), [uadmin\(2\)](#), [attributes\(5\)](#)

警告 Shutting down or halting the system by means of uadmin does not update the boot archive. Avoid using this command after

- editing of files such as /etc/system
- installing new driver binaries or kernel binaries
- updating existing driver binaries or kernel binaries.

Use [reboot\(1M\)](#) or [halt\(1M\)](#) instead.

引用名	ucodeadm – update processor microcode								
用法概要	<pre>/usr/sbin/ucodeadm -v /usr/sbin/ucodeadm -u <i>microcode-text-file</i> /usr/sbin/ucodeadm -i [-R <i>path</i>] <i>microcode-text-file</i></pre>								
描述	The ucodeadm utility can be used to report running microcode revision on the processors, update microcode, or install microcode on the target system to be used during the boot process. The <i>microcode-text-file</i> can be obtained from processor vendors.								
选项	<table><tr><td>-v</td><td>Report microcode revision.</td></tr><tr><td>-u <i>microcode-text-file</i></td><td>Update microcode on all cross-call interrupt ready processors.</td></tr><tr><td>-i <i>microcode-text-file</i></td><td>Install microcode files on target system to be used during the next boot cycle. The text file name must have the vendor name prefix, such as “intel” or “amd”.</td></tr></table> By default the microcode files will be installed at: <pre>/platform/i86pc/ucode/\$VENDORSTR/</pre> where VENDORSTR is either “GenuineIntel” or “AuthenticAMD”. <table><tr><td>-R <i>alternate path</i></td><td>Install <i>microcode</i> path in the <i>alternate path</i>.</td></tr></table>	-v	Report microcode revision.	-u <i>microcode-text-file</i>	Update microcode on all cross-call interrupt ready processors.	-i <i>microcode-text-file</i>	Install microcode files on target system to be used during the next boot cycle. The text file name must have the vendor name prefix, such as “intel” or “amd”.	-R <i>alternate path</i>	Install <i>microcode</i> path in the <i>alternate path</i> .
-v	Report microcode revision.								
-u <i>microcode-text-file</i>	Update microcode on all cross-call interrupt ready processors.								
-i <i>microcode-text-file</i>	Install microcode files on target system to be used during the next boot cycle. The text file name must have the vendor name prefix, such as “intel” or “amd”.								
-R <i>alternate path</i>	Install <i>microcode</i> path in the <i>alternate path</i> .								
示例	示例 1 Reporting the Microcode Revision The following example displays the microcode revision that is currently running: <pre># ucodeadm -v</pre> 示例 2 Updating the Processor Microcode The following example updates the processor microcode to intel-ucode.txt: <pre># ucodeadm -u intel-ucode.txt</pre> 示例 3 Installing the Microcode on the Target System The following example installs the microcode on the target system, /export/ucode-path: <pre># ucodeadm -i -R /export/ucode-path intel-ucode.txt</pre> If an alternate path is used when installing the microcode on the target system, the installed microcode file is not used on the next boot cycle.								
退出状态	The following exit values are returned: <table><tr><td>0</td><td>Successful completion.</td></tr></table>	0	Successful completion.						
0	Successful completion.								

>0 An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [psradm\(1M\)](#), [psrinfo\(1M\)](#), [attributes\(5\)](#)

引用名	ufsdump – incremental file system dump
用法概要	<code>/usr/sbin/ufsdump [options] [arguments] files_to_dump</code>
描述	ufsdump backs up all files specified by <i>files_to_dump</i> (usually either a whole file system or files within a file sytem changed after a certain date) to magnetic tape or disk file. The ufsdump command can only be used on unmounted file systems, or those mounted read-only. Attempting to dump a mounted, read-write file system might result in a system disruption or the inability to restore files from the dump. Consider using the fssnap(1M) command to create a file system snapshot if you need a point-in-time image of a file system that is mounted. If a filesystem was mounted with the logging option, it is strongly recommended that you run ufsdump as the root user. Running the command as a non-root user might result in the creation of an inconsistent dump. <i>options</i> is a single string of one-letter ufsdump options. <i>arguments</i> may be multiple strings whose association with the options is determined by order. That is, the first argument goes with the first option that takes an argument; the second argument goes with the second option that takes an argument, and so on. <i>files_to_dump</i> is required and must be the last argument on the command line. See OPERANDS for more information. With most devices ufsdump can automatically detect the end-of-media. Consequently, the <i>d</i>, <i>s</i>, and <i>t</i> options are not necessary for multi-volume dumps, unless ufsdump does not understand the way the device detects the end-of-media, or the files are to be restored on a system with an older version of the <i>restore</i> command.
选项	The following options are supported: 0–9 The “dump level.” All files specified by <i>files_to_dump</i> that have been modified since the last ufsdump at a lower dump level are copied to the <i>dump_file</i> destination (normally a magnetic tape device). For instance, if a “level 2” dump was done on Monday, followed by a “level 4” dump on Tuesday, a subsequent “level 3” dump on Wednesday would contain all files modified or added since the “level 2” (Monday) backup. A “level 0” dump copies the entire file system to the <i>dump_file</i>. a <i>archive_file</i> Archive file. Archive a dump table-of-contents in the specified <i>archive_file</i> to be used by ufsrestore(1M) to determine whether a file is in the dump file that is being restored. b <i>factor</i> Blocking factor. Specify the blocking factor for tape writes. The default is 20 blocks per write for tapes of density less than 6250BPI (bytes-per-inch). The default blocking factor for tapes of density 6250BPI and greater is 64. The default blocking factor for cartridge

tapes (`c` option) is 126. The highest blocking factor available with most tape drives is 126. Note: the blocking factor is specified in terms of 512-byte blocks, for compatibility with [tar\(1\)](#).

c

Cartridge. Set the defaults for cartridge instead of the standard half-inch reel. This sets the density to 1000BPI and the blocking factor to 126. Since `ufsdump` can automatically detect the end-of-media, only the blocking parameter normally has an effect. When cartridge tapes are used, and this option is *not* specified, `ufsdump` will slightly miscalculate the size of the tape. If the `b`, `d`, `s` or `t` options are specified with this option, their values will override the defaults set by this option.

d bpi

Tape density. Not normally required, as `ufsdump` can detect end-of-media. This parameter can be used to keep a running tab on the amount of tape used per reel. The default density is 6250BPI except when the `c` option is used for cartridge tape, in which case it is assumed to be 1000BPI per track. Typical values for tape devices are:

1/2 inch tape

6250 BPI

1/4 inch cartridge

1000 BPI The tape densities and other options are documented in the [st\(7D\)](#) man page.

D

Diskette. Obsolete option.

f dump_file

Dump file. Use *dump_file* as the file to dump to, instead of `/dev/rmt/0`. If *dump_file* is specified as `-`, dump to standard output.

If the name of the file is of the form *machine:device*, the dump is done from the specified machine over the network using [rmt\(1M\)](#). Since `ufsdump` is normally run by root, the name of the local machine must appear in the `/etc/hosts` file of the remote machine. If the file is specified as *user@machine:device*, `ufsdump` will attempt to execute as the specified user on the remote machine. The specified user must have a `/etc/hosts` file on the remote machine that allows the user invoking the command from the local machine to access the remote machine.

l

Autoload. When the end-of-tape is reached before the dump is complete, take the drive offline and wait up to two minutes for the tape drive to be ready again. This gives autoloading (stackloader) tape drives a chance to load a new tape. If the drive is ready within two minutes, continue. If it is not, prompt for another tape and wait.

L *string*

Sets the tape label to *string*, instead of the default none. *string* may be no more than sixteen characters long. If it is longer, it is truncated and a warning printed; the dump will still be done. The tape label is specific to the ufsdump tape format, and bears no resemblance to IBM or ANSI-standard tape labels.

n

Notify all operators in the sys group that ufsdump requires attention by sending messages to their terminals, in a manner similar to that used by the [wall\(1M\)](#) command. Otherwise, such messages are sent only to the terminals (such as the console) on which the user running ufsdump is logged in.

N *device_name*

Use *device_name* when recording information in /etc/dumpdates (see the u option) and when comparing against information in /etc/dumpdates for incremental dumps. The *device_name* provided can contain no white space as defined in [scanf\(3C\)](#) and is case-sensitive.

o

Offline. Take the drive offline when the dump is complete or the end-of-media is reached and rewind the tape. In the case of some autoloading 8mm drives, the tape is removed from the drive automatically. This prevents another process which rushes in to use the drive, from inadvertently overwriting the media.

s *size*

Specify the *size* of the volume being dumped to. Not normally required, as ufsdump can detect end-of-media. When the specified size is reached, ufsdump waits for you to change the volume. ufsdump interprets the specified size as the length in feet for tapes and cartridges. The values should be a little smaller than the actual physical size of the media (for example, 425 for a 450-foot cartridge). Typical values for tape devices depend on the c option for cartridge devices:

1/2 inch tape

2300 feet

60-Mbyte 1/4 inch cartridge

425 feet

150-Mbyte 1/4 inch cartridge

700 feet

S

Size estimate. Determine the amount of space that is needed to perform the dump without actually doing it, and display the estimated number of bytes it will take. This is useful with incremental dumps to determine how many volumes of media will be needed.

t *tracks*

Specify the number of tracks for a cartridge tape. Not normally required, as `ufsdump` can detect end-of-media. The default is 9 tracks. The `t` option is not compatible with the `D` option. Values for Sun-supported tape devices are:

60-Mbyte 1/4 inch cartridge

9 tracks

150-Mbyte 1/4 inch cartridge

18 tracks

T *time_wait[hms]*

Sets the amount of time to wait for an `autoload` command to complete. This option is ignored unless the `l` option has also been specified. The default time period to wait is two minutes. Specify time units with a trailing `h` (for hours), `m` (for minutes), or `s` (for seconds). The default unit is minutes.

u

Update the dump record. Add an entry to the file `/etc/dumpdates`, for each file system successfully dumped that includes the file system name (or *device_name* as specified with the `N` option), date, and dump level.

v

Verify. After each tape is written, verify the contents of the media against the source file system. If any discrepancies occur, prompt for new media, then repeat the dump/verification process. The file system *must* be unmounted. This option cannot be used to verify a dump to standard output.

w

Warning. List the file systems that have not been backed up within a day. This information is gleaned from the files `/etc/dumpdates` and `/etc/vfstab`. When the `w` option is used, all other options are ignored. After reporting, `ufsdump` exits immediately.

W

Warning with highlight. Similar to the `w` option, except that the `W` option includes all file systems that appear in `/etc/dumpdates`, along with information about their most recent dump dates and levels. File systems that have not been backed up within a day are highlighted.

操作数

The following operand is supported:

files_to_dump

Specifies the files to dump. Usually it identifies a whole file system by its raw device name (for example, `/dev/rdisk/c0t3d0s6`). Incremental dumps (levels 1 to 9) of files changed after a certain date only apply to a whole file system. Alternatively, *files_to_dump* can identify individual files or directories. All named directories that may be examined by the user running `ufsdump`, as well as any explicitly-named files, are dumped. This dump is equivalent to a level 0 dump of the indicated portions of the filesystem, except that

`/etc/dumpdates` is not updated even if the `-u` option has been specified. In all cases, the files must be contained in the same file system, and the file system must be local to the system where `ufsdump` is being run.

`files_to_dump` is required and must be the last argument on the command line.

If no *options* are given, the default is `9uf /dev/rmt/0 files_to_dump`.

用法

See [largefile\(5\)](#) for the description of the behavior of `ufsdump` when encountering files greater than or equal to 2 Gbyte (2^{31} bytes).

示例

示例1 Using `ufsdump`

The following command makes a full dump of a root file system on `c0t3d0`, on a 150-MByte cartridge tape unit 0:

```
example# ufsdump 0cfu /dev/rmt/0 /dev/rdisk/c0t3d0s0
```

The following command makes and verifies an incremental dump at level 5 of the `usr` partition of `c0t3d0`, on a 1/2 inch reel tape unit 1, :

```
example# ufsdump 5fuv /dev/rmt/1 /dev/rdisk/c0t3d0s6
```

退出状态

While running, `ufsdump` emits many verbose messages. `ufsdump` returns the following exit values:

- 0
Normal exit.
- 1
Startup errors encountered.
- 3
Abort – no checkpoint attempted.

文件

```
/dev/rmt/0
    default unit to dump to

/etc/dumpdates
    dump date record

/etc/group
    to find group sys

/etc/hosts
    to gain access to remote system with drive

/etc/vfstab
    list of file systems
```

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [cpio\(1\)](#), [tar\(1\)](#), [dd\(1M\)](#), [devnm\(1M\)](#), [fssnap\(1M\)](#), [prtvto\(1M\)](#), [rmt\(1M\)](#), [shutdown\(1M\)](#), [ufsrestore\(1M\)](#), [volcopy\(1M\)](#), [wall\(1M\)](#), [scanf\(3C\)](#), [ufsdump\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#), [st\(7D\)](#)

附注

Read Errors Fewer than 32 read errors on the file system are ignored.

Process Per Reel Because each reel requires a new process, parent processes for reels that are already written hang around until the entire tape is written.

Operator Intervention [ufsdump](#) requires operator intervention on these conditions: end of volume, end of dump, volume write error, volume open error or disk read error (if there are more than a threshold of 32). In addition to alerting all operators implied by the `n` option, [ufsdump](#) interacts with the operator on [ufsdump](#)'s control terminal at times when [ufsdump](#) can no longer proceed, or if something is grossly wrong. All questions [ufsdump](#) poses *must* be answered by typing yes or no, as appropriate.

Since backing up a disk can involve a lot of time and effort, [ufsdump](#) checkpoints at the start of each volume. If writing that volume fails for some reason, [ufsdump](#) will, with operator permission, restart itself from the checkpoint after a defective volume has been replaced.

Suggested Dump Schedule It is vital to perform full, “level 0”, dumps at regular intervals. When performing a full dump, bring the machine down to single-user mode using [shutdown\(1M\)](#). While preparing for a full dump, it is a good idea to clean the tape drive and heads. Incremental dumps should be performed with the system running in single-user mode.

Incremental dumps allow for convenient backup and recovery of active files on a more frequent basis, with a minimum of media and time. However, there are some tradeoffs. First, the interval between backups should be kept to a minimum (once a day at least). To guard against data loss as a result of a media failure (a rare, but possible occurrence), capture active files on (at least) two sets of dump volumes. Another consideration is the desire to keep unnecessary duplication of files to a minimum to save both operator time and media storage. A third consideration is the ease with which a particular backed-up version of a file can be located and restored. The following four-week schedule offers a reasonable tradeoff between these goals.

	Sun	Mon	Tue	Wed	Thu	Fri
Week 1:	Full	5	5	5	5	3
Week 2:		5	5	5	5	3
Week 3:		5	5	5	5	3
Week 4:		5	5	5	5	3

Although the Tuesday through Friday incrementals contain “extra copies” of files from Monday, this scheme assures that any file modified during the week can be recovered from the previous day's incremental dump.

Process Priority of
ufsdump

ufsdump uses multiple processes to allow it to read from the disk and write to the media concurrently. Due to the way it synchronizes between these processes, any attempt to run dump with a nice (process priority) of ‘-5’ or better will likely make ufsdump run *slower* instead of faster.

Overlapping Partitions

Most disks contain one or more overlapping slices because slice 2 covers the entire disk. The other slices are of various sizes and usually do not overlap. For example, a common configuration places root on slice 0, swap on slice 1, /opt on slice 5 and /usr on slice 6.

It should be emphasized that ufsdump dumps one ufs file system at a time. Given the above scenario where slice 0 and slice 2 have the same starting offset, executing ufsdump on slice 2 with the intent of dumping the entire disk would instead dump only the root file system on slice 0. To dump the entire disk, the user must dump the file systems on each slice separately.

已知问题

The /etc/vfstab file does not allow the desired frequency of backup for file systems to be specified (as /etc/fstab did). Consequently, the w and W options assume file systems should be backed up daily, which limits the usefulness of these options.

引用名

ufsrestore – incremental file system restore

用法概要

```
/usr/sbin/ufsrestore i | r | R | t | x [abdcfhlmstvyLT]
    [archive_file] [factor] [dumpfile] [n] [label]
    [timeout] [filename]...
```

描述

The `ufsrestore` utility restores files from backup media created with the `ufsdump` command. `ufsrestore`'s actions are controlled by the *key* argument. The *key* is exactly one *function letter* (`i`, `r`, `R`, `t`, or `x`) and zero or more *function modifiers* (letters). The *key* string contains no SPACE characters. Function modifier arguments are listed on the command line in the same order as their corresponding function modifiers appear in the *key* string.

filename arguments which appear on the command line, or as arguments to an interactive command, are treated as shell glob patterns by the `x` and `t` functions; any files or directories matching the patterns are selected. The metacharacters `*`, `?`, and `[ ]` must be protected from the shell if they appear on the command line. There is no way to quote these metacharacters to explicitly match them in a filename.

The temporary files `rstdir*` and `rstmode*` are placed in `/tmp` by default. If the environment variable `TMPDIR` is defined with a non-empty value, that location is used instead of `/tmp`.

选项

Function Letters

You must specify one (and only one) of the function letters listed below. Note that `i`, `x`, and `r` are intended to restore files into an empty directory. The `R` function is intended for restoring into a populated directory.

- `i` Interactive. After reading in the directory information from the media, `ufsrestore` invokes a shell-like interface that allows you to browse through the dump file's directory hierarchy and select individual files to be extracted. Restoration has the same semantics as `x` (see below). See Interactive Commands, below, for a description of available commands.
- `r` Recursive. Starting with an empty directory and a level 0 dump, the `r` function recreates the filesystem relative to the current working directory, exactly as it appeared when the dump was made. Information used to restore incremental dumps on top of the full dump (for example, `restoresymtable`) is also included. Several `ufsrestore` runs are typical, one for each higher level of dump (0, 1, ..., 9). Files that were deleted between the level 0 and a subsequent incremental dump will not exist after the final restore. To completely restore a file system, use the `r` function restore the level 0 dump, and again for each incremental dump. Although this function letter is intended for a complete restore onto a new file system (one just created with `newfs(1M)`), if the file system contains files not on the backup media, they are preserved.
- `R` Resume restoring. If an `r`-mode `ufsrestore` was interrupted, this function prompts for the volume from which to resume restoring and continues the restoration from where it was left off. Otherwise identical to `r`.

- t

Table of contents. List each *filename* that appears on the media. If no *filename* argument is given, the root directory is listed. This results in a list of all files on the media, unless the *h* function modifier is in effect. The table of contents is taken from the media or from the specified archive file, when the *a* function modifier is used. The *a* function modifier is mutually exclusive with the *x* and *r* function letters.
- x

Extract the named files from the media. Files are restored to the same relative locations that they had in the original file system.

If the *filename* argument matches a directory whose contents were written onto the media, and the *h* modifier is not in effect, the directory is recursively extracted, relative to the current directory, which is expected to be empty. For each file, the owner, modification time, and mode are restored (if possible).

If you omit the *filename* argument or specify *.*, the root directory is extracted. This results in the entire tape being extracted, unless the *h* modifier is in effect. *.* With the *x* function, existing files are overwritten and *ufs restore* displays the names of the overwritten files. Overwriting a currently-running executable can have unfortunate consequences.

Use the *x* option to restore partial file system dumps, as they are (by definition) not entire file systems.

Function Modifiers	<i>a archive_file</i>	Read the table of contents from <i>archive_file</i> instead of the media. This function modifier can be used in combination with the <i>t</i> , <i>i</i> , or <i>x</i> function letters, making it possible to check whether files are on the media without having to mount the media. When used with the <i>x</i> and interactive (<i>i</i>) function letters, it prompts for the volume containing the file(s) before extracting them.
	<i>b factor</i>	Blocking factor. Specify the blocking factor for tape reads. For variable length SCSI tape devices, unless the data was written with the default blocking factor, a blocking factor at least as great as that used to write the tape must be used; otherwise, an error will be generated. Note that a tape block is 512 bytes. Refer to the man page for your specific tape driver for the maximum blocking factor.
	<i>c</i>	Convert the contents of the media in 4.1BSD format to the new <i>ufs</i> file system format.
	<i>d</i>	Debug. Turn on debugging output.
	<i>f dump_file</i>	Use <i>dump_file</i> instead of <i>/dev/rmt/0</i> as the file to restore from. Typically <i>dump_file</i> specifies a tape drive. If <i>dump_file</i> is specified as <i>'-'</i> , <i>ufs restore</i> reads from the standard input. This allows ufsdump(1M) and <i>ufs restore</i> to be used in a pipeline to copy a file system:

```
example# ufsdump 0f - /dev/rdisk/c0t0d0s7 \
| (cd /home;ufsrestore xf -)
```

If the name of the file is of the form *machine:device*, the restore is done from the specified machine over the network using [rmt\(1M\)](#). Since `ufsrestore` is normally run by root, the name of the local machine must appear in the `/etc/hosts` file of the remote machine. If the file is specified as *user@machine:device*, `ufsrestore` will attempt to execute as the specified user on the remote machine. The specified user must have a `/etc/hosts` file on the remote machine that allows the user invoking the command from the local machine to access the remote machine.

- h Extract or list the actual directory, rather than the files that it references. This prevents hierarchical restoration of complete subtrees from the tape.
- l Autoload. When the end-of-tape is reached before the restore is complete, take the drive off-line and wait up to two minutes (the default, see the `T` function modifier) for the tape drive to be ready again. This gives autoloading (stackloader) tape drives a chance to load a new tape. If the drive is ready within two minutes, continue. If it is not, prompt for another tape and wait.
- L *label* The label that should appear in the header of the dump file. If the labels do not match, `ufsrestore` issues a diagnostic and exits. The tape label is specific to the `ufsdump` tape format, and bears no resemblance to IBM or ANSI-standard tape labels.
- m Extract by inode numbers rather than by filename to avoid regenerating complete pathnames. Regardless of where the files are located in the dump hierarchy, they are restored into the current directory and renamed with their inode number. This is useful if only a few files are being extracted.
- o Offline. Take the drive off-line when the restore is complete or the end-of-media is reached and rewind the tape. In the case of some autoloading 8mm drives, the tape is removed from the drive automatically.
- s *n* Skip to the *n*th file when there are multiple dump files on the same tape. For example, the command:


```
example# ufsrestore xfs /dev/rmt/0hn 5
```

 would position you to the fifth file on the tape when reading volume 1 of the dump. If a dump extends over more than one volume, all volumes except the first are assumed to start at position 0, no matter what “*s n*” value is specified.

If “s *n*” is specified, the backup media must be at BOT (beginning of tape). Otherwise, the initial positioning to read the table of contents will fail, as it is performed by skipping the tape forward *n*-1 files rather than by using absolute positioning. This is because on some devices absolute positioning is very time consuming.

- T timeout [hms] Sets the amount of time to wait for an autoloading command to complete. This function modifier is ignored unless the l function modifier has also been specified. The default timeout period is two minutes. The time units may be specified as a trailing h (hours), m (minutes), or s (seconds). The default unit is minutes.
- v Verbose. ufsrestore displays the name and inode number of each file it restores, preceded by its file type.
- y Do not ask whether to abort the restore in the event of tape errors. ufsrestore tries to skip over the bad tape block(s) and continue as best it can.

- Interactive Commands ufsrestore enters interactive mode when invoked with the i function letters. Interactive commands are reminiscent of the shell. For those commands that accept an argument, the default is the current directory. The interactive options are:
- add [filename] Add the named file or directory to the list of files to extract. If a directory is specified, add that directory and its files (recursively) to the extraction list (unless the h modifier is in effect).
- cd directory Change to *directory* (within the dump file).
- delete [filename] Delete the current directory, or the named file or directory from the list of files to extract. If a directory is specified, delete that directory and all its descendents from the extraction list (unless the h modifier is in effect). The most expedient way to extract a majority of files from a directory is to add that directory to the extraction list, and then delete specific files to omit.
- extract Extract all files on the extraction list from the dump media. ufsrestore asks which volume the user wishes to mount. The fastest way to extract a small number of files is to start with the last volume and work toward the first. If “s *n*” is given on the command line, volume 1 will automatically be positioned to file *n* when it is read.
- help Display a summary of the available commands.
- ls [directory] List files in *directory* or the current directory, represented by a ‘.’ (period). Directories are appended with a ‘/’ (slash). Entries marked for extraction are prefixed with a ‘*’ (asterisk). If the verbose option is in effect, inode numbers are also listed.

marked [<i>directory</i>]	Like <code>ls</code> , except only files marked for extraction are listed.
pager	Toggle the pagination of the output from the <code>ls</code> and <code>marked</code> commands. The pager used is that defined by the <code>PAGER</code> environment variable, or <code>more(1)</code> if that <code>env</code> is not defined. The <code>PAGER</code> <code>env</code> may include white-space-separated arguments for the pagination program.
pwd	Print the full pathname of the current working directory.
quit	<code>ufsrestore</code> exits immediately, even if the extraction list is not empty.
setmodes	Prompts: set owner/mode for '.' (period). Type <code>y</code> for yes to set the mode (permissions, owner, times) of the current directory '.' (period) into which files are being restored equal to the mode of the root directory of the file system from which they were dumped. Normally, this is what you want when restoring a whole file system, or restoring individual files into the same locations from which they were dumped. Type <code>n</code> for no, to leave the mode of the current directory unchanged. Normally, this is what you want when restoring part of a dump to a directory other than the one from which the files were dumped.
setpager <i>command</i>	Sets the command to use for paginating output instead of the default or that inherited from the environment. The <i>command</i> string may include arguments in addition to the command itself.
verbose	Toggle the status of the <code>v</code> modifier. While <code>v</code> is in effect, the <code>ls</code> command lists the inode numbers of all entries, and <code>ufsrestore</code> displays information about each file as it is extracted.
what	Display the dump header on the media.

操作数 The following operands are supported.

<i>filename</i>	Specifies the pathname of files (or directories) to be restored to disk. Unless the <code>h</code> function modifier is also used, a directory name refers to the files it contains, and (recursively) its subdirectories and the files they contain. <i>filename</i> is associated with either the <code>x</code> or <code>t</code> function letters, and must come last.

用法 See [largefile\(5\)](#) for the description of the behavior of `ufsrestore` when encountering files greater than or equal to 2 Gbyte (2³¹ bytes).

退出状态 The following exit values are returned:

- 0 Successful completion.
- 1 An error occurred. Verbose messages are displayed.

环境变量 PAGER The command to use as a filter for paginating output. This can also be used to specify the options to be used. Default is [more\(1\)](#).

- 文件

TMPDIR

Selects the directory for temporary files. Defaults to /tmp if not defined in the environment.
- /dev/rmt/0

the default tape drive
- \$TMPDIR/rstdir*

file containing directories on the tape
- \$TMPDIR/rstmode*

owner, mode, and timestamps for directories
- ./restoresymtable

information passed between incremental restores
- 属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [more\(1\)](#), [mkfs\(1M\)](#), [mount\(1M\)](#), [rmt\(1M\)](#), [ufsdump\(1M\)](#), [ufsdump\(4\)](#), [attributes\(5\)](#), [largefile\(5\)](#)

诊断 `ufsrestore` complains about bad option characters.

Read errors result in complaints. If `y` has been specified, or the user responds `y`, `ufsrestore` will attempt to continue.

If the dump extends over more than one tape, `ufsrestore` asks the user to change tapes. If the `x` or `i` function letter has been specified, `ufsrestore` also asks which volume the user wishes to mount. If the `s` modifier has been specified, and volume 1 is mounted, it is automatically positioned to the indicated file.

There are numerous consistency checks that can be listed by `ufsrestore`. Most checks are self-explanatory or can “never happen”. Common errors are given below.

Converting to new file system format
A dump tape created from the old file system has been loaded. It is automatically converted to the new file system format.

filename: not found on tape
The specified file name was listed in the tape directory, but was not found on the tape. This is caused by tape read errors while looking for the file, using a dump tape created on an active file system, or restoring a partial dump with the `r` function.

expected next file *inumber*, got *inumber*
A file that was not listed in the directory showed up. This can occur when using a dump tape created on an active file system.

Incremental tape too low
When doing an incremental restore, a tape that was written before the previous incremental tape, or that has too low an incremental level has been loaded.

Incremental tape too high

When doing incremental restore, a tape that does not begin its coverage where the previous incremental tape left off, or one that has too high an incremental level has been loaded.

media read error: invalid argument

Blocking factor specified for read is smaller than the blocking factor used to write data.

Tape read error while restoring

Tape read error while skipping over inode inumber

Tape read error while trying to resynchronize

A tape read error has occurred

If a file name is specified, then its contents are probably partially wrong. If an inode is being skipped or the tape is trying to resynchronize, then no extracted files have been corrupted, though files may not be found on the tape.

resync ufsrestore, skipped *num*

After a tape read error, ufsrestore may have to resynchronize itself. This message lists the number of blocks that were skipped over.

Incorrect tape label. Expected 'foo', got 'bar'.

The L option was specified, and its value did not match what was recorded in the header of the dump file.

附注

ufsrestore can get confused when doing incremental restores from dump tapes that were made on active file systems.

A level 0 dump must be done after a full restore. Because ufsrestore runs in user mode, it has no control over inode allocation. This means that ufsrestore repositions the files, although it does not change their contents. Thus, a full dump must be done to get a new set of directories reflecting the new file positions, so that later incremental dumps will be correct.

引用名	unshare – make local resource unavailable for mounting by remote systems				
用法概要	unshare [-F <i>protocol</i>] [-a [-t] [<i>pathname</i> <i>sharename</i>]]				
描述	The unshare command unpublishes a file system share, which makes a shared local file system unavailable for file sharing <i>protocol</i> . If the -F <i>protocol</i> option is omitted, then the first file sharing protocol listed in the /etc/dfs/fstypes file used as the default sharing protocol.				
选项	-F <i>protocol</i> Identify the file sharing protocol. -a Unpublish all active shares. -t Temporarily unpublish a share. By default, all shares are permanently unpublished.				
文件	/etc/dfs/fstypes List of file sharing protocols. NFS is the default file-sharing protocol. /etc/dfs/sharetab Contains a table of local resources published by the share command.				
属性	See attributes(5) for descriptions of the following attributes:				
<table><tr><th>ATTRIBUTETYPE</th><th>ATTRIBUTEVALUE</th></tr><tr><td>Availability</td><td>SUNWcs</td></tr></table>		ATTRIBUTETYPE	ATTRIBUTEVALUE	Availability	SUNWcs
ATTRIBUTETYPE	ATTRIBUTEVALUE				
Availability	SUNWcs				
另请参见	share(1M) , shareall(1M) , attributes(5)				
附注	If <i>pathname</i> or <i>resourcename</i> is not found in the share information, an error message is sent to standard error. When an unshare command completes successfully, a client mounting a file system specified in that unshare command no longer has access to that file system.				

引用名	unshare_nfs – make local NFS file systems unavailable for mounting by remote systems
用法概要	unshare [-F nfs] <i>pathname</i>
描述	The unshare command makes local file systems unavailable for mounting by remote systems. The shared file system must correspond to a line with NFS as the <i>FSType</i> in the file <i>/etc/dfs/sharetab</i> .
选项	The following options are supported: -F This option may be omitted if NFS is the first file system type listed in the file <i>/etc/dfs/fstypes</i> .
文件	<i>/etc/dfs/fstypes</i> <i>/etc/dfs/sharetab</i>
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/file-system/nfs

另请参见	nfsd(1M) , share(1M) , attributes(5)
附注	If the file system being unshared is a symbolic link to a valid pathname, the canonical path (the path which the symbolic link follows) will be unshared. For example, if <i>/export/foo</i> is a symbolic link to <i>/export/bar</i> (<i>/export/foo</i> -> <i>/export/bar</i>), the following unshare command will result in <i>/export/bar</i> as the unshared pathname (and not <i>/export/foo</i>): <pre>example# unshare -F nfs /export/foo</pre> For file systems that are accessed by NFS Version 4 clients, once the unshare is complete, all NFS Version 4 state (open files and file locks) are released and unrecoverable by the clients. If the intent is to share the file system after some administrative action, the NFS daemon (<i>nfsd</i>) should first be stopped and then the file system unshared. After the administrative action is complete, the file system would then be shared and the NFS daemon restarted. See nfsd(1M)

参 考 文 档

系统管理命令—第 3 单元

引用名

update_drv – modify device driver attributes

用法概要

```
update_drv [-f | -v] [-n] driver_module

update_drv [-b basedir] [-f | -v] [-n] -a [-m 'permission']
          [-i 'identify-name'] [-P 'privilege'] [-p 'policy'] driver_module

update_drv [-b basedir] [-f | -v] [-n] -d [-m 'permission']
          [-i 'identify-name'] [-P 'privilege'] [-p 'policy'] driver_module
```

描述

The `update_drv` command informs the system about attribute changes to an installed device driver. It can be used to re-read the [driver.conf\(4\)](#) file, or to add, modify, or delete a driver's minor node permissions or aliases.

Without options, `update_drv` reloads the `driver.conf` file.

Upon successfully updating the aliases, the driver binding takes effect upon reconfig boot or hotplug of the device.

Upon successfully updating the permissions, only the new driver minor nodes get created with the modified set of file permissions. Existing driver minor nodes do not get modified.

选项

The following options are supported:

-a

Add a *permission*, *aliases*, *privilege* or *policy* entry.

With the `-a` option specified, a permission entry (using the `-m` option), or a driver's aliases entry (using the `-i` option), a device privilege (using the `-P` option) or a device policy (using the `-p` option), can be added or updated. If a matching minor node permissions entry is encountered (having the same driver name and the minor node), it is replaced. If a matching aliases entry is encountered (having a different driver name and the same alias), an error is reported.

The `-a` and `-d` options are mutually exclusive.

-b *basedir*

Installs or modifies the driver on the system with a root directory of *basedir* rather than installing on the system executing `update_drv`.

注 – The root file system of any non-global zones must not be referenced with the `-b` option. Doing so might damage the global zone's file system, might compromise the security of the global zone, and might damage the non-global zone's file system. See [zones\(5\)](#).

-d

Delete a *permission*, *aliases*, *privilege* or *policy* entry.

The `-m permission`, `-i identify-name`, `-P privilege` or the `-p policy` option needs to be specified with the `-d` option.

The `-d` and `-a` options are mutually exclusive.

If the entry doesn't exist `update_drv` returns an error.

-f

Force the system to reread the `driver.conf` file even if the driver module cannot be unloaded. See NOTES section for details.

Without this option, when removing an alias for a driver, `update_drv` updates the binding files for the next boot, but returns an error if one or more devices that reference the driver-alias binding remains. *With* the -f option, `update_drv` does not return an error if such devices remain.

-i '*identify-name*'

A white-space separated list of aliases for the driver. If -a or -d option is not specified then this option is ignored. The *identify-name* string is mandatory. If all aliases need to be removed, `rem_drv(1M)` is recommended.

-m '*permission*'

Specify a white-space separated list of file system permissions for the device node of the device driver. If -a or -d option is not specified then, this option is ignored. The permission string is mandatory.

-n

Do not try to load and attach *device_driver*, just modify the system configuration files for that driver.

-p '*policy*'

With the -a option, policy is a white-space separated list of complete device policies. For the -d option, policy is a white space separated list of minor device specifications. The minor device specifications are matched exactly against the entries in `/etc/security/device_policy`, that is., no wildcard matching is performed.

-P '*privilege*'

With the -a option, privilege is a comma separated list of additional driver privileges. For the -d option, privilege is a single privilege. The privileges are added to or removed from the `/etc/security/extra_privs` file.

-v

Verbose.

示例

示例 1 Adding or Modifying an Existing Minor Permissions Entry

The following command adds or modifies the existing minor permissions entry of the clone driver:

```
example# update_drv -a -m 'llc1 777 joe staff' clone
```

示例 2 Removing All Minor Permissions Entries

The following command removes all minor permission entries of the `usbprn` driver, the USB printer driver:

示例 2 Removing All Minor Permissions Entries (续)

```
example# update_drv -d -m '* 0666 root sys' usbprn
```

示例 3 Adding a Driver Aliases Entry

The following command adds a driver aliases entry of the ugen driver with the identity string of usb459,20:

```
example# update_drv -a -i 'usb459,20' ugen
```

示例 4 Re-reading the driver.conf File For the ohci Driver

The following command re-reads the driver.conf(4) file.

```
example# update_drv ohci
```

示例 5 Requiring a Self-defined Privilege to Open a tcp Socket

The following command requires a self-defined privilege to open a tcp socket:

```
example# update_drv -a -P net_tcp -p \
    'write_priv_set=net_tcp read_priv_set=net_tcp' tcp
```

示例 6 Establishing a Path-oriented Alias

The following command establishes a path-oriented alias to force a specific driver, qlt, to be used for a particular device path:

```
example# update_drv -a -i '/pci@8,600000/SUNW,qlc@4' qlt
```

退出状态

The following exit values are returned:

- 0
Successful completion.
- >0
An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[add_drv\(1M\)](#), [modunload\(1M\)](#), [rem_drv\(1M\)](#), [driver.conf\(4\)](#), [attributes\(5\)](#), [privileges\(5\)](#)

附注

If -a or -d options are specified, update_drv does not reread the driver.conf file.

A forced update of the `driver.conf` file reloads the `driver.conf` file without reloading the driver binary module. In this case, devices which cannot be detached reference driver global properties from the old `driver.conf` file, while the remaining driver instances reference global properties in the new `driver.conf` file.

It is possible to add an alias, which changes the driver binding of a device already being managed by a different driver. A force update with the `-a` option tries to bind to the new driver and report error if it cannot. If you specify more than one of the `-m`, `-i`, `-P` or `-p` options, a force flag tries to modify aliases or permissions. This is done even if the other operation fails and vice-versa.

引用名

useradd – administer a new user login on the system

用法概要

```
useradd [-A authorization [,authorization...]]
        [-b base_dir] [-c comment] [-d dir] [-e expire]
        [-f inactive] [-g group] [-G group [,group]...]
        [-K key=value] [-m [-k skel_dir]] [-p projname]
        [-P profile [,profile...]] [-R role [,role...]]
        [-s shell] [-S repository] [-u uid [-o]] login

useradd -D [-A authorization [,authorization...]]
        [-b base_dir] [-s shell [-k skel_dir]] [-e expire]
        [-f inactive] [-g group] [-K key=value] [-p projname]
        [-P profile [,profile...]] [-R role [,role...]]
```

描述

useradd adds a new user to the `passwd`, `shadow`, and `user_attr` databases in the `files` and `ldap` repositories. The `-A` and `-P` options respectively assign authorizations and profiles to the user. The `-R` option assigns roles to a user. The `-p` option associates a project with a user. The `-K` option adds a `key=value` pair to `user_attr` entry for the user. Multiple `key=value` pairs may be added with multiple `-K` options.

useradd also creates supplementary group memberships for the user (`-G` option) and creates the home directory (`-m` option) for the user if requested. The new login remains locked until the [passwd\(1\)](#) command is executed.

Specifying useradd `-D` with the `-s`, `-k`, `-g`, `-b`, `-f`, `-e`, `-A`, `-P`, `-p`, `-R`, or `-K` option (or any combination of these options) sets the default values for the respective fields. See the `-D` option, below. Subsequent useradd commands without the `-D` option use these arguments.

The system file entries created with this command have a limit of 2048 characters per line. Specifying long arguments to several options can exceed this limit.

useradd requires that usernames be in the format described in [passwd\(4\)](#). A warning message is displayed if these restrictions are not met. See [passwd\(4\)](#) for the requirements for usernames.

An administrator must be granted the User Management Profile to be able to create a new user. The authorizations required to set the various fields in `passwd`, `shadow` and `user_attr` can be found in [passwd\(4\)](#), [shadow\(4\)](#), and [user_attr\(4\)](#). The authorizations required to assign groups and projects can be found in [group\(4\)](#) and [project\(4\)](#).

选项

The following options are supported:

-A authorization

One or more comma-separated authorizations defined in [auth_attr\(4\)](#). Only a user or role who has grant rights to the authorization can assign it to an account.

-b base_dir

The base directory for new login home directories (see the `-d` option below. When a new user account is being created, `base_dir` must already exist unless the `-m` option or the `-d` option is also specified.

-c *comment*

Any text string. It is generally a short description of the login, and is currently used as the field for the user's full name. This information is stored in the user's `passwd` entry.

-d *dir*

Specifies the `auto_home` entry for the new user. The path `/home/username` is entered in `/etc/passwd`. When the user subsequently references `/home/username`, the automounter will mount the directory specified here on `/home/username`.

The argument to the `-d` option can be specified as `server:dir` where *server* is the hostname of the machine on which the home directory resides and *dir* is the path to the user's home directory. If the server is a remote host, then the home directory needs to be created on the remote host for the system to mount it, when the user logs in. If no server name is specified then the home directory will be created on the host where the command is executed, when the `-m` option is used.

-D

Display the default values for `group`, `base_dir`, `skel_dir`, `shell`, `inactive`, `expire`, `proj`, `projname` and `key=value` pairs. When used with the `-g`, `-b`, `-f`, `-e`, `-A`, `-P`, `-p`, `-R`, or `-K` options, the `-D` option sets the default values for the specified fields. The default values are:

```
group
  other (GID of 1)
```

```
base_dir
  /export/home
```

```
skel_dir
  /etc/skel
```

```
shell
  /usr/bin/bash
```

```
inactive
  0
```

```
expire
  null
```

```
auths
  null
```

```
profiles
  null
```

```
proj
  3
```

```
projname
  default
```

key=value (pairs defined in [user_attr\(4\)](#))
not present

roles
null

-e *expire*

Specify the expiration date for a login. After this date, no user will be able to access this login. The expire option argument is a date entered using one of the date formats included in the template file `/etc/datemsk`. See [getdate\(3C\)](#).

If the date format that you choose includes spaces, it must be quoted. For example, you can enter `10/6/90` or `October 6, 1990`. A null value (`" "`) defeats the status of the expired date. This option is useful for creating temporary logins.

-f *inactive*

The maximum number of days allowed between uses of a login ID before that ID is declared invalid. Normal values are positive integers. A value of `0` defeats the status.

-g *group*

An existing group's integer ID or character-string name. Without the `-D` option, it defines the new user's primary group membership and defaults to the default group. You can reset this default value by invoking `useradd -D -g group`. GIDs 0-99 are reserved for allocation by the Solaris Operating System.

-G *group*

An existing group's integer ID or character-string name. It defines the new user's supplementary group membership. Duplicates between *group* with the `-g` and `-G` options are ignored. No more than `NGROUPS_MAX` groups can be specified. GIDs 0-99 are reserved for allocation by the Solaris Operating System.

-K *key=value*

A *key=value* pair to add to the user's attributes. Multiple `-K` options may be used to add multiple *key=value* pairs. The generic `-K` option with the appropriate key may be used instead of the specific implied key options (`-A`, `-P`, `-R`, `-p`). See [user_attr\(4\)](#) for a list of valid *key=value* pairs. The "type" key is not a valid key for this option. Keys may not be repeated.

-k *skel_dir*

A directory that contains skeleton information (such as `.profile`) that can be copied into a new user's home directory. This directory must already exist. The system provides the `/etc/skel` directory that can be used for this purpose.

-m

Create the new user's home directory if it does not already exist. If the directory already exists, it must have read, write, and execute permissions by *group*, where *group* is the user's primary group. If the server name specified to the `-d` option is a remote host then the system will not attempt to create the home directory.

If the directory does not already exist, a new ZFS dataset will be created. In the global zone, the dataset is created as `rpool/export/home/username`. For non-global zones, the dataset will be created as `ROOT-dataset/export/home/username`. The mountpoint for the ZFS dataset is `/export/home/username` by default. If `-d path` is specified and it is a path on the local machine, the dataset will be mounted at the specified location. The user is delegated permissions to create ZFS snapshots and promote them. The newly created dataset will inherit the encryption setting from its parent. If it is encrypted, the user is granted permission to change its wrapping key.

- o
This option allows a UID to be duplicated (non-unique).
- P *profile*
One or more comma-separated execution profiles defined in [prof_attr\(4\)](#).
- p *projname*
Name of the project with which the added user is associated. See the *projname* field as defined in [project\(4\)](#).
- R *role*
One or more comma-separated execution profiles defined in [user_attr\(4\)](#). Roles cannot be assigned to other roles.
- s *shell*
Full pathname of the program used as the user's shell on login. If unspecified, it will default to any value previously configured with the `-D -s` option. If no default has been set with `-D -s`, then `/usr/bin/bash` will be used. The value of *shell* must be a valid executable file.
- S *repository*
The valid repositories are `files`, `ldap`. The repository specifies which name service will be updated. The default repository is `files`. When the repository is `files`, the authorizations, profiles, and roles can be present in other name service repositories and can be assigned to a user in the `files` repository. When the repository is `ldap`, all the assignable attributes must be present in the `ldap` repository.
- u *uid*
The UID of the new user. This UID must be a non-negative decimal integer below `MAXUID` as defined in `<sys/param.h>`. The UID defaults to the next available (unique) number above the highest number currently assigned. For example, if UIDs 100, 105, and 200 are assigned, the next default UID number will be 201. UIDs 0-99 are reserved for allocation by the Solaris Operating System.

文件

`/etc/datemsk`

`/etc/passwd`

`/etc/shadow`

`/etc/group`

```
/etc/skel
/usr/include/limits.h
/etc/user_attr
```

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [passwd\(1\)](#), [profiles\(1\)](#), [roles\(1\)](#), [users\(1B\)](#), [groupadd\(1M\)](#), [groupdel\(1M\)](#), [groupmod\(1M\)](#), [grpck\(1M\)](#), [logins\(1M\)](#), [pwck\(1M\)](#), [userdel\(1M\)](#), [usermod\(1M\)](#), [getdate\(3C\)](#), [auth_attr\(4\)](#), [group\(4\)](#), [passwd\(4\)](#), [prof_attr\(4\)](#), [project\(4\)](#), [user_attr\(4\)](#), [attributes\(5\)](#)

诊断 In case of an error, useradd displays an error message and exits with a non-zero status.

The following indicates that login specified is already in use:

UX: useradd: ERROR: login is already in use. Choose another.

The following indicates that the *uid* specified with the -u option is not unique:

UX: useradd: ERROR: uid *uid* is already in use. Choose another.

The following indicates that the *group* specified with the -g option has not yet been created:

UX: useradd: ERROR: group *group* does not exist. Choose another.

The following indicates that the *uid* specified with the -u option is in the range of reserved UIDs (from 0-99):

UX: useradd: WARNING: uid *uid* is reserved.

The following indicates that the *uid* specified with the -u option exceeds MAXUID as defined in <sys/param.h>:

UX: useradd: ERROR: uid *uid* is too big. Choose another.

The following indicates that the /etc/passwd or /etc/shadow files do not exist:

UX: useradd: ERROR: Cannot update system files - login cannot be created.

The following indicates that the user executing the command does not have sufficient authorization to perform the operation:

UX: roleadd: ERROR: Permission denied.

The following indicates that an invalid directory was specified in a useradd command:

UX: *invalid_directory* is not a valid directory. Choose another.

附注

The `useradd` utility adds definitions to the `passwd`, `shadow`, `group`, `project`, and `user_attr` databases in the scope (default or specified). It will verify the uniqueness of the user name (or role) and user id and the existence of any group names specified against the external name service.

引用名	userdel – delete a user's login from the system												
用法概要	userdel [-r] [-S repository] login												
描述	The userdel utility deletes a user account from the system and makes the appropriate account-related changes to the system file and file system. An administrator must be granted the User Management Profile to be able to delete a user.												
选项	The following options are supported: <table><tr><td>-r</td><td>Remove the user's home directory from the system. This directory must exist. The files and directories under the home directory will no longer be accessible following successful execution of the command. A ZFS dataset that was created for the user's home directory will be removed. An auto_home entry that was added for the user will be deleted.</td></tr><tr><td>-S repository</td><td>The valid repositories are files, ldap. The repository specifies which name service will be updated. The default repository is files.</td></tr></table>	-r	Remove the user's home directory from the system. This directory must exist. The files and directories under the home directory will no longer be accessible following successful execution of the command. A ZFS dataset that was created for the user's home directory will be removed. An auto_home entry that was added for the user will be deleted.	-S repository	The valid repositories are files, ldap. The repository specifies which name service will be updated. The default repository is files.								
-r	Remove the user's home directory from the system. This directory must exist. The files and directories under the home directory will no longer be accessible following successful execution of the command. A ZFS dataset that was created for the user's home directory will be removed. An auto_home entry that was added for the user will be deleted.												
-S repository	The valid repositories are files, ldap. The repository specifies which name service will be updated. The default repository is files.												
操作数	The following operands are supported: <table><tr><td>login</td><td>An existing login name to be deleted.</td></tr></table>	login	An existing login name to be deleted.										
login	An existing login name to be deleted.												
退出状态	The following exit values are returned: <table><tr><td>0</td><td>Successful completion.</td></tr><tr><td>2</td><td>Invalid command syntax. A usage message for the userdel command is displayed.</td></tr><tr><td>6</td><td>The account to be removed does not exist.</td></tr><tr><td>8</td><td>The account to be removed is in use.</td></tr><tr><td>10</td><td>Cannot update the /etc/group or /etc/user_attr file but the login is removed from the /etc/passwd file.</td></tr><tr><td>12</td><td>Cannot remove or otherwise modify the home directory.</td></tr></table>	0	Successful completion.	2	Invalid command syntax. A usage message for the userdel command is displayed.	6	The account to be removed does not exist.	8	The account to be removed is in use.	10	Cannot update the /etc/group or /etc/user_attr file but the login is removed from the /etc/passwd file.	12	Cannot remove or otherwise modify the home directory.
0	Successful completion.												
2	Invalid command syntax. A usage message for the userdel command is displayed.												
6	The account to be removed does not exist.												
8	The account to be removed is in use.												
10	Cannot update the /etc/group or /etc/user_attr file but the login is removed from the /etc/passwd file.												
12	Cannot remove or otherwise modify the home directory.												
文件	<table><tr><td>/etc/passwd</td><td>system password file</td></tr><tr><td>/etc/shadow</td><td>system file contain users' encrypted passwords and related information</td></tr><tr><td>/etc/group</td><td>system file containing group definitions</td></tr><tr><td>/etc/user_attr</td><td>system file containing additional user attributes</td></tr></table>	/etc/passwd	system password file	/etc/shadow	system file contain users' encrypted passwords and related information	/etc/group	system file containing group definitions	/etc/user_attr	system file containing additional user attributes				
/etc/passwd	system password file												
/etc/shadow	system file contain users' encrypted passwords and related information												
/etc/group	system file containing group definitions												
/etc/user_attr	system file containing additional user attributes												
属性	See attributes(5) for descriptions of the following attributes:												

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见	auths(1) , passwd(1) , profiles(1) , roles(1) , users(1B) , groupadd(1M) , groupdel(1M) , groupmod(1M) , logins(1M) , roleadd(1M) , rolemod(1M) , useradd(1M) , usermod(1M) , passwd(4) , prof_attr(4) , user_attr(4) , attributes(5)
附注	The userdel utility deletes an account definition that is in the group, passwd, shadow, and user_attr databases in the files or ldap repository.

引用名

usermod – modify a user's login information on the system

用法概要

```
usermod [-u uid [-o]] [-g group] [-G [+|-]group [, group...]]
        [-d dir [-m]] [-s shell] [-c comment] [-l new_name]
        [-f inactive] [-e expire]
        [-A [+|-]authorization [,authorization]]
        [-P [+|-]profile [,profile]] [-R [+|-]role [,role]]
        [-K key[+|-]=value] [-S repository] login
```

描述

The usermod utility modifies a user's login definition on the system. It changes the definition of the specified login and makes the appropriate login-related system file and file system changes.

The system file entries created with this command have a limit of 512 characters per line. Specifying long arguments to several options might exceed this limit.

An administrator must be granted the User Security Profile to modify the security attributes for an existing user. To be able to modify the non-security attributes of an existing user requires the User Management Profile. The authorizations required to set the various fields in `passwd`, `shadow` and `user_attr` can be found in [passwd\(4\)](#), [shadow\(4\)](#), and [user_attr\(4\)](#). The authorizations required to assign groups can be found in [group\(4\)](#).

选项

The following options are supported:

-A [+|-]*authorization*

One or more comma separated authorizations as defined in [auth_attr\(4\)](#). Only a user or role who has grant rights to the *authorization* can assign it to an account. This replaces any existing authorization setting. If no authorization list is specified, the existing setting is removed.

A prefix + adds the authorization to the existing authorization; a prefix - removes the authorization from the existing authorization. With no prefix, *authorization* replaces the existing authorization.

-c *comment*

Specify a comment string. *comment* can be any text string. It is generally a short description of the login, and is currently used as the field for the user's full name. This information is stored in the user's `passwd` entry.

-d *dir*

Specify the new home directory of the user. It defaults to `base_dir/login`, where `base_dir` is the base directory for new login home directories, and `login` is the new login. This creates or modifies an `auto_home` entry for the user.

The argument to the option can be specified as `server:dir` where *server* is the hostname of the machine on which the home directory resides and *dir* is the path to the user's home directory. If the server is a remote host then the home directory needs to be created on the

remote host for the system to mount it, when the user logs in. If no server name is specified then the home directory will be created on the host where the command is executed, when the `-m` option is used.

-e *expire*

Specify the expiration date for a login. After this date, no user will be able to access this login. The expire option argument is a date entered using one of the date formats included in the template file `/etc/datemsk`. See [getdate\(3C\)](#).

For example, you may enter `10/6/90` or `October 6, 1990`. A value of “ ” defeats the status of the expired date.

-f *inactive*

Specify the maximum number of days allowed between uses of a login ID before that login ID is declared invalid. Normal values are positive integers. A value of `0` defeats the status.

-g *group*

Specify an existing group's integer ID or character-string name. It redefines the user's primary group membership.

-G *[+|-]group*

An existing group's integer ID or character-string name. It defines the new user's supplementary group membership. Duplicates between group with the `-g` and `-G` options are ignored. No more than `NGROUPS_MAX` groups can be specified. GIDs 0-99 are reserved for allocation by the Solaris Operating System.

A prefix `+` adds the group to the existing group; a prefix `-` removes the group from the existing group. With no prefix, *group* replaces the existing group.

-K *key[+|-]=value*

Replace existing or add to a user's *key=value* pair attributes. Multiple `-K` options can be used to replace or add multiple *key=value* pairs. However, keys must not be repeated. The generic `-K` option with the appropriate key can be used instead of the specific implied key options (`-A`, `-P`, `-R`, `-p`). See [user_attr\(4\)](#) for a list of valid keys. Values for these keys are usually found in man pages or other sources related to those keys. For example, see [project\(4\)](#) for guidance on values for the `project` key. Use the command [ppriv\(1\)](#) with the `-v` and `-l` options for a list of values for the keys `defaultpriv` and `limitpriv`. If no value is specified, the existing key is removed.

The keyword type can be specified with the value `role` or the value `normal`. When using the value `role`, the account changes from a normal user to a role; using the value `normal` keeps the account a normal user.

As a `role` account, no roles (`-R` or *roles=value*) can be present.

A prefix `+` adds the value to the existing value; a prefix `-` removes the value from the existing value. With no prefix, *value* replaces the existing value.

The prefix `+/ -` operation is applicable only to the following keys: `auths`, `profiles`, `roles`, `project`, `limitpriv`, and `defaultpriv`.

-l *new_logname*

Specify the new login name for the user. See [passwd\(4\)](#) for the requirements for usernames.

-m

Move the user's home directory to the new directory specified with the `-d` option. If the directory already exists, it must have permissions read/write/execute by *group*, where *group* is the user's primary group. If the server name specified to the `-d` option is a remote host then the system will not attempt to create the home directory.

If the directory does not already exist, a new ZFS dataset will be created. In the global zone, the dataset is created as `rpool/export/home/rolename`. For non-global zones, the dataset will be created as `ROOT-dataset/export/home/rolename`. The mountpoint for the ZFS dataset is `/export/home/rolename` by default. If `-d path` is specified and it is a path on the local machine, the dataset will be mounted at the specified location. The role is delegated permissions to create ZFS snapshots and promote them. The newly created dataset will inherit the encryption setting from its parent. If it is encrypted, the role is granted permission to change its wrapping key.

-o

This option allows the specified UID to be duplicated (non-unique).

-P [+|-] *profile*

One or more comma-separated rights profiles defined in [prof_attr\(4\)](#). This replaces any existing profile setting in [user_attr\(4\)](#). If an empty profile list is specified, the existing setting is removed.

A prefix `+` adds the profile to the existing profile; a prefix `-` removes the profile from the existing profile. With no prefix, *profile* replaces the existing profile.

-R [+|-] *role*

One or more comma-separated roles (see [roleadd\(1M\)](#)). This replaces any existing role setting. If no role list is specified, the existing setting is removed.

A prefix `+` adds the role to the existing role; a prefix `-` removes the role from the existing role. With no prefix, *role* replaces the existing role.

-s *shell*

Specify the full pathname of the program that is used as the user's shell on login. The value of *shell* must be a valid executable file.

-S *repository*

The valid repositories are `files`, `ldap`. The repository specifies which name service will be updated. The default repository is `files`. When the repository is `files`, the authorizations, profiles, and roles can be present in other name service repositories and can be assigned to a user in the `files` repository. When the repository is `ldap`, all the assignable attributes must be present in the `ldap` repository.

-u *uid*

Specify a new UID for the user. It must be a non-negative decimal integer less than MAXUID as defined in `<param.h>`. The UID associated with the user's home directory is not modified with this option; a user will not have access to their home directory until the UID is manually reassigned using [chown\(1\)](#).

操作数

The following operands are supported:

login

An existing login name to be modified.

示例

示例 1 Assigning Privileges to a User

The following command adds the privilege that affects high resolution times to a user's initial, inheritable set of privileges.

```
# usermod -K defaultpriv=basic,proc_clock_highres jdoe
```

This command results in the following entry in `user_attr`:

```
jdoe:::type=normal;defaultpriv=basic,proc_clock_highres
```

示例 2 Removing a Privilege from a User's Limit Set

The following command removes the privilege that allows the specified user to create hard links to directories and to unlink directories.

```
# usermod -K limitpriv=all,!sys_linkdir jdoe
```

This command results in the following entry in `user_attr`:

```
jdoe:::type=normal;defaultpriv=basic,limitpriv=all,!sys_linkdir
```

示例 3 Removing a Privilege from a User's Basic Set

The following command removes the privilege that allows the specified user to examine processes outside the user's session.

```
# usermod -K defaultpriv=basic,!proc_session jdoe
```

This command results in the following entry in `user_attr`:

```
jdoe:::type=normal;defaultpriv=basic,!proc_session;limitpriv=all
```

示例 4 Assigning a Role to a User

The following command assigns a role to a user. The role must have been created prior to this command through use of [roleadd\(1M\)](#).

```
# usermod -R mailadm jdoe
```

This command results in the following entry in `user_attr`:

```
jdoe:::type=normal;roles=mailadm;defaultpriv=basic;limitpriv=all
```

示例 5 Removing All Profiles from a User

The following command removes all profiles that were granted to a user directly. The user will still have any rights profiles that are granted by means of the PROFS_GRANTED key in [policy.conf\(4\)](#).

```
# usermod -P "" jdoe
```

退出状态

In case of an error, usermod prints an error message and exits with one of the following values:

- 2
The command syntax was invalid. A usage message for the usermod command is displayed.
- 3
An invalid argument was provided to an option.
- 4
The *uid* given with the -u option is already in use.
- 5
The password files contain an error. [pwconv\(1M\)](#) can be used to correct possible errors. See [passwd\(4\)](#).
- 6
The login to be modified does not exist, the *group* does not exist, or the login shell does not exist.
- 8
The login to be modified is in use.
- 9
The *new_logname* is already in use.
- 10
Cannot update the /etc/group or /etc/user_attr file. Other update requests will be implemented.
- 11
Insufficient space to move the home directory (-m option). Other update requests will be implemented.
- 12
Unable to complete the move of the home directory to the new home directory.

文件

- /etc/group
system file containing group definitions
- /etc/datemsk
system file of date formats
- /etc/passwd
system password file

`/etc/shadow`
system file containing users' encrypted passwords and related information

`/etc/user_attr`
system file containing additional user and role attributes

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见 [chown\(1\)](#), [passwd\(1\)](#), [users\(1B\)](#), [groupadd\(1M\)](#), [groupdel\(1M\)](#), [groupmod\(1M\)](#), [logins\(1M\)](#), [pwconv\(1M\)](#), [roleadd\(1M\)](#), [roledel\(1M\)](#), [rolemod\(1M\)](#), [useradd\(1M\)](#), [userdel\(1M\)](#), [getdate\(3C\)](#), [auth_attr\(4\)](#), [passwd\(4\)](#), [policy.conf\(4\)](#), [prof_attr\(4\)](#), [user_attr\(4\)](#), [attributes\(5\)](#)

附注 The usermod utility modifies definitions in the passwd, shadow, group, project, and user_attr databases in the scope (default or specified). The utility will verify the uniqueness of user name and user ID against the external name service.

The usermod utility uses the `/etc/datemsk` file, available with SUNWaccr, for date formatting.

引用名	utmpd – utmpx monitoring daemon
用法概要	utmpd [-debug]
描述	The utmpd daemon monitors the /var/adm/utmpx file. See utmpx(4) (and utmp(4) for historical information). utmpd receives requests from pututxline(3C) by way of a named pipe. It maintains a table of processes and uses poll(2) on /proc files to detect process termination. When utmpd detects that a process has terminated, it checks that the process has removed its utmpx entry from /var/adm/utmpx. If the process' utmpx entry has not been removed, utmpd removes the entry. By periodically scanning the /var/adm/utmpx file, utmpd also monitors processes that are not in its table.
选项	-debug Run in debug mode, leaving the process connected to the controlling terminal. Write debugging information to standard output.
退出状态	The following exit values are returned: 0 Successful completion. >0 An error occurred.
文件	/etc/default/utmpd You can set default values for the flags listed below. For example: SCAN_PERIOD=600 SCAN_PERIOD The number of seconds that utmpd sleeps between checks of /proc to see if monitored processes are still alive. The default is 300. MAX_FDS The maximum number of processes that utmpd attempts to monitor. The default value is 4096. WTMPX_UPDATE_FREQ The number of seconds that utmpd sleeps between read accesses of the wtmpx file. The wtmpx file's last access time is used by init(1M) on reboot to determine when the operating system became unavailable. The default is 60. /var/adm/utmpx File containing user and accounting information for commands such as who(1), write(1), and login(1). /proc Directory containing files for processes whose utmpx entries are being monitored.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见 [svcs\(1\)](#), [init\(1M\)](#), [svcadm\(1M\)](#), [poll\(2\)](#), [pututxline\(3C\)](#), [proc\(4\)](#), [utmp\(4\)](#), [utmpx\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 If the filesystem holding `/var/adm/wtmpx` is mounted with options which inhibit or defer access time updates, an unknown amount of error will be introduced into the `utmp DOWN_TIME` record's timestamp in the event of an uncontrolled shutdown (for example, a crash or loss of power). Controlled shutdowns will update the modify time of `/var/adm/wtmpx`, which will be used on the next boot to determine when the previous shutdown occurred, regardless of access time deferral or inhibition.

The `utmpd` service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/system/utmp:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	uuccheck – check the uucp directories and permissions file				
用法概要	<code>/usr/lib/uucp/uuccheck [-v] [-x <i>debug-level</i>]</code>				
描述	uuccheck checks for the presence of the uucp system required files and directories. uuccheck also does error checking of the Permissions file (<code>/etc/uucp/Permissions</code>). uuccheck is executed during package installation. uuccheck can only be used by the super-user or uucp.				
选项	The following options are supported: <table><tr><td>-v</td><td>Give a detailed (verbose) explanation of how the uucp programs will interpret the Permissions file.</td></tr><tr><td>-x <i>debug-level</i></td><td>Produce debugging output on the standard output. <i>debug-level</i> is a number from 0 to 9. Higher numbers give more detailed debugging information.</td></tr></table>	-v	Give a detailed (verbose) explanation of how the uucp programs will interpret the Permissions file.	-x <i>debug-level</i>	Produce debugging output on the standard output. <i>debug-level</i> is a number from 0 to 9. Higher numbers give more detailed debugging information.
-v	Give a detailed (verbose) explanation of how the uucp programs will interpret the Permissions file.				
-x <i>debug-level</i>	Produce debugging output on the standard output. <i>debug-level</i> is a number from 0 to 9. Higher numbers give more detailed debugging information.				
文件	<code>/etc/uucp/Devices</code> <code>/etc/uucp/Limits</code> <code>/etc/uucp/Permissions</code> <code>/etc/uucp/Systems</code> <code>/var/spool/locks/*</code> <code>/var/spool/uucp/*</code> <code>/var/spool/uucppublic/*</code>				
属性	See attributes(5) for descriptions of the following attributes:				
	<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>service/network/uucp</td></tr></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	service/network/uucp
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	service/network/uucp				
另请参见	uucp(1C) , uustat(1C) , uux(1C) , uucico(1M) , uusched(1M) , attributes(5)				
已知问题	The program does not check file/directory modes or some errors in the Permissions file such as duplicate login or machine name.				

引用名	uucico – file transport program for the uucp system
用法概要	<pre>/usr/lib/uucp/uucico [-f] [-c type] [-d spool-directory] [-i interface] [-r role-number] [-s system-name] [-x debug-level]</pre>
描述	uucico is the file transport program for uucp work file transfers.
选项	The following options are supported: -f This option is used to "force execution" of uucico by ignoring the limit on the maximum number of uucicos defined in the /etc/uucp/Limits file. -c type The first field in the Devices file is the "Type" field. The -c option forces uucico to only use entries in the "Type" field that match the user specified type. The specified type is usually the name of a local area network. -d spool-directory This option specifies the directory <i>spool-directory</i> that contains the uucp work files to be transferred. The default spool directory is /var/spool/uucp. -i interface This option defines the <i>interface</i> used with uucico. The interface only affects slave mode. Known interfaces are UNIX (default), TLI (basic Transport Layer Interface), and TLIS (Transport Layer Interface with Streams modules, read/write). -r role-number The <i>role-number</i> 1 is used for master mode. <i>role-number</i> 0 is used for slave mode (default). When uucico is started by a program or cron, <i>role-number</i> 1 should be used for master mode. -s system-name The -s option defines the remote system (<i>system-name</i>) that uucico will try to contact. It is required when the role is master; <i>system-name</i> must be defined in the Systems file. -x debug-level Both uux and uucp queue jobs that will be transferred by uucico. These jobs are normally started by the uusched scheduler, for debugging purposes, and can be started manually. For example, the shell Uutry starts uucico with debugging turned on. The <i>debug-level</i> is a number between 0 and 9. Higher numbers give more detailed debugging information.

文件

`/etc/uucp/Devconfig`
`/etc/uucp/Devices`
`/etc/uucp/Limits`
`/etc/uucp/Permissions`
`/etc/uucp/Sysfiles`
`/etc/uucp/Systems`
`/var/spool/locks/*`
`/var/spool/uucp/*`
`/var/spool/uucppublic/*`

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/uucp

另请参见 [uucp\(1C\)](#), [uustat\(1C\)](#), [uux\(1C\)](#), [Uutry\(1M\)](#), [cron\(1M\)](#), [uusched\(1M\)](#), [attributes\(5\)](#)

引用名	uucleanup – uucp spool directory clean-up																
用法概要	<pre>/usr/lib/uucp/uucleanup [-Ctime] [-Dtime] [-mstring] [-otime] [-ssystem] [-Wtime] [-xdebug-level] [-Xtime]</pre>																
描述	uucleanup will scan the spool directories for old files and take appropriate action to remove them in a useful way: ▪ Inform the requester of send/receive requests for systems that can not be reached. ▪ Return undeliverable mail to the sender. ▪ Deliver rnews files addressed to the local system. ▪ Remove all other files. In addition, there is a provision to warn users of requests that have been waiting for a given number of days (default 1 day). Note: uucleanup will process as if all option times were specified to the default values unless time is specifically set. This program is typically started by the shell uudemond.cleanup, which should be started by cron(1M).																
选项	<table> <tr> <td>-Ctime</td><td>Remove any C. files greater or equal to time days old and send appropriate information to the requester (default 7 days).</td></tr> <tr> <td>-Dtime</td><td>Remove any D. files greater or equal to time days old, make an attempt to deliver mail messages, and execute rnews when appropriate (default 7 days).</td></tr> <tr> <td>-mstring</td><td>Include string in the warning message generated by the -W option. The default line is "See your local administrator to locate the problem".</td></tr> <tr> <td>-otime</td><td>Delete other files whose age is more than time days (default 2 days).</td></tr> <tr> <td>-ssystem</td><td>Execute for system spool directory only.</td></tr> <tr> <td>-Wtime</td><td>Any C. files equal to time days old will cause a mail message to be sent to the requester warning about the delay in contacting the remote. The message includes the JOBID, and in the case of mail, the mail message. The administrator may include a message line telling whom to call to check the problem (-m option) (default 1 day).</td></tr> <tr> <td>-xdebug-level</td><td>Produce debugging output on standard output. debug-level is a single digit between 0 and 9; higher numbers give more detailed debugging information. (This option may not be available on all systems.)</td></tr> <tr> <td>-Xtime</td><td>Any X. files greater or equal to time days old will be removed. The D. files are probably not present (if they were, the X. could get executed). But if there are D. files, they will be taken care of by D. processing (default 2 days).</td></tr> </table>	-Ctime	Remove any C. files greater or equal to time days old and send appropriate information to the requester (default 7 days).	-Dtime	Remove any D. files greater or equal to time days old, make an attempt to deliver mail messages, and execute rnews when appropriate (default 7 days).	-mstring	Include string in the warning message generated by the -W option. The default line is "See your local administrator to locate the problem".	-otime	Delete other files whose age is more than time days (default 2 days).	-ssystem	Execute for system spool directory only.	-Wtime	Any C. files equal to time days old will cause a mail message to be sent to the requester warning about the delay in contacting the remote. The message includes the JOBID, and in the case of mail, the mail message. The administrator may include a message line telling whom to call to check the problem (-m option) (default 1 day).	-xdebug-level	Produce debugging output on standard output. debug-level is a single digit between 0 and 9; higher numbers give more detailed debugging information. (This option may not be available on all systems.)	-Xtime	Any X. files greater or equal to time days old will be removed. The D. files are probably not present (if they were, the X. could get executed). But if there are D. files, they will be taken care of by D. processing (default 2 days).
-Ctime	Remove any C. files greater or equal to time days old and send appropriate information to the requester (default 7 days).																
-Dtime	Remove any D. files greater or equal to time days old, make an attempt to deliver mail messages, and execute rnews when appropriate (default 7 days).																
-mstring	Include string in the warning message generated by the -W option. The default line is "See your local administrator to locate the problem".																
-otime	Delete other files whose age is more than time days (default 2 days).																
-ssystem	Execute for system spool directory only.																
-Wtime	Any C. files equal to time days old will cause a mail message to be sent to the requester warning about the delay in contacting the remote. The message includes the JOBID, and in the case of mail, the mail message. The administrator may include a message line telling whom to call to check the problem (-m option) (default 1 day).																
-xdebug-level	Produce debugging output on standard output. debug-level is a single digit between 0 and 9; higher numbers give more detailed debugging information. (This option may not be available on all systems.)																
-Xtime	Any X. files greater or equal to time days old will be removed. The D. files are probably not present (if they were, the X. could get executed). But if there are D. files, they will be taken care of by D. processing (default 2 days).																

文件 /usr/lib/uucp directory with commands used by uucleanup internally
 /var/spool/uucp spool directory

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/uucp

另请参见 [uucp\(1C\)](#), [uux\(1C\)](#), [cron\(1M\)](#), [attributes\(5\)](#)

引用名	uusched – uucp file transport program scheduler				
用法概要	<code>/usr/lib/uucp/uusched [-u <i>debug-level</i>] [-x <i>debug-level</i>]</code>				
描述	uusched is the uucp(1C) file transport scheduler. It is usually started by the daemon <i>uudemon.hour</i> that is started by cron(1M) from an entry in user uucp's crontab file: <pre>11,41 * * * * /etc/uucp/uucp/uudemon.hour</pre>				
选项	The options are for debugging purposes only. <i>debug-level</i> is a number between 0 and 9. Higher numbers give more detailed debugging information: The following options are supported: <code>-u <i>debug-level</i></code> Passes the <code>-u <i>debug-level</i></code> option uucico(1M) as <code>-x <i>debug-level</i></code>. <code>-x <i>debug-level</i></code> Outputs debugging messages from uusched.				
文件	<code>/etc/uucp/Devices</code> <code>/etc/uucp/Permissions</code> <code>/etc/uucp/Systems</code> <code>/var/spool/locks/*</code> <code>/var/spool/uucp/*</code> <code>/var/spool/uucppublic/*</code>				
属性	See attributes(5) for descriptions of the following attributes: <table><thead><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr></thead><tbody><tr><td>Availability</td><td>service/network/uucp</td></tr></tbody></table>	ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	service/network/uucp
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	service/network/uucp				
另请参见	uucp(1C) , uustat(1C) , uux(1C) , cron(1M) , uucico(1M) , attributes(5)				

引用名	Uutry, uutry – attempt to contact remote system with debugging on						
用法概要	<code>/usr/lib/uucp/Uutry [-r] [-c type] [-x debug-level] system-name</code>						
描述	Uutry is a shell script that is used to invoke uucico(1M) to call a remote site. Debugging is initially turned on and is set to the default value of 5. The debugging output is put in file <code>/tmp/system-name</code> .						
选项	The following options are supported: <table><tr><td><code>-r</code></td><td>This option overrides the retry time that is set in file <code>/var/uucp/.Status/system-name</code>.</td></tr><tr><td><code>-c type</code></td><td>The first field in the Devices file is the "Type" field. The <code>-c</code> option forces <code>uucico</code> to use only entries in the "Type" field that match the user-specified type. The specified type is usually the name of a local area network.</td></tr><tr><td><code>-x debug-level</code></td><td><code>debug-level</code> is a number from 0 to 9. Higher numbers give more detailed debugging information.</td></tr></table>	<code>-r</code>	This option overrides the retry time that is set in file <code>/var/uucp/.Status/system-name</code> .	<code>-c type</code>	The first field in the Devices file is the "Type" field. The <code>-c</code> option forces <code>uucico</code> to use only entries in the "Type" field that match the user-specified type. The specified type is usually the name of a local area network.	<code>-x debug-level</code>	<code>debug-level</code> is a number from 0 to 9. Higher numbers give more detailed debugging information.
<code>-r</code>	This option overrides the retry time that is set in file <code>/var/uucp/.Status/system-name</code> .						
<code>-c type</code>	The first field in the Devices file is the "Type" field. The <code>-c</code> option forces <code>uucico</code> to use only entries in the "Type" field that match the user-specified type. The specified type is usually the name of a local area network.						
<code>-x debug-level</code>	<code>debug-level</code> is a number from 0 to 9. Higher numbers give more detailed debugging information.						
文件	<code>/etc/uucp/Devices</code> <code>/etc/uucp/Limits</code> <code>/etc/uucp/Permissions</code> <code>/etc/uucp/Systems</code> <code>/tmp/system-name</code> <code>/var/spool/locks/*</code> <code>/var/spool/uucp/*</code> <code>/var/spool/uucppublic/*</code>						
属性	See attributes(5) for descriptions of the following attributes:						

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/uucp

另请参见 [uucp\(1C\)](#), [uux\(1C\)](#), [uucico\(1M\)](#), [attributes\(5\)](#)

引用名	uuxqt – execute remote command requests
用法概要	<code>/usr/lib/uucp/uuxqt [-s <i>system</i>] [-x <i>debug-level</i>]</code>
描述	uuxqt is the program that executes remote job requests from remote systems generated by the use of the uux command. (<i>mail</i> uses uux for remote mail requests). uuxqt searches the spool directories looking for execution requests. For each request, uuxqt checks to see if all the required data files are available, accessible, and the requested commands are permitted for the requesting system. The <code>Permissions</code> file is used to validate file accessibility and command execution permission. There are two environment variables that are set before the uuxqt command is executed: ■ <code>UU_MACHINE</code> is the machine that sent the job (the previous one). ■ <code>UU_USER</code> is the user that sent the job. These can be used in writing commands that remote systems can execute to provide information, auditing, or restrictions.
选项	The following options are supported: <code>-s <i>system</i></code> Specifies the remote <i>system</i> name. <code>-x <i>debug-level</i></code> <i>debug-level</i> is a number from 0 to 9. Higher numbers give more detailed debugging information.
文件	<code>/etc/uucp/Limits</code> <code>/etc/uucp/Permissions</code> <code>/var/spool/locks/*</code> <code>/var/spool/uucp/*</code>
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/uucp

另请参见 [mail\(1\)](#), [uucp\(1C\)](#), [uustat\(1C\)](#), [uux\(1C\)](#), [uucico\(1M\)](#), [attributes\(5\)](#)

引用名 vbiosd – vbios daemon

用法概要 vbiosd

描述 The vbiosd daemon is started at boot time to provide an execution environment for adapter-supplied VBIOS code that is used in graphics mode switching.

This mode switching occurs when exiting X, switching virtual terminals (see [vtdaemon\(1M\)](#) for more information), and during fast reboot.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Interface Stability	Uncommitted

另请参见 [svcs\(1\)](#), [svcadm\(1M\)](#), [vtdaemon\(1M\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注 The vbios daemon is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/system/vbiosd:default`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	vdiskadm – 创建和管理虚拟磁盘
用法概要	<pre> vdiskadm create -s size [-t type[:opt],[opt] [-c comment] vdbname vdiskadm destroy [-r] vdbname snapshot vdiskadm snapshot vdbname@snapname vdiskadm rollback [-r] snapshot vdiskadm clone [-c comment] vdbname snapshot clone_vdbname vdiskadm move vdbname dir vdiskadm rename vdbname snapshot vdbname snapshot vdiskadm list [-fp]vdbname vdiskadm verify vdbname vdiskadm prop-get [-l] -p property vdbname vdiskadm prop-set -p property=value vdbname vdiskadm prop-add -p property=value vdbname vdiskadm prop-del -p property vdbname vdiskadm import [-fnpqm] [-x type] -d file zvol dsk [-t type[:opt]] vdbname vdiskadm export -x type[:opt] -d file zvol dsk vdbname vdiskadm convert [-t type[:opt]] vdbname vdiskadm translate [-i type[:opt]] -I input_file -x type[:opt] -d output_file vdiskadm help [command] </pre>
描述	vdiskadm 命令用于管理 dom0 中的虚拟磁盘。在上述用法概要格式中，vdbname 是虚拟磁盘的路径名；其最大长度为 MAXPATHLEN（1024 字节）。 vdiskadm 是作为一组子命令实现的，其中许多子命令具有自己的选项和操作数。下文中的“子命令”部分介绍了这些子命令。 下面的各个小节介绍了与虚拟磁盘相关的概念。
快照	快照是虚拟磁盘的只读副本。快照可以极快地创建，并且最初只占用一小部分空间。随着活动虚拟磁盘内数据的变化，快照将使用比与活动虚拟磁盘共享的数据更多的数据。
克隆	克隆是虚拟磁盘的可写入副本。缺省类型的克隆是原始虚拟磁盘的合并（接合）副本。例如，当虚拟磁盘包含数个快照时，就会出现合并克隆；后续克隆操作将生成不包含快照的新虚拟磁盘。克隆与原始虚拟磁盘属于相同的类

型（即，`vmdk:fixed`）。创建合并克隆后，不存在返回至原始虚拟磁盘或其任何快照的链接。因为没有链接，这使得合并克隆可以移动到另一台物理计算机。

数字值 在指定数字属性的值时可以使用适合人阅读的后缀（如 `k`、`KB`、`M`、`Gb` 等等，包括表示 `ZB` 的 `z`）。下面都是有效（且等效）的指定：

`1536M 1.5g 1.50GB`

虚拟磁盘的类型 下面是受支持的虚拟磁盘类型和选项：

- `vmdk:fixed`
- `vmdk:sparse`
- `vdi:fixed`
- `vdi:sparse`
- `vhd:fixed`
- `vhd:sparse`
- `raw:fixed`

其中，`vmdk` 是本机 VMware 格式，`vdi` 是本机 VirtualBox 格式，`vhd` 是本机 Microsoft 格式，`raw` 描述的是看起来像原始磁盘的文件。`raw` 磁盘始终是 `fixed` 格式，因此可以显式设置或隐式推断该选项。

如果未指定类型，则缺省值为 `vmdk`。如果未指定选项，则对于类型 `raw`，缺省值为 `fixed`，对于类型 `vmdk`、`vdi` 和 `vhd`，缺省值为 `sparse`。

本机属性和用户自定义属性 属性分为两种类型：本机属性和用户自定义属性。本机属性用于导出内部统计信息或控制 `vdiskadm` 行为。此外，本机属性分为可编辑属性和只读属性。用户自定义属性是不会对 `vdiskadm` 行为产生影响的任意字符串。您可以使用它们来注释虚拟磁盘，但应使用在您的环境中有意义的方式。用户自定义属性名称必须包含一个冒号（`:`）字符，以与本机属性相区别。

属性仅与虚拟磁盘关联，与各个快照无关。

每个虚拟磁盘都有一组本机属性，用于导出有关虚拟磁盘的统计信息并控制各种行为。

以下是虚拟磁盘的本机属性：

cdrom

布尔属性，如果虚拟磁盘是 CDROM，则其值为 `true`。

removable

布尔属性，如果虚拟磁盘是可移除介质，则其值为 `true`。

readonly

布尔属性，如果虚拟磁盘是只读的，则其值为 `true`。此属性为只读。

sectors

数字属性，包含给定虚拟磁盘中的磁盘扇区数。此属性为只读。

name

字符串属性，为虚拟磁盘的名称。此属性为只读。

max-size

数字属性，包含虚拟磁盘的最大大小（以字节为单位）。此属性为只读。

effective-size

数字属性，包含虚拟磁盘的有效大小（以字节为单位）。有效大小包括数据文件和所有快照的大小。有效大小可以超过最大大小。此属性为只读。

creation-time

字符串属性，包含虚拟磁盘的创建日期和时间。此属性为只读。

creation-time-epoch

数字属性，以自新纪元以来的秒数（自 1970 年 1 月 1 日 00:00:00 UTC 以来的秒数）描述 **creation-time** 属性。此属性为只读。

modification-time

字符串属性，包含上次修改虚拟磁盘的日期和时间。此属性为只读。

modification-time-epoch

数字属性，以自新纪元以来的秒数（自 1970 年 1 月 1 日 00:00:00 UTC 以来的秒数）描述 **modification-time** 属性。此属性为只读。

description

字符串属性，其中包含在创建或克隆虚拟磁盘时提供的注释。此属性为只读。

type

字符串属性，其中包含虚拟磁盘的类型：**vmdk**、**vhd**、**vdi** 或 **raw**。此属性为只读。

sparse

布尔属性，如果虚拟磁盘是 **sparse** 格式的，则其值为 **true**。此属性为只读。

owner

字符串属性，其中包含虚拟磁盘所有者的用户名。此属性是可编辑的。

子命令

下面的各个小节介绍了 **vdiskadm** 的各个子命令及其参数。

```
vdiskadm create      vdiskadm create -s size [-t type[:opt],[opt]]
                    [-c comment] vdbname
```

在 **vdbname** 指定的位置创建指定大小的新虚拟磁盘。如果 **vdbname** 包含虚拟磁盘的路径，则在创建虚拟磁盘期间将创建基于该路径的目录。此子命令具有下列选项。

-t type[:opt],[opt]

指定要创建的虚拟磁盘的类型。缺省类型为 **vmdk**。对于 **vmdk** 和 **vdi** 类型，缺省选项为 **sparse**。对于 **raw** 类型，缺省选项为 **fixed**。

-c comment

可以添加到虚拟磁盘的注释。

vdiskadm destroy	<code>vdiskadm destroy [-r] <i>vdname</i> <i>snapshot</i></code> 销毁指定的虚拟磁盘或快照。缺省情况下，如果指定的虚拟磁盘包含快照，则销毁操作将失败。此子命令具有下列选项。 -r 递归销毁虚拟磁盘，包括与该虚拟磁盘关联的所有快照。
vdiskadm snapshot	<code>vdiskadm snapshot <i>vdname</i>@<i>snapshotname</i></code> 使用指定的 <i>snapshotname</i> 创建虚拟磁盘的快照。此子命令没有选项。
vdiskadm rollback	<code>vdiskadm rollback [-r] <i>snapshot</i></code> 将虚拟磁盘回滚到以前的某个快照。回滚虚拟磁盘后，自创建快照以来更改过的所有数据都将丢弃，虚拟磁盘将恢复到创建快照时的状态。缺省情况下，该命令无法回滚到除最新快照以外的快照。要进一步回滚，必须通过指定 -r 选项销毁中间的所有快照。此子命令具有下列选项。 -r 递归销毁比指定的快照更新的快照。
vdiskadm clone	<code>vdiskadm [-c <i>comment</i>] <i>vdname</i> <i>snapshot</i> clone <i>vdname</i></code> 创建指定快照或虚拟磁盘的克隆。将使用被克隆的虚拟磁盘的类型、选项以及大小创建克隆。如果 <i>clone_vdname</i> 包含路径，则创建克隆虚拟磁盘期间将创建基于该路径的目录。缺省情况下，将创建合并克隆映像。此子命令具有下列选项。 -c <i>comment</i> 可以添加到克隆虚拟磁盘的注释。
vdiskadm move	<code>vdiskadm move <i>vdname</i> <i>dir</i></code> 将指定的虚拟磁盘移至指定目录。虚拟磁盘保留原来的名称。新目录必须存在。此子命令没有选项。
vdiskadm rename	<code>vdiskadm rename <i>vdname</i> <i>snapshot</i> <i>vdname</i> <i>snapshot</i></code> 重命名虚拟磁盘或快照。此子命令没有选项。
vdiskadm list	<code>vdiskadm list [-fp]<i>vdname</i></code> 列出指定的虚拟磁盘及其快照。此子命令具有下列选项。 -f 列出与虚拟磁盘关联的所有文件。该列表包括存储文件和范围。 -p 以易于解析的格式列出文件，在文件前加上标签 file:、snapshot: 或 store。

vdiskadm verify	<code>vdiskadm verify <i>vdname</i></code> 如果 Solaris 无法识别或打开虚拟磁盘，则会返回错误。此子命令没有选项。
vdiskadm prop-get	<code>vdiskadm prop-get [-l] -p <i>property</i> <i>vdname</i></code> 为指定的虚拟磁盘返回属性的值。属性值 <code>all</code> 将为虚拟磁盘显示所有本机属性和用户自定义属性。此子命令具有下列选项。 -l 给出更多属性信息，例如属性的可写入状态。 -p <i>property</i> 指定要查询的属性并显示该属性的值。对于属性 <code>all</code>，在属性值之前将显示属性的名称、一个冒号和一个空格。
vdiskadm prop-set	<code>vdiskadm prop-set -p <i>property=value</i> <i>vdname</i></code> 为指定的虚拟磁盘设置指定属性的值。<i>property</i> 可以是本机属性或用户自定义属性，但必须是可写的。用来更改使用 <code>prop-add</code> 子命令添加的属性的值。此子命令具有下列选项。 -p <i>property=value</i> 指定要设置的属性。
vdiskadm prop-add	<code>vdiskadm -p <i>property=value</i> <i>vdname</i></code> 将具有指定值的用户自定义属性添加到指定的虚拟磁盘。如果该属性已存在，则会返回错误。用户自定义属性名称必须包含一个冒号字符 (:)。此子命令具有下列选项。 -p <i>property=value</i> 指定要添加的属性。
vdiskadm prop-del	<code>vdiskadm prop-del -p <i>property</i> <i>vdname</i></code> 从指定的虚拟磁盘删除用户自定义属性。此子命令具有下列选项。 -p <i>property</i> 指定要删除的属性。
vdiskadm import	<code>vdiskadm import [-fnpqm] [-x <i>type</i>] -d <i>file zvol dsk \</i> <i>[-t <i>type[:opt]</i>] <i>vdname</i></i></code> 使用来自文件或块设备的数据新建虚拟磁盘。文件的格式可以是 <code>vdi</code>、<code>vhd</code>、<code>vmdk</code> 或 <code>raw</code>。块设备始终假设为原始格式。此子命令具有以下选项。 -f 返回将在导入进程中使用的文件的列表。 -n 显示来自导入的输出但不实际运行导入。 -p 以易于解析的格式显示文件。

- q
以无任何输出的静默模式运行。
- m
将导入的文件移至虚拟磁盘，不复制数据。
- x *type*
指定要导入的虚拟磁盘数据的类型。如果 `vdiskadm` 无法检测导入的文件类型，则必须指定 -x。
- d *file|zvol|dsk*
包含要导入数据的文件或块设备。
- t *type[:opt]*
指定将在导入时创建的虚拟磁盘的类型。缺省类型为 `vmdk`。对于 `vmdk`、`vdi` 和 `vhd` 类型，缺省的 *opt* 为 `sparse`。对于 `raw` 类型，缺省的 *opt* 为 `fixed`。

`vdiskadm export` `vdiskadm export -x type[:opt] -d file|zvol|dsk vname`

将数据从虚拟磁盘导出到文件或块设备。此子命令具有以下选项。

- x *type[:opt]*
指定要导出的虚拟磁盘数据的类型。
- d *file|zvol|dsk*
接收要导出数据的文件或块设备。

`vdiskadm convert` `vdiskadm convert [-t type[:opt]] vname`

将虚拟磁盘转换为其他类型的虚拟磁盘。此子命令具有以下选项。

- t *type[:opt]*
指定将在转换时创建的虚拟磁盘的类型。缺省类型为 `vmdk`。对于 `vmdk`、`vdi` 和 `vhd` 类型，缺省的 *opt* 为 `sparse`。对于 `raw` 类型，缺省的 *opt* 为 `fixed`。

`vdiskadm translate` `vdiskadm translate [-i type[:opt]] -I input_file \`
 `-x type[:opt] -d output_file`

将数据从一种虚拟磁盘数据类型转换为另一种，但不创建虚拟磁盘。此子命令具有以下选项。

- i *type[:opt]*
指定要转换的虚拟磁盘数据的输入类型。如果 `vdiskadm` 无法检测输入文件类型，则必须指定 -i。
- I *input_file*
包含要转换数据的文件或块设备。
- x *type[:opt]*
指定要转换的虚拟磁盘数据的输出类型。对于 `vmdk`、`vdi` 和 `vhd` 类型，缺省的 *opt* 为 `sparse`。对于 `raw` 类型，缺省的 *opt* 为 `fixed`。

-d *output_file*

接收要转换数据的文件或块设备。*output_file* 不能已经存在。将在转换期间创建该文件。

vdiskadm help

vdiskadm help [*command*]

显示常规的或特定于命令的帮助消息。此子命令只能包含命令名称这一可选参数。

示例

示例1 创建 vmdk 稀疏文件

以下命令将在目录 `/guests/disks` 中创建一个大小为 8 GB、名为 `disk1` 的虚拟磁盘。

```
# vdiskadm create -s 8g -t vmdk:sparse /guests/disks/disk1
```

示例2 创建快照

以下命令将为 `/guests/disks/disk1` 中的虚拟磁盘创建快照。快照名为 `install`。

```
# vdiskadm snapshot /guests/disks/disk1@install
```

示例3 创建和销毁快照

以下命令将为 `/guests/disks/disk1` 中的虚拟磁盘创建两个快照（名为 `install` 和 `bfu`）。第三个命令将销毁新创建的快照 `install`。

```
# vdiskadm snapshot /guests/disks/disk1@install
# vdiskadm snapshot /guests/disks/disk1@bfu
# vdiskadm destroy /guests/disks/disk1@install
```

示例4 回滚虚拟磁盘

以下命令将虚拟磁盘的内容恢复到名为 `install` 的快照，删除中间的所有快照。

```
# vdiskadm rollback -r /guests/disks/disk1@install
```

示例5 列出虚拟磁盘和快照

以下命令将列出与虚拟磁盘 `/guests/disks/disk1` 关联的所有映像。

```
# vdiskadm list /guests/disks/disk1
disk1@install
disk1@bfu
disk1
```

示例6 创建克隆

以下命令将创建一个新的虚拟磁盘，该虚拟磁盘是虚拟磁盘 `/guests/disks/disk1` 的接合副本。该克隆是以与原始虚拟磁盘相同的格式（即 `vmdk:sparse`）创建的。

```
# vdiskadm clone /guests/disks/disk1 /guests/clone/clone_disk1
```

示例7 添加用户自定义属性

以下命令将向虚拟磁盘添加用户自定义属性并为其赋予指定值。可选择此属性来提供此虚拟磁盘数据的来源和要求，使用必需的冒号来分隔字段。

```
# vdiskadm prop-add -p com.sun:required-nic=2 /guests/disks/disk1
```

示例8 导入现有的 vmdk 格式文件

以下命令接受现有的 vmdk 格式文件并将其导入到虚拟磁盘。

```
# vdiskadm import -d /downloads/appliance.vmdk /guests/import/disk1
```

示例9 导入 vmdk 文件并将其转换为 vhd

以下命令接受现有的 vmdk 格式文件，并在导入时将其转换为 vhd 类型的虚拟磁盘。

```
# vdiskadm import -d /downloads/appliance.vmdk -t vhd \  
/guests/import/disk1
```

示例10 从 zvol 导入数据

以下命令将从 zvol 导入虚拟磁盘数据，并在导入时将其转换为 vmdk:fixed 类型的虚拟磁盘。

```
# vdiskadm import -d /dev/zvol/dsk/pool1/disk1 -t vmdk:fixed \  
/guests/import/disk1
```

示例11 将数据导出到块设备

以下命令接受现有的虚拟磁盘，并在导出时将其转换为 raw 类型的磁盘分片。

```
# vdiskadm export -d /dev/dsk/c0t1d0s3 -x raw /guests/disks/disk1
```

示例12 转换虚拟磁盘类型

以下命令接受现有的虚拟磁盘并将其（就地）转换为其他格式类型。

```
# vdiskadm convert -t vdi:fixed /guests/disks/disk1
```

示例13 将数据从一种类型转换为另一种类型

以下命令可在不创建虚拟磁盘的情况下，将虚拟磁盘格式文件中的数据转换为写入 zvol 的原始数据。

```
# vdiskadm translate -I /downloads/appliance.vmdk -x raw \  
-d /dev/zvol/dsk/pool1/disk1
```

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/storage/vdiskadm

属性类型	属性值
接口稳定性	Uncommitted（未确定）

另请参见

[attributes\(5\)](#)

引用名	virt-convert – convert virtual machines between formats
用法概要	virt-convert [<i>option</i>]... <i>input.vmx</i> <i>input.ovf</i> <i>input-dir</i> [<i>output.xml</i> <i>output-dir</i>]
描述	The <code>virt-convert</code> program is a command line tool for converting virtual machines (VMs) from one format to another. It accepts either a VM definition file (such as VMware <code>vmx</code> format) or a directory containing a VM. By default, a new VM definition file, and converted disk images, will be placed in a new output directory. If an output directory is specified, it will be created if necessary, and the output VM definition placed within the new directory, along with any disk images, as needed. If an output VM definition file is specified, it will be created alongside any disks in the same directory.
选项	Any of the following options can be omitted, in which case <code>virt-convert</code> will use defaults when required. An input VM definition or containing directory must be provided. By default, an output directory is generated based upon the name of the VM. The default input format is VMware <code>vmx</code>, and the default output format is a <code>libvirt</code> “image” XML definition. <code>-a, --arch=<i>arch</i></code> Architecture of the virtual machine (<code>i686</code>, <code>x86_64</code>, <code>ppc</code>). Defaults to that of the host machine. <code>-D, --disk-format=<i>format</i></code> Output disk format, or none if no conversion should be performed. <i>format</i> is one of: none No disks are converted or copied. <code>vmdk</code> VMWare VMDK format <code>raw</code> raw file <code>vdisk</code> vdisk format (see vdiskadm(1M)) <code>-d, --debug</code> Display debugging information. <code>-h, --help</code> Display the help message and exit. <code>-i, --input-format <i>format</i></code> Input format. Currently, <code>vmx</code> and <code>ovf</code> are supported. <code>--noacpi</code> Override the OS type and variant to disable the ACPI setting for fully virtualized guest.

- noapic
Override the OS type and variant to disable the APIC setting for fully virtualized guest.
- o, --output-format *format*
Output format. Currently, the supported output formats are *virt-instance* and *virt-image*. *virt-instance* is the recommended format for Solaris.
- os-type=*os_type*
Optimize the guest configuration for a type of operating system. This will attempt to pick the most suitable ACPI and APIC settings, optimally supported mouse drivers, and generally accommodate other operating system quirks.
- os-variant=*os_variant*
Further optimize the guest configuration for a specific operating system variant. This parameter is optional.
- p, --paravirt
Create a paravirtualized guest image. Convert machine to a paravirtualized Xen-based image.
- q, --quiet
Avoid verbose output.
- v, --hvm
Create a fully virtualized guest image. Convert machine to a hvm/qemu-based image (this is the default if --paravirt is not specified).

示例

示例1 Converting a VMware VMX appliance
The following sequence of commands converts a VMware VMX appliance and imports it into Solaris xVM.

```
# virt-convert -o virt-instance /guests/vmx-appliance/ \  
/guests/xvm-appliance/  
# virsh define --relative-path=/guests/xvm-appliance/ \  
\guests/xvm-appliance/appliance.xml
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/xvm/header-xvm
Interface Stability	Volatile

另请参见

[vdiskadm\(1M\)](#), [attributes\(5\)](#)

Caveats Not all conversions will result in a working guest installation. If the source OS image is configured to use SCSI disks, the use of IDE disks may cause the OS boot to fail. Some images may be configured to use the VMware drivers such as `vlance`. In the `vlance` case and in general, device emulation support may not be sufficient for all OS installations.

Authors Written by Joey Boggs and John Levon.

See the `AUTHORS` file in the source distribution for the complete list of credits.

引用名	virtinfo – virtualization domain information												
用法概要	virtinfo [-acdpstu]												
描述	The <code>virtinfo</code> utility provides virtualization information returned by <code>libv12n(3LIB)</code> about the current domain. When options are specified, symbols that represent one or more virtualization domain characteristics are written to standard output. If no options are specified, <code>virtinfo</code> shows only the virtualization domain type.												
选项	The <code>virtinfo</code> command supports the following options: -a Shows all information that is currently available from the system. -c Shows the network node name of the control domain. This name is identical to the string shown by running the <code>uname -n</code> command on the control domain. -d Shows the domain name for this virtual domain. -p Shows information in a parseable format. The parseable output for Logical Domains (LDoms) has the following format: <pre>VERSION 1.0 DOMAINROLE impl=LDoms control={true false} io={true false} service={true false} root={true false} DOMAINNAME name=domain-name DOMAINUUID uuid=uuid DOMAINCONTROL name=control-domain-nodename DOMAINCASSIS serialno=serial-no</pre> -s Shows the platform serial number. -t Shows the domain type. The domain type for Logical Domains is a string of one or more of the following blank-separated values: <table><tr><td>LDoms</td><td>LDoms implementation</td></tr><tr><td>control</td><td>Control domain</td></tr><tr><td>guest</td><td>Guest domain</td></tr><tr><td>I/O</td><td>I/O domain</td></tr><tr><td>service</td><td>Service domain</td></tr><tr><td>root</td><td>Root I/O domain</td></tr></table> Note that the <code>control</code> and <code>guest</code> values are mutually exclusive.	LDoms	LDoms implementation	control	Control domain	guest	Guest domain	I/O	I/O domain	service	Service domain	root	Root I/O domain
LDoms	LDoms implementation												
control	Control domain												
guest	Guest domain												
I/O	I/O domain												
service	Service domain												
root	Root I/O domain												

-u Shows the universally unique identifier (UUID) of the domain. See the [libuuid\(3LIB\)](#) man page.

退出状态

The following exit values are returned:

- 0 Successful completion.
- >0 An error occurred.

示例

示例 1 Viewing Default Information

The following example shows the default information about the domain:

```
$ virtinfo
Domain role: LDoms control I/O service root
```

The current domain is the control domain, which is also an I/O domain, the service domain and a root I/O domain.

示例 2 Viewing All Information

The following example shows all available information about the domain:

```
$ virtinfo -a
Domain role: LDoms control I/O service root
Domain name: primary
Domain UUID: 8e0d6ec5-cd55-e57f-ae9f-b4cc050999a4
Control domain: san-t2k-6
Chassis serial#: 0704RB0280
```

示例 3 Viewing All Information In a Parsable Format

The following example shows all available information about the domain in a parsable format:

```
$ virtinfo -ap
VERSION 1.0
DOMAINROLE|impl=LDoms|control=true|io=true|service=true|root=true
DOMAINNAME|name=primary
DOMAINUUID|uuid=8e0d6ec5-cd55-e57f-ae9f-b4cc050999a4
DOMAINCONTROL|name=san-t2k-6
DOMAINCHASSIS|serialno=0704RB0280
```

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Uncommitted

另请参见

[libuuid\(3LIB\)](#), [libv12n\(3LIB\)](#)

引用名

vmstat – report virtual memory statistics

用法概要

vmstat [-ipqsS] [-T u | d] [*disks*] [*interval*] [*count*]

描述

vmstat reports virtual memory statistics regarding kernel thread, virtual memory, disk, trap, and CPU activity.

On MP (multi-processor) systems, vmstat averages the number of CPUs into the output. For per-processor statistics, see [mpstat\(1M\)](#).

vmstat only supports statistics for certain devices. For more general system statistics, use [sar\(1\)](#), [iostat\(1M\)](#), or [sar\(1M\)](#).

Without options, vmstat displays a one-line summary of the virtual memory activity since the system was booted.

During execution of the kernel status command, the state of the system can change. If relevant, a state change message is included in the vmstat output, in one of the following forms:

```
<<device added: sd0>>
<<device removed: sd0>>
<<processors added: 1, 3>>
<<processors removed: 1, 3>>
```

See 《[Oracle Solaris 管理：常见任务](#)》 for device naming conventions for disks.

选项

The following options are supported:

- i Report the number of interrupts per device. *count* and *interval* does not apply to the -i option.
- p Report paging activity in details. This option will display the following, respectively:
 - epi Executable page-ins.
 - epo Executable page-outs.
 - epf Executable page-frees.
 - api Anonymous page-ins.
 - apo Anonymous page-outs.
 - apf Anonymous page-frees.
 - fpi File system page-ins.
 - fpo File system page-outs.
 - fpf File system page-frees.

- When executed in a zone and if the pools facility is active, all of the above only report activity on the processors in the processor set of the zone's pool.
- q Suppress messages related to state changes.
 - s Display the total number of various system events since boot. *count* and *interval* does not apply to the -s option.
 - S Report on swapping rather than paging activity. This option will change two fields in vmstat's "paging" display: rather than the "re" and "mf" fields, vmstat will report "si" (swap-ins) and "so" (swap-outs).
 - T u | d Specify u for a printed representation of the internal representation of time. See [time\(2\)](#). Specify d for standard date format. See [date\(1\)](#).

操作数

The following operands are supported:

- count* Specifies the number of times that the statistics are repeated. *count* does not apply to the -i and -s options.
- disks* Specifies which disks are to be given priority in the output (only four disks fit on a line). Common disk names are id, sd, xd, or xy, followed by a number (for example, sd2, xd0, and so forth).
- interval* Specifies the last number of seconds over which vmstat summarizes activity. This number of seconds repeats forever. *interval* does not apply to the -i and -s options.

示例

示例1 Using vmstat

The following command displays a summary of what the system is doing every five seconds.

example% **vmstat 5**

```
kthr  memory          page          disk          faults          cpu
r  b  w  swap  free re  mf  pi  p  fr  de  sr  s0  s1  s2  s3  in  sy  cs  us  sy  id
0  0  0  11456  4120 1  41 19 1  3  0  2  0  4  0  0  48 112 130  4 14 82
0  0  1  10132  4280 0   4 44 0  0  0  0  0 23  0  0 211 230 144  3 35 62
0  0  1  10132  4616 0   0 20 0  0  0  0  0 19  0  0 150 172 146  3 33 64
0  0  1  10132  5292 0   0  9 0  0  0  0  0 21  0  0 165 105 130  1 21 78
1  1  1  10132  5496 0   0  5 0  0  0  0  0 23  0  0 183  92 134  1 20 79
1  0  1  10132  5564 0   0 25 0  0  0  0  0 18  0  0 131 231 116  4 34 62
1  0  1  10124  5412 0   0 37 0  0  0  0  0 22  0  0 166 179 118  1 33 67
1  0  1  10124  5236 0   0 24 0  0  0  0  0 14  0  0 109 243 113  4 56 39
^C
```

example%

The fields of vmstat's display are

示例 1 Using vmstat (续)

kthr	Report the number of kernel threads in each of the three following states: r the number of kernel threads in run queue b the number of blocked kernel threads that are waiting for resources I/O, paging, and so forth w the number of swapped out lightweight processes (LWPs) that are waiting for processing resources to finish.
memory	Report on usage of virtual and real memory. swap available swap space (Kbytes) free size of the free list (Kbytes)
page	Report information about page faults and paging activity. The information on each of the following activities is given in units per second. re page reclaims — but see the -S option for how this field is modified. mf minor faults — but see the -S option for how this field is modified. pi kilobytes paged in po kilobytes paged out fr kilobytes freed de anticipated short-term memory shortfall (Kbytes) sr pages scanned by clock algorithm When executed in a zone and if the pools facility is active, all of the above (except for “de”) only report activity on the processors in the processor set of the zone's pool.
disk	Report the number of disk operations per second. There are slots for up to four disks, labeled with a single letter and number. The letter indicates the type of disk (s = SCSI, i = IPI, and so forth); the number is the logical unit number.
faults	Report the trap/interrupt rates (per second). in interrupts sy system calls cs CPU context switches When executed in a zone and if the pools facility is active, all of the above only report activity on the processors in the processor set of the zone's pool.

示例 1 Using vmstat (续)

cpu Give a breakdown of percentage usage of CPU time. On MP systems, this is an average across all processors.

us user time

sy system time

id idle time

When executed in a zone and if the pools facility is active, all of the above only report activity on the processors in the processor set of the zone's pool.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	See below.

Invocation is evolving. Human readable output is unstable.

另请参见

[date\(1\)](#), [sar\(1\)](#), [iostat\(1M\)](#), [mpstat\(1M\)](#), [sar\(1M\)](#), [time\(2\)](#), [attributes\(5\)](#)

附注

The sum of CPU utilization might vary slightly from 100 because of rounding errors in the production of a percentage figure.

引用名	vntsd – 用于 Logical Domains 的虚拟网络终端服务器守护进程
用法概要	/usr/lib/ldoms/vntsd
描述	vntsd 守护进程是一个通过使用 telnet(1) 支持到 Logical Domains (LDoms) 控制台的连接的服务器。当 telnet 会话启动时，vntsd 会向客户机发送 telnet 选项来表明要远程回显字符并抑制前进的意愿。 这些控制台由 LDoms Manager 分为若干个组。每个控制台组都分配有唯一的组名和 TCP 端口号。vntsd 使用组的端口号来导出对该组内的控制台的访问权限。要与某个控制台或控制台组建立连接，用户需要使用对应组的端口号启动 telnet(1) 会话。根据该组内控制台的数量，vntsd 将执行下面的两个操作之一： ■ 如果组内只有一个控制台，则 vntsd 将会话连接至该 LDoms 控制台。 ■ 如果组内有多个控制台，则 vntsd 将提示用户选择要连接到的控制台，如下面的“多控制台选项”部分所示。 对于每个控制台，vntsd 仅向连接到该控制台的第一个用户提供写访问权限。之后连接到该控制台的用户只允许从控制台进行读取，并等待写访问权限。当第一个用户断开连接后，写特权将移交到在队列中等待的下一个用户。如果没有写特权的用户尝试向控制台进行写入，则 vntsd 将显示以下消息： <pre>You do not have write access</pre> 没有写访问权限的用户可以通过使用 <code>~w</code> 特殊控制台命令（在下面的“特殊控制台命令”部分中介绍）强行获取写访问权限。 vntsd 只可通过超级用户特权调用，或者由具有主管管理员角色的用户调用。
选项	vntsd 的选项分为多控制台选项和控制台命令。
多控制台选项	当一个组中有多个 LDoms 控制台时，支持下面列出的选项。这些选项的使用语法为： <pre><hostname>-vnts-<group-name>: <option></pre> 例如： <pre>myhost-vnts-salesgroup: h</pre> h 选项调用帮助，如下所述。 <pre>h 显示以下帮助文本： h -- this help l -- list of consoles q -- quit c{id}, n{name} -- connect to console of domain {id} or domain name</pre> l 列出组内的所有控制台。例如：

DOMAIN ID	DOMAIN NAME	DOMAIN STATE
0	ldg1	online
1	ldg2	connected
...	...	...

两个域状态及其含义为：

online (联机)

无任何用户连接至该控制台。

connected (已连接)

至少一个用户已经连接至该控制台。

q

从 vntsd 断开连接。

c{id}, n{name}

连接到指定的控制台。连接时，将显示以下消息：

```
Connecting to console <domain-name> in group <group-name>
Press ~? for control options ....
```

特殊控制台命令

作为行的第一个字符出现的波浪号 (~) 是一个转义信号，指示 vntsd 执行特殊控制台命令。双波浪号 (~~) 序列输出一个波浪号。与开头的波浪号结合使用时，vntsd 接受以下特殊控制台命令：

~.

从控制台或控制台组断开连接。

~w

强制获取对控制台的写访问权限。

~p

从该控制台断开连接，并连接到控制台列表中紧邻该控制台的前一个控制台。

~n

从该控制台断开连接，并连接到控制台列表中紧邻该控制台的下一个控制台。

~#

发送中断信号。

~^B

发送替代中断信号。

~?

显示 vntsd 帮助，如下所示：

```
~# - Send break
~^B - Send alternate break
~. - Exit from this console
~w - Force write access
~n - Console next
```

~p - Console previous
~? - Help

文件 /usr/lib/ldoms/vntsd
二进制可执行 vntsd 文件。

/usr/lib/ldoms/vntsd.xml
vntsd 的服务管理工具 (smf(5)) 清单文件。

属性 有关以下属性的说明，请参见 attributes(5)：

属性类型	属性值
可用性	system/ldoms
接口稳定性	Committed（已确定）

另请参见 telnet(1)、svccfg(1M)、usermod(1M)、auth_attr(4)、attributes(5)、smf(5)

附注 vntsd 由服务管理工具 smf(5) 管理，其服务标识符为：

svc:/ldoms/vntsd

可以使用 svccfg(1M) 命令更改以下属性：

vntsd/vcc_device
设置 vntsd 将连接到的虚拟控制台集中器 (virtual console concentrator, vcc) 驱动程序的一个实例。

vntsd/listen_addr
使用以下语法设置 vntsd 将侦听的 IP 地址：

vntsd/listen_addr:"xxx.xxx.xxx.xxx"

...其中 xxx.xxx.xxx.xxx 是一个有效的 IP 地址。该属性的缺省值设置是侦听 IP 地址 127.0.0.1。如果该值设置为控制域的 IP 地址，则用户可以通过网络连接到客控制台。

注-启用对控制台的网络访问有安全问题。任何用户都将能够连接到控制台，因此，缺省情况下它被禁用。

vntsd/timeout_minutes
设置超时（以分钟为单位）。如果控制台上没有任何活动（输入或输出），vntsd 将使telnet 连接超时（关闭）。缺省值为 0，即禁用超时。

vntsd/authorization
为正被访问的域控制台启用用户和角色的授权检查。该属性的缺省值为 false 以维护向后兼容性。要启用授权检查，可使用 svccfg(1M) 命令将该属性的值设置为 true。启用该选项后，vntsd 将在 localhost 上侦听并接受连接。当启用了该选项时，如果 listen_addr 属性指定了替代 IP 地址，则 vntsd 将忽略替代 IP 地址并继续

在 `localhost` 上侦听。从其他主机启动的连接也将失败。授权可用于访问所有控制台或控制台组，或访问特定控制台或控制台组。当启用 `vntsd` 服务时，以下授权会添加到授权描述数据库 `auth_attr(4)`：

```
solaris.vntsd.consoles:::Access All LDoms Guest Consoles::
```

添加基于控制台组名称的细粒度授权。例如，如果要被授权的控制台组的名称为 `ldg1`，则会将以下条目添加到 `auth_attr(4)` 文件中：

```
solaris.vntsd.console-ldg1:::Access Specific LDoms Guest Console::
```

缺省情况下，会为 `root` 用户或角色分配可访问所有控制台的授权。主管人员（超级用户）可以使用 `usermod(1M)` 命令为其他用户或角色分配所需的授权。

以下示例向用户 `user1` 授予可访问所有域控制台的授权。

```
# usermod -A "solaris.vntsd.consoles" user1
```

以下示例向用户 `user1` 授予可访问名为 `ldg1` 的控制台组的授权：

```
# usermod -A "solaris.vntsd.console-ldg1" user1
```

引用名	volcopy – make an image copy of file system	
用法概要	volcopy [-F <i>FSType</i>] [-V] [<i>generic_options</i>] [-o <i>FSType-specific_options</i>] <i>operands</i>	
描述	volcopy makes a literal copy of the file system. This command may not be supported for all <i>FSTypes</i> .	
选项	The following options are supported:	
	-F <i>FSType</i>	Specify the <i>FSType</i> on which to operate. The <i>FSType</i> should either be specified here or be determinable from <code>/etc/vfstab</code> by matching the <i>operands</i> with an entry in the table. Otherwise, the default file system type specified in <code>/etc/default/fs</code> will be used.
	-V	Echo the complete command line, but do not execute the command. The command line is generated by using the options and arguments provided by the user and adding to them information derived from <code>/etc/vfstab</code> . This option should be used to verify and validate the command line.
	<i>generic_options</i>	Options that are commonly supported by most <i>FSType</i> -specific command modules. The following options are available: -a Require the operator to respond “yes” or “no” instead of simply waiting ten seconds before the copy is made. -s (Default) Invoke the DEL if wrong verification sequence.
	-o <i>FSType-specific_options</i>	Specify <i>FSType</i> -specific options in a comma separated (without spaces) list of suboptions and keyword-attribute pairs for interpretation by the <i>FSType</i> -specific module of the command.
操作数	The following operands are supported:	
	<i>operands</i>	generally include the device and volume names and are file system specific. A detailed description of the <i>operands</i> can be found on the <i>FSType</i> -specific man pages of volcopy.
退出状态	The following exit values are returned:	
	0	Successful file system copy
	1	An error has occurred.

文件

/etc/vfstab

list of default parameters for each file system

/etc/default/fs

default local file system type. Default values can be set for the following flags in /etc/default/fs. For example: LOCAL=ufs.

LOCAL:

The default partition for a command if no *FSType* is specified.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	system/core-os

另请参见

[labelit\(1M\)](#), [vfstab\(4\)](#), [attributes\(5\)](#) Manual pages for the FSType-specific modules of volcopy.

引用名	volcopy_ufs – make an image copy of a ufs file system				
用法概要	volcopy [-F ufs] [<i>generic_options</i>] <i>fsname srcdevice volname1 destdevice volname2</i>				
描述	volcopy makes a literal copy of the ufs file system using a blocksize matched to the device.				
选项	The following option is supported: <i>generic_options</i> options supported by the generic volcopy command. See volcopy(1M).				
操作数	The following operands are supported: <i>fsname</i> represents the mount point (for example, root, u1, etc.) of the file system being copied. <i>srcdevice</i> or <i>destdevice</i> the disk partition specified using the raw device (for example, /dev/rdisk/cld0s8, /dev/rdisk/cld1s8, etc.). <i>srcdevice</i> and <i>volname1</i> the device and physical volume from which the copy of the file system is being extracted. <i>destdevice</i> and <i>volname2</i> the target device and physical volume. <i>fsname</i> and <i>volname</i> are limited to six or fewer characters and recorded in the superblock. <i>volname</i> may be ‘-’ to use the existing volume name.				
退出状态	The following exit values are returned: 0 Successful file system copy. non-zero An error has occurred.				
文件	/var/adm/filesave.log a record of file systems/volumes copied				
属性	See attributes(5) for descriptions of the following attributes:				
<table><tr><th>ATTRIBUTE TYPE</th><th>ATTRIBUTE VALUE</th></tr><tr><td>Availability</td><td>system/core-os</td></tr></table>		ATTRIBUTE TYPE	ATTRIBUTE VALUE	Availability	system/core-os
ATTRIBUTE TYPE	ATTRIBUTE VALUE				
Availability	system/core-os				
另请参见	cpio(1) , dd(1M) , labelit(1M) , volcopy(1M) , attributes(5) , ufs(7FS)				
附注	volcopy does not support copying to tape devices. Use dd(1M) for copying to and from tape devices.				

引用名	vrrpadm – VRRP 管理工具
用法概要	<pre> vrrpadm create-router -V vrid -l link -A inet inet6 [-p priority] [-i adv_interval] [-o flags] router_name vrrpadm delete-router router_name vrrpadm disable-router router_name vrrpadm enable-router router_name vrrpadm modify-router [-p priority] [-i adv_interval] [-o flags] [router_name] vrrpadm show-router [-P -x] [-p] [-o field[,...]] [router_name] </pre>
描述	vrrpadm 命令用来管理系统中的 VRRP（Virtual Router Redundancy Protocol，虚拟路由器冗余协议）服务。 VRRP 指定一种选择协议，用以将对某个虚拟路由器的职责动态分配到 LAN 内的 VRRP 路由器之一。在任一给定时刻，只有一个 VRRP 路由器控制着与某个虚拟路由器关联的 IPv4 或 IPv6 虚拟地址（该路由器称为主路由器），并转发发送到这些 IP 地址的包。如果主路由器变得不可用，则选择进程可以动态地对转发职责进行故障转移。 每个 vrrpadm 子命令对由管理员给定的名称标识的 VRRP 路由器执行操作。同一 LAN 内具有相同的 VRID 和地址族的 VRRP 路由器构成了一个虚拟路由器，它保护着一组虚拟 IP 地址。 一个系统可以有多个 VRRP 路由器，每个属于一个不同的虚拟路由器。
子命令	所支持的子命令如下。请注意，所有子命令（show-router 除外）都需要 solaris.network.vrrp 授权。show-router 子命令不需要特殊授权。 <pre> vrrpadm create-router -V vrid -l link -A inet inet6 [-p priority] [-i adv_interval] [-o flags] router_name </pre> 创建具有指定配置的 VRRP 路由器。 -A inet inet6、--address_family=inet inet6 地址族。可以是 IPv4 或 IPv6。 -i adv_interval、--adv_interval=adv_interval 通告时间间隔（以毫秒为单位）。缺省值为 1000（一秒）。有效的时间间隔范围为 10-40950。 -l link、--link=link 在其上配置 VRRP 路由器的数据链路。这决定了该 VRRP 路由器在哪个 LAN 中运行。该数据链路可以是物理链路、VLAN 或聚合。 -o flags、--flags=flags 抢先模式和接受模式，以逗号分隔。值可以是： ■ preempt

- un_preempt
- accept
- noaccept

缺省情况下，两种模式均设为 true。

preempt 模式控制已启用的较高优先级备份路由器是否可以抢先于较低优先级主路由器。如果 **preempt** 模式为 true，则允许抢先；否则，禁止抢先。请注意，如果 VRRP 路由器是虚拟 IP 地址的所有者，则 **preempt** 模式必须为 true。

accept 模式控制虚拟 IP 地址的本地包接受情况。如果 **accept** 模式为 true，主路由器必须接受发送到虚拟 IP 地址的包。如果 **accept** 模式为 false，则主路由器不接受这些包，尽管它会为这些不被接受的虚拟 IP 地址对 ARP 请求或 ND 请求和通告进行响应。它还必须为在此子命令中指定的路由器转发包。请注意，如果 VRRP 路由器是虚拟 IP 地址的所有者，则 **accept** 模式必须为 true。此选项的语法示例：

```
-o preempt,no_accept
```

-p priority、**--priority=priority**

指定的 VRRP 路由器的优先级（用于主路由器选择）。该值越高，路由器被选作主路由器的可能性越大。

缺省值为 255，表示指定的 VRRP 路由器是 IP 地址所有者并拥有所有虚拟 IP 地址。IP 地址所有者将响应为进行 ICMP 强制回应、TCP 连接等等而发送到某个虚拟 IP 地址的包。

范围 1-254 供充当虚拟路由器的备份的 VRRP 路由器使用。选择主路由器时将优先选择具有较高优先级的 VRRP 路由器。

-V vrid、**--VRID=vrid**

虚拟路由器标识符 (VRID)。它与地址族共同标识 LAN 内的虚拟路由器。

router_name

VRRP 路由器的名称。在其他 **vrrpadm** 子命令中，该名称用来标识 VRRP 路由器。

有效路由器名称的最大长度为 31 个字符。合法字符包括字母数字（a-z、A-Z、0-9）和下划线（'_'）。

vrrpadm delete-router router_name

删除由 **router_name** 标识的 VRRP 路由器。

vrrpadm disable-router router_name

禁用由 **router_name** 标识的虚拟路由器。路由器被禁用后，它将不再参与主虚拟路由器的选择过程。

vrrpadm enable-router router_name

重新启用由 **router_name** 标识的之前被禁用的虚拟路由器。该路由器将会重新参与主虚拟路由器的选择过程。

```
vrrpadm modify-router [-p priority] [-i adv_interval] [-o flags] [router_name]
```

修改由 *router_name* 标识的 VRRP 路由器的配置。仅可修改优先级、通告时间间隔、抢先模式和接受模式。

<code>-p priority</code> 、 <code>--priority=priority</code>	该 VRRP 路由器的新优先级。
<code>-i adv_interval</code> 、 <code>--adv_interval=adv_interval</code>	新的通告时间间隔。
<code>-o flags</code> 、 <code>--flags=flags</code>	新的抢先模式和接受模式。可以指定一个，也可以两者都指定。如果二者都指定，需用逗号分隔它们。例如：
<code>-o preempt,no_accept</code>	

```
vrrpadm show-router [-P | -x] [-p] [-o field,...] [router_name]
```

显示由 *router_name* 标识的 VRRP 路由器的信息。如果未指定 *router_name*，则会显示系统上所有 VRRP 路由器的信息。

缺省情况下（无选项），显示以下字段：

NAME

VRRP 路由器的名称。

VRID

VRRP 路由器的 VRID。

LINK

在其上创建 VRRP 路由器的数据链路。

AF

VRRP 路由器的地址族，可以是 IPv4 或 IPv6。

PRIO

该 VRRP 路由器的优先级（用于主路由器选择）。

ADV_INTV

通告时间间隔（以毫秒为单位）。

STATE

VRRP 路由器的当前状态：INIT（初始化）、BACK（备份）或 MAST（主）。

MODE

与 VRRP 路由器关联的一组标志。可能的值为：

- e 路由器已启用。
- p Preempt（抢先）模式为 true。
- a Accept（接受）模式为 true。
- o 虚拟地址所有者。

VNIC

为该 VRRP 路由器创建的 VRRP VNIC。

请注意，VNIC 的名称可能会随时间推移而改变，除非该路由器已启用。

`show-router` 子命令具有以下选项：

-x、--extended

显示给定 VRRP 路由器的其他信息：

PRIMARY_IP

由 VRRP 路由器选择的主 IP 地址。

VIRTUAL_IPS

在 VRRP 路由器上配置的虚拟 IP 地址。

PRV_STAT

VRRP 路由器的前一个状态。

STAT_LAST

自上次状态转换以来已过去的时间。

-P、--peer

显示备份 VRRP 路由器的信息。仅当 VRRP 路由器处于备份状态时此选项才有意义。

将显示以下字段：

NAME

VRRP 路由器的名称。

PEER

对等 VRRP 路由器的主 IP 地址。

P_PRIO

接收自对等方的通告中包含的对等 VRRP 路由器的优先级。

P_INTV

接收自对等方的通告中包含的通告时间间隔（以毫秒为单位）。

P_ADV_LAST

自上次从对等方接收通告以来已过去的时间。

MASTER_DOWN_INT

声明主路由器已关闭之前需等待的时间间隔（以毫秒为单位）。

-p、--parseable

以计算机可解析的格式显示 VRRP 路由器信息。

-o *field*[,...]、--output=*field*

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名必须是上面列出的字段之一，或是用以显示所有字段的特殊值 `all`（全部）。缺

省情况下（不带 -o），vrrpadm show 将显示所有字段。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/network/routing/vrrp
接口稳定性	Committed（已确定）

另请参见 [dladm\(1M\)](#)、[vrrpd\(1M\)](#)、[attributes\(5\)](#)

引用名	vrrpd – VRRP 守护进程
用法概要	/usr/lib/vrrpd svc:/network/vrrp:default
描述	vrrpd 是为 VRRP 路由器处理管理事件的系统守护进程程序。通过服务管理工具 (Service Management Facility, SMF) 服务实例对其进行控制： svc:/network/vrrp:default 不能直接调用该守护进程。它不构成管理接口或编程接口的一部分。用于管理 VRRP 路由器的管理接口是 vrrpadm(1M)。
选项	该守护进程没有任何选项。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	system/network/routing/vrrp
接口稳定性	Private（专用）

另请参见 [vrrpadm\(1M\)](#)、[attributes\(5\)](#)

引用名	vscanadm – vscan 服务配置实用程序
用法概要	<pre> vscanadm set -p <i>property</i>=<i>value</i> [-p <i>property</i>=<i>value</i>]... vscanadm get [-p <i>property</i>]... vscanadm import -p <i>property</i> <i>filename</i> vscanadm export -p <i>property</i> <i>filename</i> vscanadm validate -p <i>property</i> <i>filename</i> vscanadm add-engine [-p <i>property</i>=<i>value</i>]... <i>engine_id</i> vscanadm remove-engine <i>engine_id</i> vscanadm set-engine -p<i>property</i>=<i>value</i> [-p <i>property</i>=<i>value</i>]... <i>engine_id</i> vscanadm get-engine [-p <i>property</i>=<i>value</i>]... [<i>engine_id</i>] vscanadm show vscanadm stats [-z]</pre>
描述	vscanadm 命令可设置和显示 vscan 服务 (zpool(1M)) 的属性并提供扫描统计信息。 可以使用合适的文件系统管理命令（例如 zfs(1M)）为每个文件系统配置文件系统的免病毒扫描。 扫描引擎是位于外部主机上的第三方应用程序，它们对文件执行实际的病毒扫描操作。可以配置多个扫描引擎供 vscan 服务使用。建议至少配置两个扫描引擎。文件扫描请求将分布到所配置的各个扫描引擎中，以实现负载平衡。扫描引擎由其 <i>engine_id</i> 予以标识。<i>engine_id</i> 是用户定义的字符串，最多为 64 个字节。 vscan 服务属性分为两类：扫描引擎属性（特定于扫描引擎定义）和常规属性（应用于该服务但不特定于扫描引擎）。
子命令	vscanadm 识别以下子命令： <pre> vscanadm set -p <i>property</i>=<i>value</i> [-p <i>property</i>=<i>value</i>]... 设置 vscan 服务常规属性的值。 -p <i>property</i>=<i>value</i> 指定一个属性值 vscanadm get [-p <i>property</i>]... 显示 vscan 服务常规属性的值。如果未指定任何属性，则显示所有 vscan 服务常规属性。 -p <i>property</i> 指定一个属性值</pre> 以下属性可用于 vscanadm set 和 vscanadm get 子命令： <pre> max-size 应进行病毒扫描的文件的最大大小。不会扫描超过 <i>max-size</i> 的文件。 <i>max-size-action</i> 属性决定了应允许还是拒绝访问超过 <i>max-size</i> 的文件。</pre>

max-size 的值是一个字符串，包含一个数字（十进制）组件和一个可选的用于指定单位大小的字母组件，格式为 "N[N][KMGTP][B]"。

跟在数字组件后面的可选单位可以指定为一个或两个字符。例如，"K" 或 "KB" 都可用来指定千字节。单位说明符不区分大小写，必须紧跟在数字值后，中间不能有空格。

如果没有单位说明符，或单位说明符仅为 "B"，则假定数字值以字节为单位。缺省值为 1GB。

请注意，当 *vscan* 服务定义了扫描的最大文件大小时，扫描引擎通常还会定义其自己的最大文件大小设置。建议将 *max-size* 设置为小于等于扫描引擎所定义的最大文件大小的值。

max-size-action	指定允许还是拒绝访问大于 <i>max-size</i> 的文件。不会对大于 <i>max-size</i> 的文件进行病毒扫描。有效值包括：
allow （允许列表）	允许访问大于 <i>max-size</i> 的文件（无病毒扫描）。这是缺省值。
deny	拒绝访问大于 <i>max-size</i> 的文件（无病毒扫描）。

vscanadm import -p *property filename*

从指定的文件导入属性值。该文件必须包含指定了单个属性的值的单个行。

vscanadm export -p *property filename*

将属性值导出到指定的文件。该文件必须包含指定了单个属性的值的单个行。

vscanadm validate -p *property filename*

验证指定文件中的属性值。该文件必须包含指定了单个属性的值的单个行。

以下属性可用于 *vscanadm import*、*vscanadm export* 和 *vscanadm validate* 子命令：

types 文件类型扩展名匹配规则的逗号分隔列表。此列表定义在病毒扫描期间要扫描哪些类型的文件以及应排除哪些类型的文件。每条规则都包含规则指示符 *[+|-]*，后跟要将文件的类型扩展名与之进行比较的文件类型 *expression*。文件类型 *expression* 不区分大小写，可以包含 "*" 和 "?" 通配符。规则指示符与文件类型 *expression* 之间不应有空格。如果文件类型表达式内包含逗号，则必须使用 "\"（反斜杠）对其进行转义。文件类型扩展名不包含其前面的点。

规则指示符是单个字符，可以是下面列出的项之一：

```
+ include file type in virus scanning
- exclude file type from virus scanning
```

对某个文件进行评估并确定它是要进行病毒扫描的候选文件后，会将其文件类型将与类型中定义的规则进行比较。将应用第一个相匹配的规则。如

果未找到匹配项，将对该文件进行病毒扫描。类型字符串的总长度不能超过 4096 字节。类型列表的缺省内容为 "+"。

vscanadm add-engine [-p *property=value*]... *engine_id*

添加由 *engine_id* 标识的一个新扫描引擎。对于未指定的任何扫描引擎，将为其使用缺省值。主机名缺省设置为 *engine_id*。

-p *property=value* 指定一个属性值

vscanadm remove-engine *engine_id*

删除由 *engine_id* 标识的扫描引擎，进而删除其所有配置属性值。

vscanadm set-engine -p*property=value* [-p *property=value*]... *engine_id*

为 *engine_id* 标识的扫描引擎创建或更新配置属性值。

-p *property=value* 指定一个属性值

vscanadm get-engine [-p *property=value*]... [*engine_id*]

为 *engine_id* 标识的扫描引擎显示指定的扫描引擎属性的值。如果未指定 *engine_id*，该子命令将为所有已配置的扫描引擎显示指定的扫描引擎属性值。如果未指定属性，该子命令将显示所有 vscan 服务扫描引擎属性。

-p *property=value* 指定一个属性值

以下属性可用于 **vscanadm add-engine**、**vscanadm remove-engine**、**vscanadm set-engine** 和 **vscanadm get-engine** 子命令。

enable	指定是启用还是禁用扫描引擎。有效值为 "on"（启用）和 "off"（禁用）。缺省值为 "on"（启用）。如果扫描引擎的 host（主机）属性无效，则无法启用该扫描引擎。
host	扫描引擎的主机名或 IPv4 格式的 IP 地址。
port	扫描引擎的 ICAP 端口号。该数字值的范围是 0 到 65535。缺省的 ICAP 端口为 1344。
max-connection	可以与一个扫描引擎建立的最大并发连接数。该数字值的范围是 1 到 512。该属性的缺省值为 8。
vscanadm show	显示所有 vscan 服务常规属性和扫描引擎属性的值。
vscanadm stats [-z]	显示或重置以下 vscan 服务统计信息： ■ 已扫描的文件数 ■ 受感染的文件数 ■ 失败的扫描请求数 ■ 扫描错误（包括每个扫描引擎的错误计数） -z 将 vscan 服务统计信息计数器重置为零

示例

示例1 设置最大大小限制

要将进行病毒扫描的文件的最大大小限制设置为 128 兆字节，请输入

```
# vscanadm set -p max-size=128M
```

示例2 允许访问文件

要允许访问超过最大文件大小的文件，请输入

```
# vscanadm set -p max-size-action=allow
```

示例3 设置文件类型

要设置类型以便仅对类型为 "odt"、"exe" 和 "jpg" 的文件进行病毒扫描，请输入

```
# vscanadm set -p types=+odt,+exe,+jpg,-*
```

要设置类型以便对除 "doc" 之外的所有文件类型进行病毒扫描，请输入

```
# vscanadm set -p types=-doc,+*
```

示例4 显示文件类型列表

要显示文件类型列表，请输入

```
# vscanadm get -p types
```

示例5 添加扫描引擎

要使用缺省值添加扫描引擎 "my_eng"，请输入

```
# vscanadm add-engine my_eng
```

示例6 禁用扫描引擎

要禁用扫描引擎 "my_eng"，请输入

```
# vscanadm set-engine -p enable=off my_eng
```

示例7 显示扫描引擎属性

要显示扫描引擎 "my_eng" 的属性，请输入

```
# vscanadm get-engine my_eng
```

示例8 删除扫描引擎

要删除扫描引擎 "my_eng"，请输入

```
# vscanadm remove-engine my_eng
```

示例 9 显示 Vscan 服务常规属性和扫描引擎属性
要显示所有 vscan 服务常规属性和扫描引擎属性，请输入

```
# vscanadm show
```

退出状态 将返回下列退出值：

- 0 成功完成。
- 非零 出现错误。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/storage/virus-scan
接口稳定性	Uncommitted（未确定）
实用程序输出格式	Not-An-Interface（不是接口）

另请参见 [vscand\(1M\)](#)、[zfs\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

附注 所有用户都可以使用 vscanadm 查看 vscan 属性和统计信息。要设置属性值或重置统计信息，需要以下授权：

solaris.smf.value.vscan
更改属性值或重置统计信息

solaris.manage.vscan
刷新服务以应用属性值更改

要添加或删除属性（add-engine、remove-engine），需要以下授权：

solaris.smf.modify.application 添加或删除属性组

solaris.manage.vscan 刷新服务以应用属性值更改

所有这些授权都包含在 "VSCAN Management"（VSCAN 管理）配置文件中。

引用名 vscand – vscan 服务守护进程

用法概要 /usr/lib/vscan/vscand

描述 vscand 是一种守护进程，可处理文件系统发出的针对文件打开和文件关闭操作的病毒扫描请求。文件系统可以在每一个数据集的基础上，使用其管理命名（例如 [zfs\(1M\)](#)）来支持启用和禁用病毒扫描。

如果文件状态或扫描策略（请参见 [vscanadm\(1M\)](#)）要求扫描某个文件，vscand 将使用 Internet 内容适配协议（ICAP，RFC 3507）与外部第三方病毒扫描程序（扫描引擎）通信，以便对文件进行扫描。

如果在最后一次扫描某个文件后又对其进行了修改，或者尚未使用最新的扫描引擎配置（病毒定义）扫描该文件，则会将其提交给扫描引擎。文件的 modified 属性和 scanstamp 属性可用于存储该信息。完成文件扫描以后，将清除 modified 属性并更新 scanstamp 属性。

如果发现文件包含病毒，则会将病毒记录在 [syslogd\(1M\)](#) 中，同时编写审计记录并将文件隔离起来（只需设置文件的 quarantine 属性即可）。隔离某个文件以后，文件系统将拒绝针对该文件的读取、执行或重命名操作。[syslogd\(1M\)](#) 项和审计记录描述了受感染文件的名称以及在文件中检测到的违规。将按“ID - 威胁描述”这样的形式来描述每一项违规，其中，ID 和威胁描述均在 ICAP RFC 3507; 扩展的 X-Infection-Found-Header 中定义。

缺省情况下，vscand 通过端口 1344 连接到扫描引擎。端口和其他服务配置参数可使用 [vscanadm\(1M\)](#) 进行配置。

vscan 服务缺省情况下禁用，但可通过 [svcadm\(1M\)](#) 启用。

退出状态 将返回下列退出值：

- 0 守护进程启动成功。
- 非零 守护进程启动失败。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	service/storage/virus-scan
接口稳定性	Uncommitted（未确定）

另请参见 [ps\(1\)](#)、[svcs\(1\)](#)、[logadm\(1M\)](#)、[svcadm\(1M\)](#)、[syslogd\(1M\)](#)、[vscanadm\(1M\)](#)、[zfs\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

附注

如果使用不调用文件系统打开和关闭操作的协议（例如 NFSv3）来访问文件，则不会启动针对该文件的病毒扫描。

文件内容将以 `cleartext` 数据形式传输到扫描引擎。

可以使用 `svcadm(1M)` 来执行对 `vscan` 服务的管理操作，如启用、禁用或请求重新启动。可以使用 `svcs(1)` 命令来查询 `vscan` 服务的状态。

`vscan` 服务由服务管理工具 `smf(5)` 管理，其服务标识符为：

```
svc:/system/filesystem/vscan
```

引用名 vtdaemon – 文本虚拟控制台之间的切换功能

用法概要 /usr/lib/vtdaemon

描述 vtdaemon 提供在不同的文本虚拟控制台之间通过热键方式进行切换的安全切换功能。它在完全特权模式下运行。

该服务可以使用设备安全权限配置文件中所含的操作授权 `solaris.smf.manage.vt` 进行管理。该服务仅仅是本地服务，没有传入网络端口，且在非全局区域中或虚拟控制台功能不可用的情况下处于禁用状态。

vtdaemon 使用 PAM 进行验证。

请注意，vtdaemon 仅对文本控制台执行锁定和解除锁定操作，而对于图形 X 会话，则只调用相应的锁定程序（例如，`xscreensaver`）来处理验证和审计操作。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/core-os
接口稳定性	Private（专用）

另请参见 [svcs\(1\)](#)、[svcadm\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

附注 vtdaemon 服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

`svc:/system/vtdaemon:default`

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

vtdaemon 服务具有以下 SMF 配置属性：

- hotkeys (boolean)（热键）（布尔型）**
允许授权的用户动态启用或禁用采用热键方式的 VT 切换。其缺省值为 **TRUE**（启用）。
- secure (boolean)（安全）（布尔型）**
允许授权的用户动态启用或禁用热键的安全性。如果禁用该属性，用户不经验证即可随意切换到任何会话。其缺省值为 **TRUE**（启用）。

引用名	wall – write to all users
用法概要	<code>/usr/sbin/wall [-a] [-g <i>grpname</i>] [<i>filename</i>]</code>
描述	wall reads its standard input until an end-of-file. It then sends this message to all currently logged-in users preceded by: <pre>Broadcast Message from . . .</pre> If <i>filename</i> is given, then the message is read in from that file. Normally, pseudo-terminals that do not correspond to rlogin sessions are ignored. Thus, when using a window system, the message appears only on the console window. However, -a will send the message even to such pseudo-terminals. It is used to warn all users, typically prior to shutting down the system. The sender must be superuser to override any protections the users may have invoked See mesg(1). wall runs <code>setgid()</code> to the group ID <code>tty</code>, in order to have write permissions on other user's terminals. See setuid(2). wall will detect non-printable characters before sending them to the user's terminal. Control characters will appear as a <code>"^"</code> followed by the appropriate ASCII character; characters with the high-order bit set will appear in “meta” notation. For example, <code>\003</code> is displayed as <code>^C</code> and <code>\372</code> as <code>M-z</code>.
选项	The following options are supported: -a Broadcast message to the console and pseudo-terminals. -g <i>grpname</i> Broadcast to the users in a specified group only, per the group database (see group(4)).
环境变量	If the <code>LC_*</code> variables (<code>LC_CTYPE</code>, <code>LC_TIME</code>, <code>LC_COLLATE</code>, <code>LC_NUMERIC</code>, and <code>LC_MONETARY</code>) are not set in the environment, the operational behavior of wall for each corresponding locale category is determined by the value of the <code>LANG</code> environment variable. See environ(5). If <code>LC_ALL</code> is set, its contents are used to override both the <code>LANG</code> and the other <code>LC_*</code> variables. If none of the above variables are set in the environment, the "C" (U.S. style) locale determines how wall behaves.
文件	<code>/dev/tty*</code>
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[msg\(1\)](#), [write\(1\)](#), [setuid\(2\)](#), [attributes\(5\)](#), [environ\(5\)](#)

附注

wall displays “Cannot send to . . .” when the open on a user's tty file fails.

引用名

wanboot_keygen – create and display client and server keys for WAN booting

用法概要

```
/usr/lib/inet/wanboot/keygen -c -o net=a.b.c.d ,cid=client_ID,type=3des
/usr/lib/inet/wanboot/keygen -c -o net=a.b.c.d ,cid=client_ID,type=aes
/usr/lib/inet/wanboot/keygen -m
/usr/lib/inet/wanboot/keygen -c -o net=a.b.c.d ,cid=client_ID,type=sha1
/usr/lib/inet/wanboot/keygen -d -m
/usr/lib/inet/wanboot/keygen -c -o net=a.b.c.d ,cid=client_ID,type=keytype
```

描述

The keygen utility has three purposes:

- Using the -c flag, to generate and store per-client 3DES/AES encryption keys, avoiding any DES weak keys.
- Using the -m flag, to generate and store a “master” HMAC SHA-1 key for WAN install, and to derive from the master key per-client HMAC SHA-1 hashing keys, in a manner described in RFC 3118, Appendix A.
- Using the -d flag along with either the -c or -m flag to indicate the key repository, to display a key of type specified by *keytype*, which must be one of 3des, aes, or sha1.

The net and cid arguments are used to identify a specific client. Both arguments are optional. If the cid option is not provided, the key being created or displayed will have a per-network scope. If the net option is not provided, then the key will have a global scope. Default net and code values are used to derive an HMAC SHA-1 key if the values are not provided by the user.

选项

The following options are supported:

- c Generate and store per-client 3DES/AES encryption keys, avoiding any DES weak keys. Also generates and stores per-client HMAC SHA-1 keys. Used in conjunction with -o.
- d Display a key of type specified by *keytype*, which must be one of 3des, aes, or sha1. Use -d with -m or with -c and -o.
- m Generate and store a “master” HMAC SHA-1 key for WAN install.
- o Specifies the WANboot client and/or keytype.

示例

示例 1 Generate a Master HMAC SHA-1 Key

```
# keygen -m
```

示例 2 Generate and Then Display a Client-Specific Master HMAC SHA-1 Key

```
# keygen -c -o net=172.16.174.0,cid=010003BA0E6A36,type=sha1
# keygen -d -c -o net=172.16.174.0,cid=010003BA0E6A36,type=sha1
```

```
# keygen -c -o net=172.16.174.0,type=3des
# keygen -d -o net=172.16.174.0,type=3des
```

0	Successful operation.
>0	An error occurred.

See `attributes(5)` for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/boot/wanboot
Interface Stability	Obsolete

```
attributes(5)
```

引用名

wanboot_keymgmt – insert and extract keys

用法概要

```
/usr/lib/inet/wanboot/keymgmt -i -k key_file -s keystore -o type=keytype
/usr/lib/inet/wanboot/keymgmt -x -f outfile -s keystore -o type=keytype
```

描述

The keymgmt utility has two purposes:

- To take a raw key, stored in *key_file*, and insert it in the repository specified by *keystore*.
- To extract a key of a specified type from the repository specified by *keystore*, depositing it in *outfile*.

outfile will be created if it does not already exist. The type of key being added or extracted is specified by *keytype* and may have one of four values: *3des*, *aes*, *rsa*, or *sha1* (the last used by HMAC SHA-1). When extracting a key, the first key with an OID matching the supplied type is used.

Arguments

The following arguments are supported:

-i	Used in conjunction with -k to insert a raw key in <i>keystore</i> .
-f <i>outfile</i>	Used to specify a file to receive an extracted key.
-k <i>key_file</i>	Used in conjunction with -i to specify the file in which a raw key is stored. This key will be inserted in <i>keystore</i> .
-o type= <i>keytype</i>	Specifies the type of key being inserted or extracted. Must be one of <i>3des</i> , <i>aes</i> , <i>rsa</i> , or <i>sha1</i> .
-s <i>keystore</i>	Specifies a repository in which a key will be inserted or from which a key will be extracted.
-x	Used in conjunction with -f to extract a key of a specified type and deposit it in <i>outfile</i> .

退出状态

0	Successful operation.
>0	An error occurred.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/boot/wanboot
Interface Stability	Obsolete

另请参见

[attributes\(5\)](#)
ITU-T Recommendation X.208

引用名

wanboot_p12split – split a PKCS #12 file into separate certificate and key files

用法概要

```
/usr/lib/inet/wanboot/p12split -i p12file -c out_cert -k out_key
[-t out_trust -l id -v]
```

描述

The `p12split` utility extracts a certificate and private key from the repository specified by *p12file*, depositing the certificate in *out_cert* and the key in *out_key*. If supplied, the `-l` option specifies the value for the `LocalKeyId` that will be used in the new certificate and key files. `p12split` can optionally extract a trust certificate into the *out_trust* file if the `-t` option is specified. Use the `-v` option to get a verbose description of the split displayed to standard output.

选项

The following arguments and options are supported:

- `-c out_cert` Specifies a repository that receives a extracted certificate.
- `-i p12file` Specifies a repository from which a certificate and private key is extracted.
- `-k out_key` Specifies a repository that receives a extracted private key.
- `-l id` Specifies the value for the `LocalKeyId` that will be used in the new certificate and key files.
- `-t out_trust` Specifies a file for receiving an extracted trust certificate.
- `-v` Displays a verbose description of the split to stdout.

退出状态

0 Successful operation.

>0 An error occurred.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/boot/wanboot
Interface Stability	Uncommitted

另请参见 [attributes\(5\)](#)

引用名	wanbootutil – manage keys and certificates for WAN booting
用法概要	wanbootutil [keygen] [<i>option_specific_arguments</i>] wanbootutil [keymgmt] [<i>option_specific_arguments</i>] wanbootutil [p12split] [<i>option_specific_arguments</i>]
描述	The wanbootutil command creates and manages WANboot encryption and hashing keys and manipulates PKCS #12 files for use by WAN boot. wanbootutil has three subcommands, each covered in a separate man page: wanboot_keygen(1M) Generates encryption and hashing keys. wanboot_keymgmt(1M) Inserts and extracts keys from WAN boot key repositories. wanboot_p12split(1M) Splits a PKCS #12 file into separate certificate and key files for use by WAN boot.
选项	The options are supported for wanbootutil are the use of keygen, keymgmt, or p12split. The options for these subcommands are described in their respective man pages.
示例	示例 1 Generate a 3DES Client Key <pre># wanbootutil keygen -c -o net=172.16.174.0,cid=010003BA0E6A36,type=3des</pre> 示例 2 Insert an RSA Private Client Key <pre>wanbootutil keymgmt -i -k keyfile \ -s /etc/netboot/172.16.174.0/010003BA0E6A36/keystore -o type=rsa</pre> 示例 3 Split a PKCS #12 File into Certificate and Key Components <pre># wanbootutil p12split -i p12file -c out_cert -k out_key</pre>
退出状态	0 Successful operation. non-zero An error occurred. Writes an appropriate error message to standard error.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTETYPE	ATTRIBUTE VALUE
Availability	system/boot/wanboot
Interface Stability	Obsolete

另请参见 [wanboot_keygen\(1M\)](#), [wanboot_keymgmt\(1M\)](#), [wanboot_p12split\(1M\)](#), [attributes\(5\)](#)

引用名 `wbemadmin – start Sun WBEM User Manager`

用法概要 `/usr/sadm/bin/wbemadmin`

描述 The `wbemadmin` utility starts Sun WBEM User Manager, a graphical user interface that enables you to add and delete authorized WBEM users and to set their access privileges. Use this application to manage access to groups of managed resources, such as disks and installed software, in the Solaris operating environment.

The `wbemadmin` utility allows you to perform the following tasks:

Manage user access rights	Use the <code>wbemadmin</code> utility to add, delete, or modify an individual user's access rights to a namespace on a WBEM-enabled system.
Manage namespace access rights	Use the <code>wbemadmin</code> utility to add, delete, or modify access rights for all users to a namespace.

The Sun WBEM User Manager displays a Login dialog box. You must log in as root or a user with write access to the `root\security` namespace to grant access rights to users. By default, Solaris users have guest privileges, which grants them read access to the default namespaces.

Managed resources are described using a standard information model called Common Information Model (CIM). A CIM object is a computer representation, or model, of a managed resource, such as a printer, disk drive, or CPU. CIM objects can be shared by any WBEM-enabled system, device, or application. CIM objects are grouped into meaningful collections called schema. One or more schemas can be stored in directory-like structures called namespaces.

All programming operations are performed within a namespace. Two namespaces are created by default during installation:

- `root\cimv2` — Contains the default CIM classes that represent objects on your system.
- `root\security` — Contains the security classes used by the CIM Object Manager to represent access rights for users and namespaces.

When a WBEM client application connects to the CIM Object Manager in a particular namespace, all subsequent operations occur within that namespace. When you connect to a namespace, you can access the classes and instances in that namespace (if they exist) and in any namespaces contained in that namespace.

When a WBEM client application accesses CIM data, the WBEM system validates the user's login information on the current host. By default, a validated WBEM user is granted read access to the Common Information Model (CIM) Schema. The CIM Schema describes managed objects on your system in a standard format that all WBEM-enabled systems and applications can interpret.

You can set access privileges on individual namespaces or for a user-namespace combination. When you add a user and select a namespace, by default the user is granted read access to CIM objects in the selected namespace. An effective way to combine user and namespace access rights is to first restrict access to a namespace. Then grant individual users read, read and write, or write access to that namespace.

You cannot set access rights on individual managed objects. However you can set access rights for all managed objects in a namespace as well as on a per-user basis.

If you log in to the root account, you can set the following types of access to CIM objects:

- Read Only — Allows read-only access to CIM Schema objects. Users with this privilege can retrieve instances and classes, but cannot create, delete, or modify CIM objects.
- Read/Write — Allows full read, write, and delete access to all CIM classes and instances.
- Write — Allows write and delete, but not read access to all CIM classes and instances.
- None — Allows no access to CIM classes and instances.

Context help is displayed in the left side of the wbemadmin dialog boxes. When you click on a field, the help content changes to describe the selected field. No context help is available on the main User Manager window.

The wbemadmin security administration tool updates the following Java classes in the root\security namespace:

- Solaris_UserAcl — Updated when access rights are granted or changed for a user.
- Solaris_namespaceAcl — Updated when access rights are granted or changed for a namespace.

用法

The wbemadmin utility is not the tool for a distributed environment. It is used for local administration on the machine on which the CIM Object Manager is running.

退出状态

The wbemadmin utility terminates with exit status 0.

Warning

The root\security namespace stores access privileges. If you grant other users access to the root\security namespace, those users can grant themselves or other users rights to all other namespaces.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWwbco

另请参见 [mofcomp\(1M\)](#), [wbemlogviewer\(1M\)](#), [init.wbem\(1M\)](#), [attributes\(5\)](#)

引用名	wbemconfig – convert a JavaSpaces datastore to the newer Reliable Log datastore format
用法概要	<code>/usr/sadm/lib/wbem/wbemconfig convert</code>
描述	A Reliable Log directory is created that contains the converted data. This directory is named <code>/var/sadm/wbem/logr</code>. The <code>convert</code> argument is the only supported option of this command. You should only run this command after stopping WBEM (CIM Object Manager) with the <code>init.wbem stop</code> command. Otherwise your data may be corrupted. This command successfully converts any proprietary custom MOFs you have created in the datastore, but not any CIM or Solaris MOFs you have modified. These will be destroyed. To recompile any modified CIM or Solaris MOFs into the new datastore, run the <code>mofcomp</code> command on the MOF files containing the class definitions. Because the <code>wbemconfig convert</code> command invokes the JVM (Java Virtual Machine) to perform conversion of the JavaSpaces datastore, you must be running the same version of the JVM as when the original JavaSpaces storage was created. After the <code>wbemconfig convert</code> command is completed, you can change to any version of the JVM you want. To see what version of the JVM you are running, issue the <code>java -version</code> command.
选项	The following options are supported: <code>convert</code> Convert a JavaSpaces datastore to the newer Reliable Log datastore format.
属性	See attributes(5) for descriptions of the following attributes:
另请参见	init.wbem(1M) , wbemadmin(1M) , wbemlogviewer(1M) , mofcomp(1M) , attributes(5)

ATTRIBUTETYPE	ATTRIBUTEVALUE
Availability	SUNWwbco

引用名

wbemlogviewer – start WBEM Log Viewer

用法概要

/usr/sadm/bin/wbemlogviewer

描述

The `wbemlogviewer` utility starts the WBEM Log Viewer graphical user interface, which enables administrators to view and maintain log records created by WBEM clients and providers. The WBEM Log Viewer displays a Login dialog box. You must log in as root or a user with write access to the `root\cimv2` namespace to view and maintain log files. Namespaces are described in [wbemadmin\(1M\)](#).

Log events can have three severity levels.

- Errors
- Warnings
- Informational

The WBEM log file is created in the `/var/sadm/wbem/log` directory, with the name `wbem_log`. The first time the log file is backed up, it is renamed `wbem_log.1`, and a new `wbem_log` file is created. Each succeeding time the `wbem_log` file is backed up, the file extension number of each backup log file is increased by 1, and the oldest backup log file is removed *if* the limit, which in turn is specified in the log service settings, on the number of logfiles is exceeded. Older backup files have higher file extension numbers than more recent backup files.

The log file is renamed with a `.1` file extension and saved when one of the following two conditions occur:

- The current file reaches the specified file size limit.
- A WBEM client application uses the `clearLog()` method in the `Solaris_LogService` class to clear the current log file.
- A WBEM client application uses the `clearLog()` method in the `Solaris_LogService` class to clear the current log file.
- A user chooses Action->Back Up Now in the Log Viewer application.

Help is displayed in the left panel of each dialog box. Context help is not displayed in the main Log Viewer window.

用法

The WBEM Log Viewer is not the tool for a distributed environment. It is used for local administration.

The WBEM Log Viewer allows you to perform the following tasks:

View the logs

Set properties of log files	Click Action->Log File Settings to specify log file parameters and the log file directory.

Back up a log file	Click Action->Back Up Now to back up and close the current log file and start a new log file.

- Open historical log files Click Action->Open Log File to open a backed-up log file.
- Delete an old log file Open the file and then click Action->Delete Log File. You can only delete backed-up log files.
- View log record details Double-click a log entry or click View->Log Entry Details to display the details of a log record.
- Sort the logs Click View->Sort By to sort displayed entries. You can also click any column heading to sort the list. By default, the log entries are displayed in reverse chronological order (new logs first).

退出状态 The `wbemlogviewer` utility terminates with exit status 0.

文件 `/var/sadm/wbem/log/wbem_log` WBEM log file

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWwbco

另请参见 [wbemadmin\(1M\)](#), [init.wbem\(1M\)](#), [mofcomp\(1M\)](#), [attributes\(5\)](#)

引用名	wcadmin – manage the configuration of the Sun Java Web Console										
用法概要	<pre>/usr/sbin/wcadmin <i>subcommand options</i> /usr/sbin/wcadmin [-h] --help] /usr/sbin/wcadmin [-V --version]</pre>										
描述	wcadmin is a Command Line Interface (CLI) — based tool for managing the configuration of the Sun Java Web Console.										
Subcommands	The following subcommands are supported: <table><tr><td>add</td><td> The <code>wcadmin add</code> subcommand adds a new shared jar file, a new JAAS login module, or a new shared service property to the console configuration. An optional identifier may be specified; if omitted, the identifier is derived from the resource name. The resource is added when the console is next started. The format of the <code>add</code> subcommand is: <pre>add -l -a <i>appname</i> [-n <i>id</i>] <i>jarpath</i> add -m -a <i>appname</i> [-n <i>id</i>] -b <i>behavior</i> -s <i>service</i> [-o <i>name=value</i>] <i>classname</i> add -p -a <i>appname</i> <i>name=value</i> [... <i>name=value</i>]</pre></td></tr><tr><td>deploy</td><td> The <code>wcadmin deploy</code> subcommand deploys the specified web application into the console's web server instance. Applications are deployed directly from their installation directories. The format of the <code>deploy</code> subcommand is: <pre>deploy [-D] -a <i>appname</i> -x <i>context</i> <i>app_path</i></pre></td></tr><tr><td>disable</td><td> The <code>wcadmin disable</code> subcommand disables access to the specified web application in the console's web server instance. The format of the <code>disable</code> subcommand is: <pre>disable -x <i>context</i></pre></td></tr><tr><td>enable</td><td> The <code>wcadmin enable</code> subcommand enables access to the specified web application in the console's web server instance. The format of the <code>enable</code> subcommand is: <pre>enable -x <i>context</i></pre></td></tr><tr><td>list</td><td> The <code>wcadmin list</code> subcommand lists the resources currently configured for the console; including deployed web applications, shared jar files, login modules, and shared service properties. If no option is specified, all resources are listed. </td></tr></table>	add	The <code>wcadmin add</code> subcommand adds a new shared jar file, a new JAAS login module, or a new shared service property to the console configuration. An optional identifier may be specified; if omitted, the identifier is derived from the resource name. The resource is added when the console is next started. The format of the <code>add</code> subcommand is: <pre>add -l -a <i>appname</i> [-n <i>id</i>] <i>jarpath</i> add -m -a <i>appname</i> [-n <i>id</i>] -b <i>behavior</i> -s <i>service</i> [-o <i>name=value</i>] <i>classname</i> add -p -a <i>appname</i> <i>name=value</i> [... <i>name=value</i>]</pre>	deploy	The <code>wcadmin deploy</code> subcommand deploys the specified web application into the console's web server instance. Applications are deployed directly from their installation directories. The format of the <code>deploy</code> subcommand is: <pre>deploy [-D] -a <i>appname</i> -x <i>context</i> <i>app_path</i></pre>	disable	The <code>wcadmin disable</code> subcommand disables access to the specified web application in the console's web server instance. The format of the <code>disable</code> subcommand is: <pre>disable -x <i>context</i></pre>	enable	The <code>wcadmin enable</code> subcommand enables access to the specified web application in the console's web server instance. The format of the <code>enable</code> subcommand is: <pre>enable -x <i>context</i></pre>	list	The <code>wcadmin list</code> subcommand lists the resources currently configured for the console; including deployed web applications, shared jar files, login modules, and shared service properties. If no option is specified, all resources are listed.
add	The <code>wcadmin add</code> subcommand adds a new shared jar file, a new JAAS login module, or a new shared service property to the console configuration. An optional identifier may be specified; if omitted, the identifier is derived from the resource name. The resource is added when the console is next started. The format of the <code>add</code> subcommand is: <pre>add -l -a <i>appname</i> [-n <i>id</i>] <i>jarpath</i> add -m -a <i>appname</i> [-n <i>id</i>] -b <i>behavior</i> -s <i>service</i> [-o <i>name=value</i>] <i>classname</i> add -p -a <i>appname</i> <i>name=value</i> [... <i>name=value</i>]</pre>										
deploy	The <code>wcadmin deploy</code> subcommand deploys the specified web application into the console's web server instance. Applications are deployed directly from their installation directories. The format of the <code>deploy</code> subcommand is: <pre>deploy [-D] -a <i>appname</i> -x <i>context</i> <i>app_path</i></pre>										
disable	The <code>wcadmin disable</code> subcommand disables access to the specified web application in the console's web server instance. The format of the <code>disable</code> subcommand is: <pre>disable -x <i>context</i></pre>										
enable	The <code>wcadmin enable</code> subcommand enables access to the specified web application in the console's web server instance. The format of the <code>enable</code> subcommand is: <pre>enable -x <i>context</i></pre>										
list	The <code>wcadmin list</code> subcommand lists the resources currently configured for the console; including deployed web applications, shared jar files, login modules, and shared service properties. If no option is specified, all resources are listed.										

	The format of the <code>list</code> subcommand is: <pre>list [-a] [-d] [-l] [-m] [-p]</pre>
password	The <code>wcadmin password</code> subcommand manages the administrator and security keystore passwords for the console. Keystore passwords should not be changed while the console is running. The format of the <code>password</code> subcommand is: <pre>password [-a] [-k] [-t]</pre> <pre>password -f <i>password_file</i></pre>
reload	The <code>wcadmin reload</code> subcommand unloads the specified web application from the console's web server instance and reloads the application from its original installation directory. The format of the <code>reload</code> subcommand is: <pre>reload -x <i>context</i></pre>
remove	The <code>wcadmin remove</code> subcommand removes a shared jar file, a login module, or a shared service property from the console configuration. The resource may be specified by its identifier or by its full jarpath or classname. The resource is removed when the console is next started. The format of the <code>remove</code> subcommand is: <pre>remove -l -a <i>appname</i> [-n <i>id</i>] <i>jarpath</i></pre> <pre>remove -m -a <i>appname</i> [-n <i>id</i>] <i>classname</i></pre> <pre>remove -p -a <i>appname</i> <i>property</i> [...]</pre>
undeploy	The <code>wcadmin undeploy</code> subcommand undeploys the specified web application from the console's web server instance. The format of the <code>undeploy</code> subcommand is: <pre>undeploy [-D] -a <i>appname</i> -x <i>context</i></pre>

选项

The following options are supported:	
<code>-h --help -?</code>	Display runtime help.
<code>-V --version</code>	Display console version information.
<code>-D --defer</code>	When used with the <code>deploy</code> and <code>undeploy</code> subcommands, defers the deployment or undeployment until the next console restart.

The operation is deferred by simply adding or removing the corresponding resource registration notification file. If defer is not specified, a runtime deployment or undeployment is performed, so that the application becomes available or unavailable in the running console. If the console instance is not currently running, the operation is automatically deferred.

- a | --adminpassword Specify that the administrator password should be changed, when used with password subcommand. You are prompted for a new password, which must be 8 to 32 characters.
- a | --application Specify the application name, when used with subcommands other than password subcommand.
- d | --detail Specify that configuration details of each resource should be displayed.
- f Specify the fully qualified path name to a file containing one or more password property values. See the description of the password_file argument.
- k | --keypassword Specify that the keystore password should be changed. You are prompted for a new password, which must be 8 to 32 characters.
- l | --library Specify that the resource is a shared jar file.
- m | --module Specify that the resource is a JAAS login module.
- n | --name Specify the short-hand identifier name for the resource. If omitted, the identifier name is derived from the full resource name.
- o | --option Specify the name and value of a login module option property, separated by the equals character.
- p | --property Specify that the resource is one or more shared service properties.
- s | --service Specify the name of the JAAS login service definition. If omitted, the default console login service definition is assumed.
- t | --trustpassword Specify that the truststore password should be changed. You are prompted for a new password, which must be 8 to 32 characters.
- x | --context Specify the web application context path name under which the application is deployed.

Arguments

- app_path* The fully qualified file system path to the web application installation directory.
- appname* The application name. The name must be unique among all web applications registered with the console. It is also used as the name of the subdirectory under the console's pre-registration directory which contains

	all the resource registration notification files for that application. Typically, the application package name, plugin identifier, or context path name is specified for the application name.
<i>behavior</i>	The JAAS login module control flag behavior. Must be one of “optional”, “required”, “requisite”, or “sufficient”.
<i>classname</i>	The fully qualified Java package class name of the JAAS login module. The specified class must be included in a shared jar file added to the console.
<i>context</i>	The web application context path name under which the application is deployed. With the .reg suffix, the context forms the file name of the registration notification file for that application.
<i>id</i>	The short-hand identifier name for a jar file or login module resource to be added or removed. The identifier name must be unique among the resources shared for a given application name. With the .reg suffix, it forms the file name of the registration notification file for that resource.
<i>jarpath</i>	The fully qualified file system path to the jar file resource. When the resource is added to the console, its path is included in the classpath of the console's shared class loader.
<i>option</i>	The JAAS login module option property name.
<i>property</i>	The shared service property name.
<i>password_file</i>	The fully qualified path to a password text file that contains the new administrator, keystore, and truststore passwords in property file format. The administrator password is specified using the “adminpassword” property. The keystore password is specified using the “keypassword” property. The truststore password is specified using the “trustpasswd” property. At least one password property must be contained in the password file.
<i>value</i>	The login module option or shared service property value. If the value contains white space, it must be quoted.

示例

The following command adds a jar file to be shared in the console:

```
wcadmin add -l -a myapp_1.0 -n wbem /usr/sadm/lib/wbem.jar
```

The following command deploys a new web application:

```
wcadmin deploy -a myapp_1.0 -x myapp /opt/SUNWmyapp/myapp
```

The following command reloads an existing web application:

```
wcadmin reload -x myapp
```

The following command undeploys a web application at the next server restart:

```
wcadmin undeploy -D -a myapp_1.0 -x myapp
```

The following command lists all the deployed web applications in the console. If the status field is “running”, the web application is available. If the status field is “stopped”, the web application is disabled and is not available. If all web applications are “stopped”, this typically indicates the console web server instance is not running.

```
wcadmin list -a
```

The following command removes a shared jar file:

```
wcadmin remove -l -a myapp_1.0 -n wbm
```

The following command changes passwords that are specified in a file:

```
wcadmin password -f /home/mydir/console-passwords
```

退出状态

The following exit values are returned:

- 0 Subcommand succeeded without error
- 1 Usage error: missing or malformed arguments
- 2 Fatal error: subcommand failed with one or more errors

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	SUNWmcon
Interface Stability	Committed

另请参见

[attributes\(5\)](#), [environ\(5\)](#)

引用名 whodo – who is doing what

用法概要 `/usr/sbin/whodo [-h] [-l] [user]`

描述 The `whodo` command produces formatted and dated output from information in the `/var/adm/utmpx` and `/proc/pid` files.

The display is headed by the date, time, and machine name. For each user logged in, device name, user-ID and login time is shown, followed by a list of active processes associated with the user-ID. The list includes the device name, process-ID, CPU minutes and seconds used, and process name.

If *user* is specified, output is restricted to all sessions pertaining to that user.

选项 The following options are supported:

- h Suppress the heading.
- l Produce a long form of output. The fields displayed are: the user's login name, the name of the tty the user is on, the time of day the user logged in (in *hours:minutes*), the idle time — that is, the time since the user last typed anything (in *hours:minutes*), the CPU time used by all processes and their children on that terminal (in *minutes:seconds*), the CPU time used by the currently active processes (in *minutes:seconds*), and the name and arguments of the current process.

示例 示例 1 Using the `whodo` Command

The command:

```
example% whodo
```

produces a display like this:

```
Tue Mar 12 15:48:03 1985
bailey
tty09      mcn          8:51
           tty09      28158    0:29 sh

tty52      bdr          15:23
           tty52      21688    0:05 sh
           tty52      22788    0:01 whodo
           tty52      22017    0:03 vi
           tty52      22549    0:01 sh

xt162      lee          10:20
           tty08       6748    0:01 layers
           xt162       6751    0:01 sh
           xt163       6761    0:05 sh
           tty08       6536    0:05 sh
```

环境变量

If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see [environ\(5\)](#)) are not set in the environment, the operational behavior of [tar\(1\)](#) for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how whodo behaves.

- LC_CTYPE

Determines how whodo handles characters. When LC_CTYPE is set to a valid value, whodo can display and handle text and filenames containing valid characters for that locale. The whodo command can display and handle Extended Unix code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. whodo can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.
- LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).
- LC_TIME

Determines how whodo handles date and time formats. In the "C" locale, date and time handling follow the U.S. rules.

退出状态

The following exit values are returned:

- 0

Successful completion.
- non-zero

An error occurred.

文件

- /etc/passwd

System password file
- /var/adm/utmpx

User access and administration information
- /proc/pid

Contains PID

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[ps\(1\)](#), [who\(1\)](#), [attributes\(5\)](#), [environ\(5\)](#)

引用名	wpad – WPA 和 WPA2 协议守护进程
用法概要	<code>/usr/lib/inet/wpad [-i interface] [-k pre_shared_key_name]</code>
描述	wpad 守护进程为 WiFi 保护访问 (WiFi Protected Access, WPA) 版本 1 和版本 2 提供了按 IEEE802.11i 标准定义的常用客户端功能。WPA 是由行业协会 WiFi 联盟制定的。WPA 实现了大部分 IEEE 802.11i 标准，其用途是在 802.11i 标准制定过程中充当一种过渡标准以取代有线等效保密 (Wired Equivalent Privacy, WEP)。WPA2 实现了整套标准。 wpad 提供以下 WPA/IEEE 802.11i 功能： ■ WPA-PSK (“WPA-个人”) ■ 针对 CCMP、TKIP、WEP104、WEP40 的密钥管理 使用 dladm(1M) 停止和启动 wpad 守护进程。用法： <pre># dladm connect-wifi</pre> ...启动 wpad 守护进程。用法： <pre># dladm disconnect-wifi</pre> ...停止该守护进程。
选项	所支持的选项如下： -i <i>interface</i> 指定一个用于启动 wpad 守护进程的 WiFi 链路接口。 -k <i>pre_shared_key_name</i> 指定用于 WiFi 链路的预共享密钥。
示例	示例1 在特定 WiFi 链路上启动 wpad 守护进程 要创建 WPA 密钥 psk，请输入以下命令： <pre># dladm create-secobj -c wpa psk</pre> 要使用密钥 psk 连接到链路 ath0 上的 ESSID wlan，请输入以下命令： <pre># dladm connect-wifi -k psk -e wlan ath0</pre> 示例2 停止特定 WiFi 链路上的 wpad 守护进程 要停止链路 ath0 上的守护进程，请输入： <pre># dladm disconnect-wifi ath0</pre>
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	service/network/wpa

属性类型	属性值
接口稳定性	Uncommitted（未确定）

另请参见 [svcs\(1\)](#)、[dladm\(1M\)](#)、[svcadm\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)

附注 `wpad` 服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

```
svc:/network/wpa:<link>
```

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcadm\(1M\)](#) 命令来查询服务的状态。

引用名	wracct – write extended accounting records for active processes and tasks
用法概要	<pre>/usr/sbin/wracct -i <i>id_list</i> [-t <i>record_type</i>] {process task}</pre>
描述	The wracct utility allows the administrator to invoke the extended accounting system, if active, to write intermediate records representing the resource usage of a selected set of processes or tasks. For tasks, a <i>record_type</i> option is also supported, allowing the administrator to request the writing of: ■ an interval record, which reflects task usage since a previous interval record (or since task creation if there is no interval record), or ■ a partial record, which reflects usage since task creation.
选项	The following options are supported: -i <i>id_list</i> Select the IDs of the tasks or processes to write records for. Specify <i>id_list</i> as a comma- or space-separated list of IDs, presented as a single argument. For some shells, this requires appropriate quoting of the argument. -t <i>record_type</i> Select type of record to write for the selected task or process. For tasks, <i>record_type</i> can be <i>partial</i> or <i>interval</i>. <i>partial</i> is the default type, and the only type available for process records.
操作数	The following operands are supported: process Treat the given ID as a process ID for the purposes of constructing and writing an extended accounting record. task Treat the given ID as a task ID for the purposes of constructing and writing an extended accounting record.
示例	示例 1 Writing a Partial Record Write a partial record for all active sendmail processes. <pre># /usr/sbin/wracct -i "pgrep sendmail" process</pre> 示例 2 Writing an Interval Record Write an interval record for the task with ID 182. <pre># /usr/sbin/wracct -t interval -i 182 task</pre>
退出状态	The following exit values are returned: 0 Successful completion. 1 An error occurred. 2 Invalid command line options were specified.

3 Pertinent components of extended accounting facility are not active.

文件

/var/adm/exacct/task

/var/adm/exacct/proc Extended accounting data files.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os

另请参见

[acctadm\(1M\)](#), [attributes\(5\)](#)

引用名	wusbadm – 管理无线 USB 主机和设备
用法概要	<pre> wusbadm list [-h -d] [-o field[,...]] wusbadm associate [-h host-id] [[-c [-f]] -n] [-o] wusbadm remove-dev [[-d dev-id] [-h host-id]] [-f] wusbadm remove-host [-h host-id] [-f] wusbadm enable-host [-h host-id] wusbadm disable-host [-h host-id] </pre>
描述	wusbadm 命令提供了一个可用于管理无线 USB 主机和设备的命令行接口，其功能包括列出主机信息和设备信息、将主机和设备关联起来、删除系统中的主机信息或设备信息以及启用或禁用主机。 首次将无线 USB 设备连接到主机之前，用户需通过运行 wusbadm 的 associate 子命令，在设备和主机之间设置关联信息。然后，用户即可通过直接打开或关闭设备射频（可能是设备上的一个按钮，具体取决于生产商）连接到设备或断开与设备的连接。设备射频的打开和关闭类似于有线 USB 设备的热插拔。 由 associate 子命令创建的关联信息保存在设备和主机的非易失性内存中。在主机上，可通过 remove-dev 或 remove-host 子命令删除这些信息。在设备上，可通过另一关联操作覆盖这些信息。对于与多个主机相关联的设备来说，其确定多个关联记录的优先级或更新这些关联记录的方式取决于生产商。 每个 wusbadm 子命令都作用于下面的某个对象： host-id 一个两位数的数字（其范围为 01 到 99），可唯一标识系统中的无线 USB 主机。该数字是在 wusb 服务（请参见“附注”部分）成功启用并首次找到主机实例的情况下生成的。在由 remove-host 子命令删除之前，该数字将一直保留。 dev-id 一个五位数的数字，可唯一标识与无线 USB 主机关联的无线 USB 设备。头两位数字是与设备关联的无线 USB 主机的 host-id。后三位数字（范围为 001 到 999）用于区分与同一主机关联的设备。在这个五位数中，头两位数与后三位数之间用点隔开。 dev-id 是在设备关联过程中生成的。通过 remove-dev 子命令删除该数字之前，或者通过同一主机和设备之间的另一关联操作对其进行更新之前，系统将始终为设备保留该数字。
子命令	所支持的子命令如下。如果在运行子命令时没有设置任何选项或操作数，则每个子命令都会显示特定于该子命令的用法信息，list 子命令除外。 <pre> list [-h -d] [-o field[,...]] </pre> 列出系统上的无线 USB 主机和设备，同时显示所有主机和设备的 ID、状态及类型。缺省情况下，list 将列出所有主机和设备及所有字段。每个主机及其设备都将显示为一个组。该子命令支持以下选项。

`-o field[,...]`、`--output=field[,...]`

要显示的输出字段的列表，这些字段不区分大小写，并且由逗号分隔。字段名称必须是下面列出的某个字段，或者是可以显示所有字段的特殊值 `all`。缺省情况下（无 `-o`），`list` 会显示所有字段。

ID

host-id 或 *dev-id*。

TYPE

`host`（主机）或 `device`（设备）类型。

对于 `host`（主机）来说，其类型包括 `whci`（板载主机）和 `hwa`（热插拔主机）。

对于 `device`（设备）来说，其类型包括 `kbd`（键盘）、`mouse`（鼠标）、`storage`（存储器）、`printer`（打印机）、`dwa`（无线 USB 集线器）、`audio`（音频）、`video`（视频）等。

STATE

主机有下述几种状态：

`enabled`（已启用）

主机准备好工作或已经在工作，包括执行关联、连接设备、执行数据通信等。

`disabled`（禁用）

主机尚未准备好与任何设备一起工作，且没有设备连接到该主机。出现这种情况的原因，可能是因为已通过 `disable-host` 子命令停用主机，或者主机因插头已拔下或驱动程序分离而无法使用。

`disconnected`（断开）

主机没有连接到系统。将 `hwa` 设备从系统的 USB 端口拔下以后，该设备将处于此状态。

设备有下述几种状态：

`connected`（已连接）

设备已与主机连接，可以打开，或已打开并处于工作状态。缺省情况下，设备会在完成关联并打开射频后尝试进入此状态。

`disconnected`（断开）

设备未连接到主机或尚未做好打开准备。设备之所以会处于这种状态，可能是因为射频超出范围、电源关闭、硬件问题，等等。

`-h, --host`

仅列出无线 USB 主机。

`-d, --device`

仅列出无线 USB 设备。

associate [-h *host-id*] [[-c [-f]] | -n] [-o]

指定主机启动关联过程。在将无线 USB 设备与无线 USB 主机连接之前，初始步骤是关联。

共有两种关联模式：

电缆关联

用户首先使用 USB 电缆连接设备和主机，然后运行此子命令，指定主机设置与设备的关联信息。关联生效以后，在同一主机和设备后续连接中将不再需要该电缆。

数字关联

用户打开设备射频并运行此命令，指定主机与设备通话。然后，主机和设备上会显示一个短数字。用户可比较数字的值并在主机和设备上进行确认。

成功关联以后，关联的 USB 主机和设备即可继续无线连接过程。缺省情况下，在被删除或覆盖之前，关联信息将一直保留在主机和设备上。

如果存在多个可关联的设备，则此子命令会列出所有这些设备以供用户选择。此子命令具有以下选项。

-h *host-id*, --host *host-id*

指定关联操作所针对的主机。如果未指定此选项，此子命令将列出所有已启用的主机供用户选择。

-c, --cable

启动电缆关联过程。用户将无线 USB 设备插入主机，然后运行带有此选项的 **associate** 子命令。

-n, --numeric

启动数字关联过程。此子命令会提示用户比较显示在主机和设备上的数字。

如果前述两个关联模式选项（-n 或 -c）均未指定，此子命令会提示用户指定下述关联模式选项之一。

-f, --force

启动电缆关联过程。用户将无线 USB 设备插入主机，然后运行带有此选项的 **associate** 子命令。

-o, --onetime

表示该关联用于一次性连接。也就是说，在关联以后，如果设备连接然后又断开，则该设备的关联信息将从主机系统中删除。要进行下一次连接，用户需再一次执行关联操作。

remove-dev [[-d *dev-id*] | [-h *host-id*]] [-f]

将无线 USB 设备的关联信息从系统中删除。删除以后，除非用户再一次对主机和设备运行 **associate** 子命令，否则设备将无法与主机连接。此子命令具有以下选项。

-d, --device=*dev-id*

删除 *dev-id* 指定的无线 USB 设备的关联信息。

-h *host-id*, --host=*host-id*

删除与 *host-id* 指定的主机关联的所有无线 USB 设备的关联信息。

-f, --force

在未经确认的情况下执行删除操作。如果设备正在与主机进行连接，则此子命令将强制设备断开连接。

remove-host [-h *host-id*] [-f]

从系统中删除主机信息，包括 *host-id* 以及与主机关联的所有设备的关联信息。此子命令最常用于移除临时使用的热插拔无线 USB 主机，例如 *hwa dongle*。该主机可通过重新枚举的方式恢复，只需以物理方式对 *hwa dongle* 执行热插拔即可。然后系统会对 *host-id* 进行更新，但设备关联信息无法恢复。不建议移除板载主机。此子命令具有以下选项。

-h *host-id*, --host=*host-id*

指定要删除的 *host-id*。

-f, --force

在未经确认的情况下执行删除操作。如果有一个或多个设备与主机建立了连接，则强制其断开连接。

enable-host [-h *host-id*]

使主机进入启用状态。缺省情况下，主机处于启用状态。此子命令具有以下选项。

-h *host-id*, --host=*host-id*

指定要启用的 *host-id*。

disable-host [-h *host-id*] [-f]

使主机进入禁用状态。主机的 *host-id* 和所有关联信息都将保留。发出 *enable-host* 子命令即可令主机恢复启用状态。此子命令具有以下选项。

-h *host-id*, --host=*host-id*

指定要禁用的 *host-id*。

-f, --force

在未经确认的情况下执行禁用操作。如果有一个或多个设备与主机建立了连接，此选项将强制其断开连接。

示例

示例1 列出所有主机和设备

以下命令列出了所有无线 USB 主机和设备。

```
# wusbadm list
01      enabled      hwa
01.001  connected     mouse
01.002  connected     kbd
02      enabled      whci
02.001  connected     printer
```

示例1 列出所有主机和设备 (续)

```
02.002 disconnected storage
03      disabled hwa
03.001 disconnected storage
03.002 disconnected dwa
```

示例2 使用电缆与设备进行关联

以下命令使用电缆关联方法将设备与特定主机 (host-id 01) 关联。

```
# wusbadm associate -h 01 -c
Associate a device with host (01) via cable.
Continue (yes/no)?
```

示例3 删除设备的关联

以下命令从主机系统中删除设备的关联信息。

```
# wusbadm remove-dev -d 01.002
Remove the information of device (01.002) from system.
This device can not be connected with the host until it is associated
again. Continue (yes/no)?
```

示例4 删除所有设备的关联

以下命令删除与特定主机关联的所有设备的关联信息。

```
# wusbadm remove-dev -h 02
Remove the information of all the devices associated with host (02)
from the system.
All the devices associated with the host cannot be connected with it
until they are associated again. Continue (yes/no)?
```

退出状态

将返回下列退出值：

- 0
操作成功。
- 1
错误：操作失败。例如，某个设备无法与主机关联。
- 2
用法错误。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

```
/usr/sbin
```

属性类型	属性值
可用性	system/io/usb
接口稳定性	Committed（已确定）

另请参见

[attributes\(5\)](#)、[hwahc\(7D\)](#)、[usba\(7D\)](#)

附注

wusb（无线 USB 管理）服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

`svc:/system/wusb:default`

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

wusb 服务可通过 wusbd 守护进程这一专用接口实现。与 wusb 服务一样，守护进程由 SMF 启动。使用服务实例指定守护进程：

`svc:/system/wusbd:default`

不应直接调用 wusbd 守护进程。

引用名 ypbind – NIS binder process

用法概要 /usr/lib/netsvcs/yp/ypbind [-broadcast | -ypset | -ypsetme]

描述

NIS provides a simple network lookup service consisting of databases and processes. The databases are stored at the machine that runs an NIS server process. The programmatic interface to NIS is described in [ypclnt\(3NSL\)](#). Administrative tools are described in [ypinit\(1M\)](#), [ypwhich\(1\)](#), and [ypset\(1M\)](#). Tools to see the contents of NIS maps are described in [ypcat\(1\)](#), and [ypmatch\(1\)](#).

ypbind is a daemon process that is activated at system startup time from the `svc:/network/nis/client:default` service. By default, it is invoked as `ypbind -broadcast`. ypbind runs on all client machines that are set up to use NIS. The function of ypbind is to remember information that lets all NIS client processes on a node communicate with some NIS server process. ypbind must run on every machine which has NIS client processes. The NIS server may or may not be running on the same node, but must be running somewhere on the network.

The SMF service `svc:/network/nis/client` has the following properties in the config property group:

```
config.use_broadcast
config.use_ypsetme
```

The information ypbind remembers is called a *binding* — the association of a domain name with a NIS server. The process of binding is driven by client requests. As a request for an unbound domain comes in, if started with the `-broadcast` option, the ypbind process broadcasts on the net trying to find an NIS server, that is, a `ypserv` process serving the domain with a name the same as (case sensitive) the name of the domain in the client request. Since the binding is established by broadcasting, there must be at least one NIS server on the net. If started without the `-broadcast` option, ypbind process steps through the list of NIS servers that was created by `ypinit -c` for the requested domain. There must be an NIS server process on at least one of the hosts in the NIS servers file. It is recommended that you list each of these NIS servers by name and numeric IP address in `/etc/hosts`. Though the practice is not recommended, NIS allows you to list servers by numeric address only, bypassing `/etc/hosts`. In such a configuration, [ypwhich\(1\)](#) returns a numeric address instead of a name.

Once a domain is bound by ypbind, that same binding is given to every client process on the node. The ypbind process on the local node or a remote node may be queried for the binding of a particular domain by using the [ypwhich\(1\)](#) command.

If ypbind is unable to speak to the NIS server process it is bound to, it marks the domain as unbound, tells the client process that the domain is unbound, and tries to bind the domain once again. Requests received for an unbound domain will wait until the requested domain is bound. In general, a bound domain is marked as unbound when the node running the NIS server crashes or gets overloaded. In such a case, ypbind will try to bind to another NIS server using the process described above. ypbind also accepts requests to set its binding for a

particular domain. The request is usually generated by the [ypset\(1M\)](#) command. In order for ypset to work, ypbind must have been invoked with flags -ypset or -ypsetme.

选项

-broadcast
Send a broadcast datagram using UDP/IP that requests the information needed to bind to a specific NIS server. This option is analogous to ypbind with no options in earlier Sun releases and is recommended for ease of use.

Enabling the SMF property `config.use_broadcast` enables -broadcast.

-ypset
Allow users from any remote machine to change the binding by means of the ypset command. By default, no one can change the binding. This option is insecure.

-ypsetme
Only allow root on the local machine to change the binding to a desired server by means of the ypset command. ypbind can verify the caller is indeed a root user by accepting such requests only on the loopback transport. By default, no external process can change the binding.

Enabling the SMF property `config.use_ypsetme` enables -ypsetme.

文件

`/var/yp/binding/ypdomain/ypservers`
Lists the servers to which the NIS client is allowed to bind.

`/etc/inet/hosts`
File in which it is recommended that NIS servers be listed.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/nis

另请参见

[svcs\(1\)](#), [ypcat\(1\)](#), [ypmatch\(1\)](#), [ypwhich\(1\)](#), [ifconfig\(1M\)](#), [svcadm\(1M\)](#), [ypinit\(1M\)](#), [ypset\(1M\)](#), [ypclnt\(3NSL\)](#), [hosts\(4\)](#), [ypfiles\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注

ypbind supports multiple domains. The ypbind process can maintain bindings to several domains and their servers, the default domain is the one specified by the [domainname\(1M\)](#) command at startup time.

The -broadcast option works only on the UDP transport. It is insecure since it trusts “any” machine on the net that responds to the broadcast request and poses itself as an NIS server.

The ypbind service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

`svc:/network/nis/client:default`

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名	ypinit – set up NIS client
用法概要	<code>/usr/sbin/ypinit [-c] [-m] [-s <i>master_server</i>]</code>
描述	ypinit can be used to set up an NIS client system. You must be the superuser to run this command. This script need not be used at all if ypbind(1M) is started with the <code>-broadcast</code> option (it is invoked with this option from the <code>svc:/network/nis/client:default</code> service). Normally, ypinit is run only once after installing the system. It may be run whenever a new NIS server is added to the network or an existing one is decommissioned. ypinit prompts for a list of NIS servers to bind the client to; this list should be ordered from the closest to the furthest server. It is recommended that you list each of these NIS servers by name and numeric IP address in <code>/etc/hosts</code>. Though the practice is not recommended, NIS allows you to list servers by numeric address only, bypassing <code>/etc/hosts</code>. In such a configuration, ypwhich(1) returns a numeric address instead of a name. ypinit stores the list of servers to which a client can bind in the file <code>/var/yp/binding/<i>domain</i>/ypservers</code>. This file is used by ypbind when run without the <code>-broadcast</code> option.
选项	<code>-c</code> Set up a ypclient system. <code>-m</code> Build a master ypserver data base. <code>-s <i>master_server</i></code> Slave data base. <i>master_server</i> must be the same master configured in the YP maps and returned by the <code>ypwhich -m</code> command.
文件	<code>/etc/hosts</code> File in which it is recommended that NIS servers be listed. <code>/var/yp/binding/<i>domain</i>/ypservers</code> Lists the servers to which the NIS client is allowed to bind.
属性	See attributes(5) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/nis

另请参见 [svcs\(1\)](#), [ypwhich\(1\)](#), [svcadm\(1M\)](#), [ypbind\(1M\)](#), [sysinfo\(2\)](#), [hosts\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

附注

The NIS client service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/nis/client:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

已知问题

ypinit sets up the list of NIS servers only for the current domain on the system when it is run, that is, the domain returned by the `SI_SRPC_DOMAIN` command to [sysinfo\(2\)](#). Care should be taken to ensure that this is the same as the desired domain for NIS client processes.

引用名 ypmake – rebuild NIS database

用法概要 cd /var/yp ; make [map]

描述 The file called Makefile in /var/yp is used by [make\(1S\)](#) to build the Network Information Service (NIS) database. With no arguments, make creates dbm databases for any NIS maps that are out-of-date, and then executes [yppush\(1M\)](#) to notify slave databases that there has been a change.

If you supply a *map* on the command line, make will update that map only. Typing make passwd will create and yppush the password database (assuming it is out of date). Likewise, make ipnodes and make networks will create and yppush the ipnodes and network files, \$(INETDIR)/ipnodes and \$(DIR)/networks.

There are four special variables used by make: DIR, which gives the directory of the source files; NOPUSH, which when non-null inhibits doing a yppush of the new database files; INETDIR, which gives the directory of the ipnodes source file; and DOM, which is used to construct a domain other than the master's default domain. The default for DIR is /etc, and the default for INETDIR is /etc/inet. The default for NOPUSH is the null string.

Refer to [ypfiles\(4\)](#) and [ypserv\(1M\)](#) for an overview of the NIS service.

If a NIS to LDAP (N2L) configuration file, /var/yp/NISLDAPmapping, is present, the NIS server components run in N2L mode. In N2L mode, the server components use a new set of map files with an LDAP-prefix, based on the LDAP DIT. In N2L mode, authoritative NIS information is obtained from the DIT. The NIS source files and ypmake have no role, and they should not be used. If ypmake is accidentally run, then the server components will detect this, and will log a warning message. For additional information, see [ypfiles\(4\)](#).

文件 /var/yp Directory containing NIS configuration files.
/etc/inet/hosts System hosts file.
/etc Default directory for source files.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/nis
Interface Stability	Committed

另请参见 [make\(1S\)](#), [makedbm\(1M\)](#), [ypbind\(1M\)](#), [yppush\(1M\)](#), [ypserv\(1M\)](#), [ypc1nt\(3NSL\)](#), [NISLDAPmapping\(4\)](#), [ypfiles\(4\)](#), [ypserv\(4\)](#)

附注

The NIS makefile is only used when running the [ypserv\(1M\)](#) server to provide NIS services. See [ypfiles\(4\)](#) for more details.

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same. Only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications plc, and may not be used without permission.

引用名	ypmap2src – convert NIS maps to NIS source files
用法概要	<pre>/usr/lib/netsvc/yp/ypmap2src [-t] [-c custom-map-name]]... [-d domain] -o output-directory [[source-file]]...</pre>
描述	Use the <code>ypmap2src</code> utility to convert standard NIS maps to approximations of the equivalent NIS source files. This utility functions like the reverse of ypmake(1M). The primary use for <code>ypmap2src</code> is to convert from a NIS server that uses the NIS to LDAP(N2L) transition mechanism, which does not use NIS source files, to traditional NIS, where source files are required. The <code>ypmap2src</code> utility is also used by NIS administrators who wish to discover the contents of NIS maps for which the sources are not available. Generally, this operation is not necessary. More often, administrators will switch from traditional NIS to N2L in anticipation of the eventual transition to LDAP naming. When this switch is made, authoritative information is moved into the LDAP DIT, and the NIS sources have no further role. N2L supports NIS clients until such time as they can be converted to LDAP, and the NIS service suspended. The <code>ypmap2src</code> utility does not guarantee that the files that are generated are identical to the original NIS source files. Some information might have been thrown away by <code>ypmake</code> and cannot be recovered. N2L also might have updated the maps to reflect changes made by LDAP clients. It is essential that the sources generated are checked to confirm no problems have occurred. Per entry comment fields, from existing source files, are not merged into source files generated by <code>ypmap2src</code>. If a user wishes N2L to maintain comment information, then the <code>NISLDAPmapping</code> configuration file should be modified so that the comment fields are mapped into LDAP. This will ensure that the comments are visible to native LDAP clients and present in the N2L map files. When <code>ypmap2src</code> is run, it will take up-to-date comments from the map file and insert them into the NIS source file generated.
Handling Custom Maps	<code>ypmap2src</code> only knows about the standard NIS maps and standard source to map conversion. If an advanced user has changed these, that is, the user has modified the NIS makefile, the equivalent changes must also be made to the <code>ypmap2src</code> script.
选项	<code>ypmap2src</code> supports the following options: -c Specifies that <i>custom-map-name</i> should be converted to a source file by running <code>makedbm -u</code> on it. This is a short cut so that simple custom maps can be handled without editing <code>ypmap2src</code>. -d <i>domain-name</i> Specifies the domain to convert. The <i>domain-name</i> can be a fully qualified file path, such as <code>/var/yp/a.b.c</code>, or just a domain name, <code>a.b.c</code>. In the latter case, <code>ypmap2src</code> looks in <code>/var/yp</code> for the domain directory.

- o *dest* Specifies the destination directory for the converted files. A directory other than /etc should be specified. The maps generated are copied to the correct location, /etc, /etc/security or other source directory, as appropriate.
- t Specifies that traditional NIS maps, without N2L's LDAP_ prefix, should be converted. By default, maps with the LDAP_ prefix are converted.

操作数 ypmap2src supports the following operands:

source-file Lists the standard source files to convert. If this option is not given, then all the standard source files, plus any custom files pecified by the -c option, are converted.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/nis
Interface Stability	Obsolete

另请参见 [ypmake\(1M\)](#), [ypserv\(1M\)](#), [NISLDAPmapping\(4\)](#), [attributes\(5\)](#)
《Oracle Solaris Administration: Naming and Directory Services》

- 引用名

yppoll – return current version of a NIS map at a NIS server host
- 用法概要

```
/usr/sbin/yppoll [-d ypdomain] [-h host] mapname
```
- 描述

The yppoll command asks a ypserv() process what the order number is, and which host is the master NIS server for the named map.
- 选项

-d ypdomain

Use ypdomain instead of the default domain.

-h host

Ask the ypserv process at host about the map parameters. If host is not specified, the NIS server for the local host is used. That is, the default host is the one returned by ypwhich(1).

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/nis

另请参见

[ypwhich\(1\)](#), [ypfiles\(4\)](#), [attributes\(5\)](#)

引用名

用法概要

描述

选项

文件

属性

yppush – force propagation of changed NIS map

```
/usr/lib/netsvc/yp/yppush [-v] [-h host] [-d domain]
    [-p #parallel-xfrs] mapname
```

yppush copies a new version of a Network Information Service (NIS) map from the master NIS server to the slave NIS servers. It is normally run only on the master NIS server by the Makefile in /var/yp after the master databases are changed. It first constructs a list of NIS server hosts by reading the NIS ypservers map within the *domain*. Keys within the ypservers map are the ASCII names of the machines on which the NIS servers run.

A “transfer map” request is sent to the NIS server at each host, along with the information needed by the transfer agent (the program that actually moves the map) to call back the yppush. When the attempt has completed (successfully or not), and the transfer agent has sent yppush a status message, the results can be printed to stdout. Messages are also printed when a transfer is not possible, for instance, when the request message is undeliverable, or when the timeout period on responses has expired.

Refer to [ypfiles\(4\)](#) and [ypserv\(1M\)](#) for an overview of the NIS service.

The following options are supported:

-d <i>domain</i>	Specifies a <i>domain</i> .
-h <i>host</i>	Propagates only to the named <i>host</i> .
-p # <i>parallel-xfrs</i>	Allows the specified number of map transfers to occur in parallel.
-v	Verbose. This prints messages when each server is called, and for each response. If this flag is omitted, only error messages are printed.

/var/yp	Directory where NIS configuration files reside.
/var/yp/ <i>domain</i> /ypservers. <i>{dir, pag}</i>	Map containing list of NIS servers to bind to when running in server mode.

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/nis

另请参见

附注

[ypserv\(1M\)](#), [ypxfr\(1M\)](#), [ypfiles\(4\)](#), [attributes\(5\)](#)

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications PLC, and must not be used without permission.

已知问题

In the current implementation (version 2 NIS protocol), the transfer agent is [ypxfr\(1M\)](#), which is started by the `ypserv` program. If `yppush` detects that it is speaking to a version 1 NIS protocol server, it uses the older protocol, sending a version 1 `YPPROC_GET` request and issues a message to that effect. Unfortunately, there is no way of knowing if or when the map transfer is performed for version 1 servers. `yppush` prints a message saying that an “old-style” message has been sent. The system administrator should later check to see that the transfer has actually taken place.

引用名	ypserv, ypxfrd – NIS server and binder processes
用法概要	<pre>/usr/lib/netsvc/yp/ypserv [-dv] [-i -I] [-r -R]</pre> <pre>/usr/lib/netsvc/yp/ypxfrd</pre>
描述	The Network Information Service (NIS) provides a simple network lookup service consisting of databases and processes. The databases are ndbm files in a directory tree rooted at /var/yp. See ndbm(3C). These files are described in ypfiles(4). The processes are ypserv, the NIS database lookup server, and ypbind, the NIS binder. The programmatic interface to the NIS service is described in ypclnt(3NSL). Administrative tools are described in ypoll(1M), yppush(1M), ypset(1M), ypxfr(1M), and ypwhich(1). Tools to see the contents of NIS maps are described in ypcat(1), and ypmatch(1). Database generation and maintenance tools are described in ypinit(1M), ypmake(1M), and makedbm(1M). The ypserv utility is a daemon process typically activated at system startup from <code>svc:/network/nis/server:default</code>. Alternatively, you can, as the root user, start NIS services using ypstart(1M) from the command-line. ypserv runs only on NIS server machines with a complete NIS database. You can halt all NIS services using the ypstop(1M) command. The ypxfrd utility transfers entire NIS maps in an efficient manner. For systems that use this daemon, map transfers are 10 to 100 times faster, depending on the map. To use this daemon, be sure ypxfrd is running on the master server. See <code>/usr/lib/netsvc/yp/ypstart</code>. ypxfr attempts to use ypxfrd first. If that fails, it prints a warning, then uses the older transfer method. The ypserv daemon's primary function is to look up information in its local database of NIS maps. The operations performed by ypserv are defined for the implementor by the <i>YP Protocol Specification</i>, and for the programmer by the header file <code><rpcsvc/yp_prot.h></code>. Communication to and from ypserv is by means of RPC calls. Lookup functions are described in ypclnt(3NSL), and are supplied as C-callable functions in the libnsl(3LIB) library. There are four lookup functions, all of which are performed on a specified map within some NIS domain: yp_match(3NSL), yp_first(3NSL), yp_next(3NSL), and yp_all(3NSL). The <code>yp_match</code> operation takes a key, and returns the associated value. The <code>yp_first</code> operation returns the first key-value pair from the map, and <code>yp_next</code> can be used to enumerate the remainder. <code>yp_all</code> ships the entire map to the requester as the response to a single RPC request. The SMF service <code>svc:/network/nis/server</code> manages the configuration of the ypserv daemon. A number of special keys in the DBM files can alter the way in which ypserv operates. The keys of interest are:

YP_INTERDOMAIN	The presence of this key causes ypserv to forward to a DNS server host lookups that cannot be satisfied by the DBM files.
YP_SECURE	This key causes ypserv to answer only questions coming from clients on reserved ports.
YP_MULTI_hostname	This is a special key in the form, YP_MULTI_hostname addr1,...,addrN. A client looking for hostname receives the closest address.

Two other functions supply information about the map, rather than map entries: [yp_order\(3NSL\)](#), and [yp_master\(3NSL\)](#). In fact, both order number and master name exist in the map as key-value pairs, but the server will not return either through the normal lookup functions. If you examine the map with [makedbm\(1M\)](#), however, they are visible. Other functions are used within the NIS service subsystem itself, and are not of general interest to NIS clients. These functions include `do_you_serve_this_domain?`, `transfer_map`, and `reinitialize_internal_state`.

On start up, ypserv checks for the existence of the NIS to LDAP (N2L) configuration file `/var/yp/NISLDAPmapping`. If it is present then a master server starts in N2L mode. If the file is not present it starts in “traditional” (non N2L) mode. Slave servers always start in traditional mode.

In N2L mode, a new set of map files, with an LDAP_ prefix, are generated, based on the contents of the LDAP DIT. The old map files, NIS source files and [ypmake\(1M\)](#) are not used.

It is possible that [ypmake\(1M\)](#) can be accidentally run in N2L mode. If the occurs, the old style map files are overwritten. That the map files are overwritten is harmless. However, any resulting [yppush\(1M\)](#) operation will push information based on the DIT rather than the source files. The user may not expect information based on the DIT. ypserv keeps track of the last modification date of the old style map files. If the map files have been updated, a warning is logged that suggests that the user call yppush directly instead of ypmake.

If a server attempts to run in N2L mode and a LDAP server cannot be contacted, it behaves as follows:

1. When ypserv is started, a warning will be logged.
2. When a NIS read access is made and the TTL entry has expired, a warning is logged. Information that is returned from the cache has not been updated.
3. When a NIS write access is made, a warning is logged. The cache will not be updated, and a NIS failure will be returned.

If `ypxfrd` is running in N2L mode and is asked to transfer a map, `ypxfrd` first checks whether the map is out of date. If the map is out of date, `ypxfrd` initiates an update from the DIT. `ypxfrd` cannot wait for the update to complete. If `ypxfrd` waited, the client end `ypxfr`

operation could time out. To prevent `ypxfrd` from timing out, the existing map is transferred from the cache. The most up to date map will be transferred on subsequent `ypxfrd` operations.

选项

ypserv

- d The NIS service should go to the DNS for more host information. This requires the existence of a correct `/etc/resolv.conf` file pointing to a DNS server. This option turns on DNS forwarding regardless of whether or not the `YP_INTERDOMAIN` flag is set in the hosts maps. See [makedbm\(1M\)](#). In the absence of an `/etc/resolv.conf` file, `ypserv` complains, but ignores the `-d` option.

If enabled, the property group/property `config.service_dns` tells `ypserv` to enable the `-d` option.
- i If in N2L mode, initialize the NIS related parts of the DIT based on the current, non LDAP_ prefixed, map files. The LDAP_ prefixed maps are not created or updated. If you require that LDAP_ prefixed maps be updated or created, then use the `-ir` option.

The `-i` option does not attempt to create any NIS domain or container objects. If any NIS domain or container objects have not already been created, then errors will occur, as entries are written to nonexistent containers.
- I Identical to `-i`, except that any missing domain and container objects are created.
- r If in N2L mode, then refresh the LDAP_ prefixed map files based on the contents of the DIT.
- ir If both `-i` and `-r` are specified in N2L mode, then the DIT will first be initialized from the current non LDAP_ prefixed map files. A new set of LDAP_ prefixed maps will then be generated from the contents of the DIT. A new set of LDAP_ prefixed maps is required when moving from traditional NIS to N2L mode NIS.
- Ir Identical to `-ir`, except that any missing domain and container objects are created.
- v Operate in the verbose mode, printing diagnostic messages to `stderr`.

When run with the `-i`, `-r`, `-I`, `-ir` or `-Ir` options, the `ypserv` command runs in the foreground and exits once map initialization has been completed. Once the `ypserv` command exits, the user knows the maps are ready and can restart `ypserv` and the other `yp` daemons by running [ypstart\(1M\)](#).

If there is a requirement to initialize the DIT from the NIS source files, which may have been modified since the maps were last remade, run `ypmake` before running `ypserv -i` or `ypserv -ir`. `ypmake` regenerated old style NIS maps. Then `ypserv -ir` dumps them into the DIT. When the `-ir` option is used, the LDAP_ prefix maps are also generated or updated. Since these maps will be more recent than the old style maps, `ypmake` will not be reported as erroneous when it is run.

文件	/var/yp/securenets Defines the hosts and networks that are granted access to information in the served domain. It is read at startup time by both ypserv and ypxfrd.
	/var/yp/ypserv.log If the /var/yp/ypserv.log file exists when ypserv starts up, log information is written to it when error conditions arise.
	/var/yp/binding/domainname/ypservers Lists the NIS server hosts that ypbind can bind to.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/nis

另请参见 [svcs\(1\)](#), [ypcat\(1\)](#), [ypmatch\(1\)](#), [ypwhich\(1\)](#), [domainname\(1M\)](#), [makedbm\(1M\)](#), [svcadm\(1M\)](#), [ypbind\(1M\)](#), [ypinit\(1M\)](#), [ypmake\(1M\)](#), [yppoll\(1M\)](#), [yppush\(1M\)](#), [ypset\(1M\)](#), [ypstart\(1M\)](#), [ypstop\(1M\)](#), [ypxfr\(1M\)](#), [ndbm\(3C\)](#), [ypclnt\(3NSL\)](#), [libnsl\(3LIB\)](#), [NISLDAPmapping\(4\)](#), [securenets\(4\)](#), [ypfiles\(4\)](#), [ypserv\(4\)](#), [attributes\(5\)](#), [smf\(5\)](#)

《Network Interface Guide》

《Oracle Solaris Administration: Naming and Directory Services》

附注 ypserv supports multiple domains. The ypserv process determines the domains it serves by looking for directories of the same name in the directory /var/yp. It replies to all broadcasts requesting yp service for that domain.

The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications PLC, and must not be used without permission.

NIS uses `ndbm( )` files to store maps. Therefore, it is subject to the 1024 byte limitations described in the USAGE and NOTES sections of the [ndbm\(3C\)](#) man page.

The NIS server service is managed by the service management facility, [smf\(5\)](#), under the service identifier:

```
svc:/network/nis/server:default
```

Administrative actions on this service, such as enabling, disabling, or requesting restart, can be performed using [svcadm\(1M\)](#). The service's status can be queried using the [svcs\(1\)](#) command.

引用名

ypserv_resolv – DNS forwarding service daemon for NIS

用法概要

ypserv_resolv [-v | -V] [-F [-C fd]] [-t xx] [-p yy]

描述

ypserv_resolv is an auxiliary process that provides DNS forwarding service for NIS hosts requests to ypserv that is running in the NIS compatibility mode. It is generally started by invoking [ypserv\(1M\)](#) with the -d option. Although it is not recommended, ypserv_resolv can also be started independently with the options shown below.

This command requires that the /etc/resolv.conf file be setup for communication with a DNS nameserver. The [dig\(1M\)](#) utility can be used to verify communication with a DNS nameserver. See [resolv.conf\(4\)](#) and [dig\(1M\)](#).

选项

-F

Run in foreground.

-C fd

Use fd for service xprt.

-v

Verbose. Send output to the syslog daemon.

-V

Verbose. Send output to stdout.

-t xx

Use transport xx.

-p yy

Use transient program# yy.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/nis

另请参见

[dig\(1M\)](#), [resolv.conf\(4\)](#), [attributes\(5\)](#)

引用名

ypset – point ypbind at a particular server

用法概要

`/usr/sbin/ypset [-d ypdomain] [-h host] server`

描述

In order to run `ypset`, `ypbind` must be initiated with the `–ypset` or `–ypsetme` options. See [ypbind\(1M\)](#). `ypset` tells `ypbind` to get NIS services for the specified *ypdomain* from the `ypserv` process running on *server*. If *server* is down, or is not running `ypserv`, this might not be discovered until an NIS client process tries to obtain a binding for the domain. At this point, the binding set by `ypset` is tested by `ypbind`. If the binding is invalid, `ypbind` attempts to rebind for the same domain.

`ypset` is useful for binding a client node that is not on a broadcast net, or is on a broadcast net that is not running an NIS server host. It is also useful for debugging NIS client applications, for instance, where an NIS map exists only at a single NIS server host.

Where several hosts on the local net are supplying NIS services, `ypbind` can rebind to another host, even while you attempt to find out if the `ypset` operation succeeded. For example, if you enter the `ypset` command below, you might get the subsequent response from `ypwhich`:

```
example% ypset host1
example% ypwhich
host2
```

The sequence shown above is a function of the NIS subsystem's attempt to load-balance among the available NIS servers, and occurs when `host1` does not respond to `ypbind` because it is not running `ypserv` (or is overloaded), and `host2`, running `ypserv`, obtains the binding.

server indicates which NIS server to bind to, and must be specified as a name or an IP address. This works only if the node has a current valid binding for the domain in question and `ypbind` has been set to allow use of `ypset`. In most cases, *server* should be specified as an IP address.

`ypset` tries to bind over a connectionless transport. The NIS library call, `yp_all()`, uses connection-oriented transport and derives the NIS server's address based on the connectionless address supplied by `ypset`.

Refer to [ypfiles\(4\)](#) for an overview of the NIS name service.

选项

- `-d ypdomain` Use *ypdomain*, instead of the default domain.
- `-h host` Set `ypbind`'s binding on *host*, instead of locally. Specify *host* as a name.

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/nis

另请参见

[ypwhich\(1\)](#), [ypfiles\(4\)](#), [attributes\(5\)](#)

引用名 ypstart, yptest – Start and stop NIS services

用法概要 /usr/lib/netsvc/yp/ypstart

/usr/lib/netsvc/yp/ypstop

描述 The ypstart command is used to start the Network Information Service (NIS). After the host has been configured using the [ypinit\(1M\)](#) command, ypstart automatically determines the NIS status of the machine and starts the appropriate daemons.

The ypstop command is used to stop the Network Information Service (NIS).

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/nis

另请参见 [ypinit\(1M\)](#), [attributes\(5\)](#)

《[Oracle Solaris 管理：常见任务](#)》

附注 The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two services remains the same; only the name has changed. The name Yellow Pages is a registered trademark in the United Kingdom of British Telecommunications PLC, and must not be used without permission.

引用名	ypxfr, ypxfr_1perday, ypxfr_1perhour, ypxfr_2perday – transfer NIS map from a NIS server to host
用法概要	<pre>/usr/lib/netsvc/yp/ypxfr [-c] [-f] [-C tid prog server] [-d ypdomain] [-h host] [-s ypdomain] mapname</pre>
描述	The <code>ypxfr</code> command moves an NIS map in the default domain for the local host to the local host by making use of normal NIS services. It creates a temporary map in the directory <code>/var/yp/ypdomain</code> (this directory must already exist; <i>ypdomain</i> is the default domain for the local host), fills it by enumerating the map's entries, fetches the map parameters (master and order number), and loads them. It then deletes any old versions of the map and moves the temporary map to the real <i>name</i>. If run interactively, <code>ypxfr</code> writes its output to the terminal. However, if it is started without a controlling terminal, and if the log file <code>/var/yp/ypxfr.log</code> exists, it appends all its output to that file. Since <code>ypxfr</code> is most often run from the privileged user's crontab file, or by <code>ypserv</code>, the log file can retain a record of what was attempted, and what the results were. For consistency between servers, <code>ypxfr</code> should be run periodically for every map in the NIS data base. Different maps change at different rates: a map might not change for months at a time, for instance, and can therefore be checked only once a day. Some maps might change several times per day. In such a case, you might want to check hourly for updates. A crontab(1) entry can be used to automatically perform periodic updates. Rather than having a separate crontab entry for each map, you can group commands to update several maps in a shell script. Examples (mnemonically named) are in <code>/usr/sbin/yp: ypxfr_1perday</code>, <code>ypxfr_2perday</code>, and <code>ypxfr_1perhour</code>. Refer to ypfiles(4) for an overview of the NIS name service.
选项	-c Do not send a “Clear current map” request to the local <code>ypserv</code> process. Use this flag if <code>ypserv</code> is not running locally at the time you are running <code>ypxfr</code>. Otherwise, <code>ypxfr</code> complains that it cannot communicate with the local <code>ypserv</code>, and the transfer fails. -f Force the transfer to occur even if the version at the master is not more recent than the local version. -C <i>tid prog server</i> This option is for use <i>only</i> by <code>ypserv</code>. When <code>ypserv</code> starts <code>ypxfr</code>, it specifies that <code>ypxfr</code> should call back a <code>yppush</code> process at the host <i>server</i>, registered as program number <i>prog</i>, and waiting for a response to transaction <i>tid</i>. -d <i>ypdomain</i> Specify a domain other than the default domain. -h <i>host</i> Get the map from <i>host</i>, regardless of the master. If <i>host</i> is not specified, <code>ypxfr</code> asks the NIS service for the name of the master, and tries to get the map from there. <i>host</i> must be a valid host name.

	<code>-s ypdomain</code>	Specify a source domain from which to transfer a map that should be the same across domains.
文件	<code>/var/yp/ypxfr.log</code>	Log file
	<code>/usr/lib/netsvc/yp/ypxfr_1perday</code>	Script to run one transfer per day, for use with cron(1M)
	<code>/usr/lib/netsvc/yp/ypxfr_2perday</code>	Script to run two transfer per day, for use with cron(1M)
	<code>/usr/lib/netsvc/yp/ypxfr_1perhour</code>	Script for hourly transfers of volatile maps
	<code>/var/yp/ypdomain</code>	NIS domain
	<code>/usr/spool/cron/crontabs/root</code>	Privileged user's crontab file

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ypxfr Only

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/network/nis

ypxfr_1perday,
ypxfr_1perhour, and
ypxfr_2perday

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	service/network/nis

另请参见 [crontab\(1\)](#), [cron\(1M\)](#), [ypinit\(1M\)](#), [yppush\(1M\)](#), [ypserv\(1M\)](#), [ypfiles\(4\)](#), [attributes\(5\)](#)

引用名	zdb – ZFS 调试器
用法概要	<code>zdb pool</code>
描述	<code>zdb</code> 命令可供支持工程师用于故障诊断以及统计信息的收集。由于磁盘上的 ZFS 文件系统必须始终保持一致且必须进行自我修复，因此应仅在支持工程师的指导下运行 <code>zdb</code>。 在不指定参数的情况下，<code>zdb</code> 将对池和关联的数据集执行基本的一致性检查，然后报告检测到的任何问题。 此命令支持的任何选项均属 Sun 的专用选项，可能会随时更改。
退出状态	将返回下列退出值： 0 池保持一致。 1 检测到错误。 2 指定的命令行选项无效。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	system/file-system/zfs
接口稳定性	Uncommitted（未确定）

另请参见 [zfs\(1M\)](#)、[zpool\(1M\)](#)、[attributes\(5\)](#)

引用名

zdump – time zone dumper

用法概要

zdump [--version] [-v] [-c [loyear,]hiyear] [zonename]...

描述

The zdump command prints the current time for each time zone (*zonename*) listed on the command line. Specify *zonename* as the name of the time zone database file relative to `/usr/share/lib/zoneinfo`.

Specifying an invalid time zone (*zonename*) to zdump does not return an error, rather zdump uses UTC. This is consistent with the behavior of the library calls; zdump reflects the same behavior of the time routines in `libc`. See [ctime\(3C\)](#) and [mktime\(3C\)](#).

选项

The following options are supported:

--version

Outputs version information and exits.

-v

Displays the entire contents of the time zone database file for *zonename*. Prints the time at the lowest possible time value; the time one day after the lowest possible time value; the times both one second before and exactly at each time at which the rules for computing local time change; the time at the highest possible time value; and the time at one day less than the highest possible time value. See [mktime\(3C\)](#) and [ctime\(3C\)](#) for information regarding time value (`time_t`). Each line of output ends with `isdst=1` if the given time is Daylight Saving Time, or `isdst=0` otherwise.

-c [loyear,]hiyear

Cuts off the verbose output near the start of the given year(s) . By default, the program cuts off verbose output near the start of the years -500 and 2500.

退出状态

The following exit values are returned:

0

Successful completion.

1

An error occurred.

文件

`/usr/share/lib/zoneinfo` Standard zone information directory

属性

See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed

另请参见

[zic\(1M\)](#), [ctime\(3C\)](#), [mktime\(3C\)](#), [attributes\(5\)](#), [environ\(5\)](#)

引用名

zfs – 配置 ZFS 文件系统

用法概要

```

zfs [-?]

zfs help subcommand | help | property property-name | permission

zfs create [-p] [-o property=value] ... filesystem

zfs create [-ps] [-b blocksize] [-o property=value] ... -V size volume

zfs destroy [-rRf] filesystem|volume

zfs destroy [-rRd] snapshot

zfs snapshot [-r] [-o property=value]...
    filesystem@snapname|volume@snapname

zfs rollback [-rRf] snapshot

zfs clone [-p] [-K] [-o property=value] ... snapshot filesystem|volume

zfs promote clone-filesystem

zfs rename filesystem|volume|snapshot
    filesystem|volume|snapshot

zfs rename [-p] filesystem|volume filesystem|volume

zfs rename -r snapshot snapshot

zfs list [-rH|-d depth][-o property[,...]] [-t type[,...]]
    [-s property] ... [-S property] ...
    [filesystem|volume|snapshot|path] ...

zfs set [-r] property=value filesystem|volume|snapshot ...

zfs get [-rHp|-d depth][-o all | field[,...]] [-s source[,...]]
    all | property[,...] filesystem|volume|snapshot ...

zfs set share=name=sharename,path=pathname,[desc=description],
    prot=smb | nfs[,property=value[,...]] filesystem

zfs set -c share=name=sharename filesystem

zfs get share [filesystem]

zfs inherit [-rS] property filesystem|volume|snapshot ...

zfs upgrade

zfs upgrade [-v]

zfs upgrade [-r] [-V version] -a | filesystem

zfs userspace [-hniHp] [-o field[,...]] [-sS field] ...
    [-t type [,...]] filesystem|snapshot

zfs groupspace [-hniHpU] [-o field[,...]] [-sS field] ...
    [-t type [,...]] filesystem|snapshot

```

```
zfs mount

zfs mount [-vOS] [-o options] -a | filesystem

zfs unmount [-f] -a | filesystem|mountpoint

zfs share -a | dataset

zfs unshare -a dataset|mountpoint

zfs send [-DvRbp] [-iI] snapshot snapshot

zfs send -r [-Dvbc] [-i snapshot] snapshot

zfs receive [-vnFu] [[-o property=value] | [-x property]] ...
    filesystem|volume|snapshot

zfs receive [-vnFu] [[-o property=value] | [-x property]] ...
    [-d | -e] filesystem

zfs allow filesystem|volume

zfs allow [-ldug] everyone|user|group[,...] perm|@setname[,...]
    filesystem|volume

zfs allow [-ld] -e perm|@setname[,...] filesystem|volume

zfs allow -c perm|@setname[,...] filesystem|volume

zfs allow -s @setname perm|@setname[,...] filesystem|volume

zfs unallow [-rldug] everyone|user|group[,...] [perm|@setname[,... ]]
    filesystem|volume

zfs unallow [-rld] -e [perm|@setname[,... ]] filesystem|volume

zfs unallow [-r] -c [perm|@setname[ ... ]] filesystem|volume

zfs unallow [-r] -s @setname [perm|@setname[,... ]] filesystem|volume

zfs hold [-r] tag snapshot...

zfs holds [-r] snapshot...

zfs release [-r] tag snapshot...

zfs diff [-FHte] [-o field] ... snapshot [snapshot|filesystem]

zfs diff -E [-Fht] [-o field] ... snapshot|filesystem

zfs key -l {-a | [-r] filesystem|volume}

zfs key -u [-f] {-a | [-r] filesystem|volume}

zfs key -c [-o keysource=value] {-a | [-r] filesystem|volume}

zfs key -K {-a | [-r] filesystem|volume}
```

描述

`zfs` 命令按照 `zpool(1M)` 中描述的方式配置 ZFS 存储池中的 ZFS 数据集。数据集由 ZFS 名称空间中的唯一路径进行标识。例如：

```
pool/{filesystem,volume,snapshot}
```

其中，数据集名称的最大长度为 `MAXNAMELEN`（256 字节）。

数据集可以为下列类型之一：

file system

`filesystem` 类型的 ZFS 数据集可以挂载在标准系统名称空间中，并且行为与其他文件系统相似。虽然 ZFS 文件系统旨在符合 POSIX 标准，但是在某些情况下仍然存在一些已知的妨碍符合性的问题。依赖于标准符合性的应用程序可能会在检查文件系统可用空间时由于非标准行为而失败。

volume

作为原始设备或块设备导出的逻辑卷。此类型的数据集只应在特殊情况下使用。文件系统通常适用于大多数环境。

snapshot

文件系统或卷在给定时间点的只读版本。它被指定为 `filesystem@name` 或 `volume@name`。

ZFS 文件系统分层结构

ZFS 存储池是为数据集提供空间的设备的逻辑集合。存储池也是 ZFS 文件系统分层结构的根。

存储池的根可作为文件系统访问，例如挂载和卸载、创建快照以及设置属性。但是，物理存储特征由 `zpool(1M)` 命令管理。

有关创建和管理存储池的更多信息，请参见 `zpool(1M)`。

快照

快照是文件系统或卷的只读副本。快照可以非常快速地创建，而且最初不占用池中的其他空间。当活动数据集中的数据发生变化时，快照占用的数据多于按其他方式与活动数据集共享的数据。

快照可以具有任意名称。卷的快照可进行克隆或回滚，但是不能单独访问。

可从文件系统根目录中的 `.zfs/snapshot` 目录下访问文件系统快照。快照会自动根据需要进行挂载，并可以按固定的间隔卸载。`.zfs` 目录的可见性可由 `snapdir` 属性控制。

克隆

克隆是可写的卷或文件系统，其初始内容与其他数据集相同。与快照一样，创建克隆也几乎是即时的，而且最初不占用其他空间。

克隆只能从快照创建。在克隆快照时，会在父项与子项之间创建隐式相关性。即使克隆是在数据集分层结构中的某个其他位置创建的，但只要存在克隆，就无法销毁原始快照。`origin` 属性显示此相关性，而 `destroy` 命令会列出任何此类相关性（如果存在）。

使用 `promote` 子命令，可以颠倒克隆父项-子项相关性关系。这将导致“原始”文件系统成为所指定文件系统的克隆，从而可能销毁创建克隆时所基于的文件系统。

挂载点

创建 ZFS 文件系统是一项简单的操作，因此每个系统可能有多个文件系统。为处理这种情况，ZFS 自动管理文件系统的挂载和卸载，而不需要编辑 `/etc/vfstab` 文件。所有自动管理的文件系统均由 ZFS 在引导时挂载。

缺省情况下，文件系统挂载在 `/path` 下，其中 `path` 是 ZFS 名称空间中文件系统的名称。可根据需要创建和销毁目录。

文件系统也可以在 `mountpoint` 属性中设置挂载点。此目录会根据需要进行创建；在调用 `zfs mount -a` 命令时，ZFS 将自动挂载文件系统（无需编辑 `/etc/vfstab`）。`mountpoint` 属性可以继承，因此如果 `pool/home` 有一个挂载点 `/export/stuff`，则 `pool/home/user` 会自动继承挂载点 `/export/stuff/user`。

通过为 `zfs mount` 命令指定 `-o mountpoint=value` 选项，可以在文件系统的持久性挂载点以外的位置临时挂载文件系统。仅允许对具有非先前挂载点的文件系统执行此操作。

文件系统的 `mountpoint` 属性 `none` 可防止挂载此文件系统。

如果需要，也可以用传统工具（`mount`、`umount`、`/etc/vfstab`）来管理 ZFS 文件系统。如果将文件系统的挂载点设置为 `legacy`，则 ZFS 不会尝试管理文件系统，并且管理员将负责挂载和卸载文件系统。

区域

添加的文件系统的物理属性由全局管理员控制。但是，区域管理员可以根据添加的文件系统的挂载方式，在该文件系统中创建、修改或销毁文件。

使用 `zonecfg add dataset` 子命令也可以将数据集委托给非全局区域。您不能将数据集委托给一个区域，也不能将同一数据集的子项委托给另一个区域。区域管理员可以更改数据集或其任何子项的属性。但是，`quota` 属性由全局管理员控制。

通过 `zonecfg add device` 子命令，可以将 ZFS 卷作为一个设备添加到非全局区域。但是，只有全局管理员才可修改其物理属性。

有关 `zonecfg` 语法的更多信息，请参见 [zonecfg\(1M\)](#)。

将数据集委托给非全局区域后，将自动设置 `zoned` 属性。在全局区域中，只能通过使用临时的 `mountpoint` 属性（参见“临时挂载点属性”）挂载区域文件系统。

全局管理员可强制清除 `zoned` 属性，但是完成此操作应非常谨慎。全局管理员应在清除该属性前确认所有挂载点是可接受的。

重复数据删除

重复数据删除是在块级别删除冗余数据以减少存储的总数据量的过程。重复数据删除面向存储池范围；每个数据集都可以使用自己的 `dedup` 属性增减内容。如果文件系统启用了 `dedup` 属性，则在写入时将同步删除重复数据块。因此，在存储池已启用 `dedup` 的所有数据集中只会存储唯一的数据，并且会在这些数据集的文件中共享公共组件。

加密	有关 ZFS 加密和 ZFS 加密语法的完整描述，请参见 zfs_encrypt(1M) 。
本机属性	属性分为两种类型：本机属性和用户定义的（或“用户”）属性。本机属性用于提供内部统计信息或控制 ZFS 行为。此外，本机属性分为可编辑属性和只读属性。用户属性对 ZFS 行为没有影响，但可通过用户环境中有意义的方式来注释数据集。有关用户属性的更多信息，请参见下文的“用户属性”一节。 每个数据集都有一组属性，用于提供有关数据集的统计信息并控制各种行为。属性继承自父项，但子项所覆盖的除外。有些属性仅适用于某些类型的数据集（文件系统、卷或快照）。 数值属性的值可使用便于阅读的后缀（例如 k、KB、M、Gb 等，直到表示泽字节的 Z）来指定。下面都是有效（且等效）的指定： 1536M, 1.5g, 1.50GB 非数值属性的值区分大小写，并且必须为小写形式，mountpoint 属性除外。 以下本机属性包含数据集的相关只读统计信息。既不能设置这些属性，也不能继承这些属性。除非另有说明，否则本机属性适用于所有数据集类型。 available 数据集及其所有子项的可用空间的数量（假定池中没有任何其他活动）。由于池中会共享空间，因此可用空间会受到许多因素的限制，包括物理池大小、配额、预留空间或池中的其他数据集。 此属性也可通过其简短列名 <code>avail</code> 来引用。 compressratio 为此数据集实现的压缩率，以倍数形式表示。通过运行 <code>zfs set compression=on dataset</code> 可以启用压缩。缺省值为 <code>off</code>。 creation 创建此数据集的时间。 defer_destroy 如果已使用 <code>zfs destroy -d</code> 命令将快照标记为延期销毁，则此属性为 <code>on</code>。否则，此属性为 <code>off</code>。 keystatus 有关更多信息，请参见 zfs_encrypt(1M)。 mounted 对于文件系统，指示当前是否已挂载文件系统。此属性可以是 <code>yes</code> 或 <code>no</code>。 origin 对于克隆的文件系统或卷，指示创建克隆时所基于的快照。只要克隆存在，便不能销毁克隆源（即使使用 <code>-r</code> 或 <code>-f</code> 选项也是如此）。

referenced

此数据集可访问的数据量，这些数据可以与池中其他数据集共享，也可以不与其他数据集共享。创建快照或克隆时，首先会引用与创建时所基于的文件系统或快照相同的空间量，因为其内容相同。

该属性也可通过其简短列名 **refer** 来引用。

rekeydate

有关更多信息，请参见 [zfs_encrypt\(1M\)](#)。

type

数据集类型：**filesystem**、**volume** 或 **snapshot**。

used

此数据集及其所有后代占用的空间量。这是根据此数据集的配额和预留空间进行检查的值。使用的空间不包括此数据集的预留空间，但是会考虑 **refreservation**（通过 **usedbyrefreservation**）以及任何后代数据集的预留空间（通过 **usedbychildren**）。数据集占用其父项的空间量以及以递归方式销毁此数据集时所释放的空间量，应为其使用空间和预留空间的较大者。

创建快照时（参见“快照”一节），其空间最初将在快照与文件系统之间共享，也可能与以前的快照共享。随着文件系统的变化，以前共享的空间将供快照专用，并会计算在快照的使用空间内。此外，删除快照可增加其他快照专用（和使用）的空间量。

使用的空间量、可用的空间量或引用的空间量不会考虑暂挂更改。通常，暂挂更改仅占用几秒钟的时间。使用 **fsync(3C)** 或 **O_SYNC** 提交对磁盘的更改，不一定可以保证空间使用情况信息会立即更新。

usedby*

usedby* 属性将 **used** 属性分解为使用该空间的各种原因。具体而言，**used** = **usedbychildren** + **usedbydataset** + **usedbyrefreservation** + **usedbysnapshots**。这些属性只能用于在“版本 13”或更高版本的池上创建的数据集。

usedbychildren

此数据集的子项使用的空间量，如果销毁此数据集的所有子项，将释放该空间。

usedbydataset

此数据集本身使用的空间量，如果销毁此数据集，将释放该空间（首次删除任何 **refreservation** 并销毁任何必需的快照或后代之后）。

usedbyrefreservation

针对此数据集设置的 **refreservation** 占用的空间量，如果删除 **refreservation**，将释放该空间。

此属性所占用的空间表示未来写操作可能消耗的空间，预留以防此数据集出现写分配故障。这可能包括未写入的数据、当前与快照共享的空间，以及卷的压缩节省空间（当用更少的可压缩数据替换时可能丢失）。当为后续写操作分配空间增加 **usedbydataset** 或 **usedbysnapshots** 时，将相应减少 **usedbyrefreservation**。

usedbysnapshots

此数据集的快照占用的空间量。特别地，此空间量是销毁此数据集的所有快照后将释放的空间量。请注意，此空间量不是快照的 **used** 属性的简单总和，因为空间可被多个快照共享。

userused@user

此数据集中指定的用户占用的空间量。空间由每个文件的所有者占用，如 **ls -l** 显示的内容所示。占用的空间量由 **du** 和 **ls -s** 显示。有关更多信息，请参见 **zfs userspace** 子命令。

未授权的用户只能访问各自的空间使用情况。**root** 用户或已用 **zfs allow** 授予 **userused** 特权的用户，可以访问每个人的空间使用情况。

zfs get all 不会显示 **userused@...** 属性。必须使用以下形式之一在 **@** 符号后加上用户名：

- *POSIX name*（例如 **joe**）
- *POSIX numeric ID*（例如 **789**）
- *SID name*（例如 **joe.smith@mydomain**）
- *SID numeric ID*（例如 **S-1-123-456-789**）

userrefs

为此快照上的用户保留量可设置此属性。通过使用 **zfs hold** 命令可设置用户保留量。

groupused@group

此数据集中指定的组占用的空间量。空间由每个文件的组所占用，如 **ls -l** 显示的内容所示。有关更多信息，请参见 **userused@user** 属性。

未授权的用户只能访问各自的组的空间使用情况。**root** 用户或已用 **zfs allow** 授予 **groupused** 特权的用户，可以访问所有组的空间使用情况。

volblocksize=blocksize

可为卷指定卷的块大小。在写入卷后无法更改 **blocksize**，因此应在创建卷时设置该值。卷的缺省 **blocksize** 是 8 KB。范围位于 512 字节到 1 MB 之间的 2 的任意次幂都有效。

该属性也可通过其简短列名 **volblock** 来引用。

使用以下本机属性可以更改 ZFS 数据集的行为。

aclmode=discard|mask|passthrough

控制 **chmod(2)** 执行期间如何修改 ACL。带有 **aclmode** 属性 **discard**（缺省值）的文件系统删除不表示文件模式的所有 ACL 条目。**aclmode** 属性 **mask** 减少用户或组的权限。除非用户项与文件或目录的属主具有相同的 UID，否则将减少权限，以使其不会大于组权限位。在这种情况下，减少 ACL 权限，以使其不会大于属主权限位。**mask** 还跨模式更改保留 ACL（两次模式更改之间没有进行显式 ACL 设置 [通过 **chmod(1)**]）。带有 **aclmode** 属性 **passthrough** 的文件系统指示除了生成必要的 ACL 条目来表示文件或目录的新模式以外，不会对 ACL 进行其他更改。

`aclinherit=discard|noallow|restricted|passthrough|passthrough-x`

控制创建文件和目录时继承 ACL 项的方法。aclinherit 属性为 discard 的文件系统不继承任何 ACL 项。aclinherit 属性值为 noallow 的文件系统仅继承指定了“拒绝”权限的可继承 ACL 项。在继承 ACL 项时，属性值 restricted（缺省值）会删除 write_acl 和 write_owner 权限。如果文件系统的 aclinherit 属性值为 passthrough，则会继承所有可继承的 ACL 项，并且继承 ACL 项时不会对其进行任何修改。如果文件系统的 aclinherit 属性值为 passthrough-x，则与 passthrough 具有同样的意义，区别在于仅当文件创建模式也需要执行位时，所有 ACE 均继承执行权限。

当属性值设置为 passthrough 时，由可继承的 ACE 确定的模式来创建文件。如果不存在会影响模式的可继承的 ACE，则按照应用程序要求的模式来设置模式。

`atime=on|off`

控制文件被读取后是否更新该文件的访问时间。禁用该属性可避免在读取文件时产生写入流量，因此可显著提高性能，但可能会使邮件程序与其他相似的实用程序感到困惑。缺省值为 on。

`canmount=on|off|noauto`

如果此属性设置为 off，则无法挂载文件系统，并且 `zfs mount -a` 将忽略此文件系统。将此属性设置为 off 与将 mountpoint 属性设置为 none 的效果相似，区别在于数据集仍有一个可继承的正常 mountpoint 属性。通过将此属性设置为 off，可以将数据集单独用作一种继承属性的机制。例如，设置 canmount=off 将使两个数据集具有相同的 mountpoint，因此两个数据集的子项会出现在同一个目录中，但是可能具有不同的继承特征。

设置 noauto 选项时，只能显式挂载和取消挂载数据集。数据集不会在创建或导入时自动挂载，也不能通过 `zfs mount -a` 命令挂载或通过 `zfs unmount -a` 命令取消挂载。

不会继承此属性。

`checksum=on|off|fletcher2,|fletcher4|sha256|sha256+mac`

控制用于验证数据完整性的校验和。缺省值是 on，它将自动选择正确的算法（当前为 fletcher4，但是在将来的版本中可能改变）。值为 off 将禁止对用户数据进行完整性检查。实际操作中建议不要禁用校验和。

更改此属性只会影响新写入的数据。

只有在启用加密时，值 sha256+mac 才可用。在启用加密时，校验和属性变为 readonly，然后始终将其设置为 sha256+mac。

`compression=on|off|lzjb|gzip|gzip-N|zle`

控制用于此数据集的压缩算法。在提供适当的数据压缩时，将为 lzjb 压缩算法优化性能。将压缩设置为 on 会使用 lzjb 压缩算法。gzip 压缩算法使用与 `gzip(1)` 命令相同的压缩。通过使用值 gzip-N 可以指定 gzip 级别，其中 N 是一个从 1（最快）至 9（最佳压缩率）的整数。当前，gzip 等效于 gzip-6（这也是 gzip(1) 的缺省值）。

该属性也可通过其简短列名 `compress` 来引用。更改此属性只会影响新写入的数据。

`copies=1|2|3`

控制为此数据集存储的数据副本数量。除了作为池提供的任何冗余，这些副本还可作为镜像或 RAID-Z。如果可能，这些副本会存储在不同磁盘上。多个副本使用的空间由相关联的文件和数据集占用，按照配额和预留空间更改 `used` 属性和计算方式。

更改此属性只会影响新写入的数据。因此，在创建文件系统时通过使用 `-o copies=N` 选项设置此属性。

在数据集上启用加密时，可设置的最大副本数量是 2。

`dedup=on|off|verify|sha256[,verify]`

控制重复数据删除对数据集是否有效。缺省值为 `off`。用于重复数据删除的缺省校验和是 `sha256`（可能会更改）。在启用 `dedup` 时，`dedup` 校验和算法会覆盖 `checksum` 属性。将值设置为 `verify` 等效于指定 `sha256,verify`。

如果将属性设置为 `verify`，则只要两个块具有相同的签名，ZFS 都会与现有块进行逐字节比较，以确保内容完全相同。

`devices=on|off`

控制是否可在此文件系统中打开设备节点。缺省值为 `on`。

`exec=on|off`

控制是否可从此文件系统中执行进程。缺省值为 `on`。

`logbias=latency|throughput`

控制 ZFS 如何为此数据集优化同步请求。如果 `logbias` 设置为 `latency`，ZFS 将使用池的不同日志设备（如有）低延迟地处理请求。如果 `logbias` 设置为 `throughput`，ZFS 将不使用池的不同日志设备。相反，ZFS 将优化同步操作，以提高池的全局吞吐量并有效使用资源。缺省值为 `latency`。

`mlslabel=label|none`

`mlslabel` 属性是一个敏感标签，用于确定能否在启用了 Trusted Extensions 的系统的区域中挂载数据集。如果有标签的数据集与有标签的区域相符，则可以从有标签的区域挂载和访问数据集。

如果未设置 `mlslabel` 属性，缺省值将为 `none`。将 `mlslabel` 属性设置为 `none` 等效于删除该属性。

只有在启用了 Trusted Extensions 并且具有相应的特权时，才能修改 `mlslabel` 属性。不可委托用于修改此属性的权限。将标签更改为较高级别标签或设置初始数据集标签时，需要 `{PRIV_FILE_UPGRADE_SL}` 特权。将标签更改为较低级别标签或缺省标签 (`none`) 时，需要 `{PRIV_FILE_DOWNGRADE_SL}` 特权。只有在未挂载数据集时，才能将数据集更改为缺省值以外的标签。将具有缺省标签的数据集挂载到有标签区域时，挂载操作会自动将 `mlslabel` 属性设置为该区域的标签。

如果未启用 Trusted Extensions，则只能挂载具有缺省标签 (`none`) 的数据集。

mountpoint=*path* | **none** | **legacy**

控制用于此文件系统的挂载点。有关如何使用此属性的更多信息，请参见“挂载点”一节。

当文件系统的 **mountpoint** 属性发生更改时，将取消挂载该文件系统以及继承挂载点的任何子项。如果新值为 **legacy**，则该文件系统和子项将保持取消挂载状态。否则，如果属性以前为 **legacy** 或 **none**，或者该文件系统和子项在属性发生更改之前处于挂载状态，则会自动在新位置重新挂载它们。此外，任何共享文件系统都将取消共享，并在新位置进行共享。

nbmand=**on** | **off**

有关更多信息，请参见 **zfs_share(1M)**。

primarycache=**all** | **none** | **metadata**

控制在主高速缓冲存储器 (ARC) 中缓存的内容。如果将此属性设置为 **all**，则会高速缓存用户数据和元数据。如果将此属性设置为 **none**，则不会高速缓存用户数据和元数据。如果将此属性设置为 **metadata**，则只会高速缓存元数据。缺省值为 **off**。

quota=*size* | **none**

限制数据集及其后代可占用的空间量。此属性对使用的空间量强制实施硬限制。这包括后代（含文件系统和快照）占用的所有空间。对已有配额的数据集的后代设置配额不会覆盖祖先的配额，但会施加额外的限制。

不能对卷设置配额，因为 **volsize** 属性可用作隐式配额。

sync=**standard** | **always** | **disabled**

确定文件系统事务的同步程度。在创建数据集时可设置此属性，也可动态设置此属性，并且属性将会立即生效。此属性可能具有以下设置之一：

standard

缺省选项。将同步文件系统事务写入到意图日志，然后刷新写入的所有设备，以确保数据稳定（即设备控制器未进行高速缓存）。

always

将每个文件系统事务写入到稳定的存储器中并刷新。此值对性能具有重大不利影响，但是适合为同步文件系统事务排除故障。

disabled

禁用同步请求。仅当下一次提交 DMU 事务组时，文件系统事务才会提交至稳定的存储器，这可能会延迟好几秒钟。此设置提供最高性能。但是，这是非常危险的，因为 ZFS 会忽略应用程序的同步事务要求，例如数据库或 NFS。而且，当此设置对当前活动的根或 **/var** 文件系统生效时，可能导致行为不合规、应用程序数据丢失以及增加反复攻击的漏洞。管理员只应在了解这些风险时才使用此选项。

`userquota@user=size | none`

限制指定的用户占用的空间量。与 `refquota` 属性相似，计算 `userquota` 空间不包括后代数据集（例如快照和克隆）使用的空间。用户空间占用由 `userspace@user` 属性标识。

用户配额的实施可能会延迟几秒钟。这种延迟意味着在系统注意到用户超出配额前，用户可以超出其配额。系统然后会开始拒绝额外的写入内容，并提示 `EDQUOT` 错误消息。有关更多信息，请参见 `zfs userspace` 子命令。

未授权的用户只能访问各自的组的空间使用情况。`root` 用户或已用 `zfs allow` 授予 `userquota` 特权的用户，可以获取和设置每个人的配额。

对于卷、第 4 版以前的文件系统或第 15 版以前的池，此属性均不可用。`zfs get all` 不会显示 `userquota@...` 属性。必须使用以下形式之一在 `@` 符号后加上用户名：

- *POSIX name*（例如 `joe`）
- *POSIX numeric ID*（例如 `789`）
- *SID name*（例如 `joe.smith@mydomain`）
- *SID numeric ID*（例如 `S-1-123-456-789`）

`groupquota@group=size | none`

限制指定的组占用的空间量。组空间占用由 `userspace@user` 属性标识。

未授权的用户只能访问各自的组的空间使用情况。`root` 用户或已用 `zfs allow` 授予 `groupquota` 特权的用户，可以获取和设置所有组的配额。

`readonly=on | off`

控制是否可以修改此数据集。缺省值为 `off`。

该属性也可通过其简短列名 `rdonly` 来引用。

`recordsize=size`

为文件系统中的文件指定建议的块大小。该属性专门设计用于对大小固定的记录中的文件进行访问的数据库工作负荷。ZFS 会根据为典型的访问模式优化的内部算法来自动调整块大小。

对于创建很大的文件但访问较小的随机块中的文件的数据库而言，这些算法可能不是最优的。将 `recordsize` 指定为大于或等于数据库的记录大小的值可以显著提高性能。强烈建议不要将该属性用于一般用途的文件系统，否则可能会对性能产生不利影响。

确实记录大小是 128 KB。指定的大小必须是 2 的若干次幂，并且必须大于或等于 512 字节同时小于或等于 1 MB。

更改文件系统的 `recordsize` 只影响以后创建的文件，现有文件和接收的数据不受影响。

该属性也可通过其简短列名 `recsize` 来引用。

refquota=size | none

限制数据集可以占用的空间量。此属性对使用的空间量强制实施硬限制。这种硬限制不包括后代以及文件系统和快照使用的空间。

refreservation=size | none

为数据集（不含其后代）预留的最小空间量。如果 **usedbydataset** 空间低于此值，则认为数据集正在使用 **refreservation** 指定的空间量。**usedbyrefreservation** 数字表示该额外空间，并添加到数据集占用的总 **used** 空间，继而占用父数据集的用量、配额和预留空间。这将通过确保提前预留用于未来写操作的空间，防止数据集过度承载池资源。

与快照共享的空间此后将被新数据替换，快照表示承诺保留两份副本。如果设置了 **refreservation**，则必须考虑此承诺将 **usedbyrefreservation** 增加到拍摄新快照时 **refreservation** 的最大空间。如果对此增量没有足够空间可用于数据集，则将拒绝创建快照。

该属性也可通过其简短列名 **refreserv** 来引用。

reservation=size | none

为数据集及其后代预留的最小空间量。如果使用的空间量低于该值，则认为数据集正在占用其预留空间指定的空间量。父数据集的使用空间中会包含预留空间，并会针对父数据集的配额和预留空间对其进行计数。

该属性也可通过其简短列名 **reserv** 来引用。

rstchown=on | off

指示文件系统是否通过 **chown(1)** 或 **chown(2)** 系统调用来限制用户放弃其文件。缺省值是限制 **chown**。如果 **rstchown** 为 **off**，则 **chown** 的行为将如同用户具有 **PRIV_FILE_CHOWN_SELF** 特权一样。

secondarycache=all | none | metadata

控制在二级高速缓存 (L2ARC) 中缓存的内容。如果将此属性设置为 **all**，则会高速缓存用户数据和元数据。如果将此属性设置为 **none**，则不会高速缓存用户数据和元数据。如果将此属性设置为 **metadata**，则只会高速缓存元数据。缺省值为 **all**。

setuid=on | off

控制文件系统是否遵照 **set-UID** 位。缺省值为 **on**。

shadow=URI | none

将 ZFS 文件系统标识为 **URI** 所描述的文件系统的“影子”。设置该属性时，数据从 **URI** 标识的文件系统迁移到影子文件系统。将迁移的文件系统必须为只读才能进行完整迁移。

对影子文件系统中尚未迁移的目录的访问会被阻止，直到迁移了整个目录。对影子文件系统中尚未迁移的文件进行访问将导致只迁移被访问文件的一部分。多个进程可以同时迁移一个文件的不同部分。

接受两种格式的 **URI**，一种用于将本地文件系统迁移到同一物理系统上的其他文件系统，另一个用于从 NFS 服务器远程迁移文件系统。这两种格式为：

```
file:///path
nfs://host/path
```

如果迁移完成时 **shadowd(1M)** 仍在运行，则文件系统将自动重新挂载（影子属性设置为 **none**）。或者，迁移完成时，您可以手动将影子属性设置为 **none**。

```
sharenfs=on|off
```

有关更多信息，请参见 **zfs_share(1M)**。

```
sharesmb=on|off
```

有关更多信息，请参见 **zfs_share(1M)**。

```
snapdir=hidden|visible
```

按照“快照”一节的讨论，控制在文件系统的根中是否隐藏或显示 **.zfs** 目录。缺省值为 **hidden**。

```
version=1|2|current
```

此文件系统的盘上版本，与池版本无关。只可以为后续支持的版本设置此属性。设置 **zfs upgrade** 命令。

```
volsize=size
```

指定卷的逻辑大小。缺省情况下，创建卷将建立稍大于实际逻辑卷大小的 **refreservation**，以用于 ZFS 源数据开销。对 **volsize** 的任何更改都会反映为对 **refreservation** 的等效更改。**volsize** 只可以设置为多个 **volblocksize**，并且不能为零。

在卷上设置 **refreservation**，以防止使用者的意外行为。如果没有预留空间，则卷可能用完空间，导致不明确的行为或数据损坏，具体取决于卷的使用方式。如果在卷的使用过程中（特别是在缩减大小时）更改卷大小，也会出现上述影响。调整卷大小时，需要格外小心。

可以通过指定 **zfs create -V** 命令的 **-s** 选项来创建“稀疏卷”（也称为“瘦置备”），但并不建议这样做。“稀疏卷”是预留空间小于卷大小的卷。因此，当池空间不足时，对稀疏卷的写操作可能因 **ENOSPC** 而失败。对于稀疏卷，预留空间中不会反映对 **volsize** 的更改。

```
vscan=on|off
```

有关更多信息，请参见 **zfs_share(1M)**。

```
xattr=on|off
```

控制是否为此文件系统启用扩展属性。缺省值为 **on**。

```
zoned=on|off
```

控制是否从非全局区域管理数据集。有关更多信息，请参见“区域”一节。缺省值为 **off**。

在创建文件系统后无法更改以下属性，因此在创建文件系统时应设置这些属性。如果未使用 **zfs create** 或 **zpool create** 命令设置属性，则从父数据集继承这些属性。如果在这些功能获得支持以前已创建的父数据集缺少这些属性，则新文件系统将为这些属性采用缺省值。

casesensitivity=sensitive | insensitive | mixed
有关更多信息，请参见 [zfs_share\(1M\)](#)。

normalization = none | formC | formD | formKC | formKD
有关更多信息，请参见 [zfs_share\(1M\)](#)。

utf8only=on | off
有关更多信息，请参见 [zfs_share\(1M\)](#)。

encryption=off | on | aes-128-ccm | aes-192-ccm | aes-256-ccm | aes-128-gcm |
aes-192-gcm | aes-256-gcm
有关更多信息，请参见 [zfs_encrypt\(1M\)](#)。

在创建时必须指定以下属性，并可以使用特殊命令修改该属性：

keysource=raw | hex | passphrase,prompt | file
有关更多信息，请参见 [zfs_encrypt\(1M\)](#)。

临时挂载点属性

通过传统 [mount\(1M\)](#) 命令或 `zfs mount` 命令挂载文件系统时，将会根据其属性设置其挂载选项。属性与挂载选项之间的相互关系如下：

PROPERTY	MOUNT OPTION
devices	devices/nodevices
exec	exec/noexec
readonly	ro/rw
setuid	setuid/nosetuid
xattr	xattr/noxattr
rstchown	rstchown/norstchown

此外，可以使用 `-o` 选项在每次挂载基础上设置这些选项，而不影响磁盘上存储的属性。命令行上指定的值会覆盖数据集中存储的值。`-nosuid` 选项是 `nodevices,nosetuid` 的别名。`zfs get` 命令将这些属性作为“临时”属性进行报告。对于 `mountpoint` 以外的属性，如果在挂载数据集时更改属性，则新设置会覆盖任何临时设置。在临时挂载点属性生效时（即在临时位置挂载数据集时），无法更改 `mountpoint` 属性。

用户属性

除了标准的本机属性外，ZFS 还支持任意用户属性。用户属性对 ZFS 行为没有影响，但是应用程序或管理员可以使用这些属性来注释数据集（文件系统、卷和快照）。

用户属性名称必须含有一个冒号 (:) 字符，以便与本机属性区分。这些名称可包含小写字母、数字和以下标点符号：冒号 (:)、破折号 (-)、句号 (.) 和下划线 (_)。预期约定是将属性名称分为两部分，例如 `module:property`，但是 ZFS 并不强制执行这种名称空间。用户属性名称最多可有 256 个字符，不能以破折号 (-) 开头。

在程序中使用用户属性时，强烈建议对属性名的 `module` 部分使用反向 DNS 域名，以避免两个独立开发的软件包将同一属性名用于不同用途。在 Oracle Solaris 发行版中，为 `beadm` 命令和库保留了 `com.oracle` 用户属性。

用户属性的值是任意字符串，始终会继承但是从不会验证这些值。对属性执行操作的所有命令（如 `zfs list`、`zfs get`、`zfs set` 等）都可用来处理本机属性和用户属性。使用 `zfs inherit` 命令可清除用户属性。如果任意父数据集集中均未定义该属性，则会将其完全删除。属性值仅限于 1024 个字符。

ZFS 卷作为交换设备或转储设备

初始安装期间，将会在 ZFS 根池的 ZFS 卷上创建交换设备和转储设备。必须将单独的 ZFS 卷用于交换区域和转储设备。在 ZFS 文件系统中，不要交换到文件。不支持 ZFS 交换文件配置。

通过为该设备指定 `encryption` 属性，以及在 `vfstab(4)` 中指定 `encrypted` 选项，可以对用作交换设备的 ZFS 卷进行加密。有关加密属性的更多信息，请参见 `zfs_encrypt.1m`。

如果您需要在安装或升级系统后更改交换区域或转储设备，请使用 `swap(1M)` 和 `dumpadm(1M)` 命令。如果需要更改交换区域或转储设备的大小，请参见《Oracle Solaris 管理：ZFS 文件系统》

子命令

修改状态的所有子命令将以原始形式永久记录到池中。

`zfs ?`

显示帮助消息。

`zfs help command | help | property property-name | permission`

显示 `zfs` 命令的用法信息。可以显示特定命令、属性或委托权限的帮助。如果显示特定命令或属性的帮助，将显示命令语法或属性值。在不带任何参数的情况下使用 `zfs help` 将显示完整的 `zfs` 命令列表。

`zfs create [-p] [-o property=value] ... filesystem`

创建新的 ZFS 文件系统。该文件系统会根据从父项继承的 `mountpoint` 属性自动挂载。

`-p`

创建所有非现有的父数据集。以这种方式创建的数据集会根据从父项继承的 `mountpoint` 属性自动挂载。在命令行上使用 `-o` 选项指定的任何属性将被忽略。如果目标文件系统已经存在，则会成功完成此操作。

`-o property=value`

与在创建数据集时调用命令 `zfs set property=value` 一样设置指定的属性。在创建时也可设置任何可编辑的 ZFS 属性。可以指定多个 `-o` 选项。如果在多个 `-o` 选项中指定了相同的属性，则会导致错误。

`zfs create [-ps] [-b blocksize] [-o property=value] ... -V size volume`

创建一个给定大小的卷。该卷作为 `/dev/zvol/{dsk,rdsk}/path` 中的一个块设备导出，其中 `path` 是 ZFS 名称空间中卷的名称。`size` 表示设备导出的逻辑大小。缺省情况下，创建一个同等大小的预留空间。

`size` 自动舍入到最接近的 128 KB，确保无论 `blocksize` 有多大，卷具有的块数量均为整数。

- p
创建所有非现有的父数据集。以这种方式创建的数据集会根据从父项继承的 `mountpoint` 属性自动挂载。在命令行上使用 `-o` 选项指定的任何属性将被忽略。如果目标文件系统已经存在，则会成功完成此操作。
- s
创建一个不带预留空间的稀疏卷。有关稀疏卷的更多信息，请参见“本机属性”一节中的 `volsize`。
- o *property=value*
与在创建数据集时调用 `zfs set property=value` 命令一样设置指定的属性。在创建时也可设置任何可编辑的 ZFS 属性。可以指定多个 `-o` 选项。如果在多个 `-o` 选项中指定了相同的属性，则会导致错误。
- b *blocksize*
相当于 `-o volblocksize=blocksize`。如果结合 `-o volblocksize` 指定此选项，则产生的行为不确定。

zfs destroy [-rRf] filesystem|volume

销毁给定的数据集。缺省情况下，此命令会取消共享当前已共享的任何文件系统，卸载当前已挂载的任何文件系统，并拒绝销毁具有活动相关项（子项或克隆）的数据集。

- r
以递归方式销毁所有子项。
- R
以递归方式销毁所有相关项，包括在目标分层结构以外克隆的文件系统。
- f
强制使用 `umount -f` 命令卸载任何文件系统。此选项对非文件系统或已卸载的文件系统没有影响。

在应用 `-r` 或 `-f` 选项时应非常谨慎，因为这些选项可以销毁大部分池，并导致正在使用的挂载文件系统出现意外行为。

zfs destroy [-rRd] snapshot

当且仅当不带 `-d` 选项的 `zfs destroy` 命令会销毁快照时，才会立即销毁给定的快照。例如快照没有克隆和用户启动的引用计数为零的情况下，会发生此类立即销毁。

如果快照不符合立即销毁的资格，则将其标记为延期删除。在此状态下，此快照作为可用的可见快照存在，直到满足上述两个先决条件时将其删除。

- d
延迟快照删除。
- r
销毁后代文件系统中具有此名称的所有快照（或标记为延期删除）。

-R
以递归方式销毁所有相关项。

zfs snapshot [-r] [-o *property=value*] ... *filesystem@snapname*|*volume@snapname*
以给定的名称创建快照。通过成功的系统调用对文件系统进行的所有以前的修改都是快照的一部分。**zfs snap** 可以用作 **zfs snapshot** 的别名。有关详细信息，请参见“快照”一节。

-r
以递归方式创建所有后代数据集的快照。以原子方式创建快照，以便所有递归快照与创建快照的时刻相对应。

-o *property=value*
设置指定的属性；有关详细信息，请参见 **zfs create**。

zfs rollback [-rRf] *snapshot*
将给定的数据集回滚到以前的快照。在回滚数据集时，将放弃自创建快照以来更改的所有数据，而且数据集会返回到创建快照时的状态。缺省情况下，该命令无法回滚到除最新快照以外的快照。为了执行此操作，必须通过指定 **-r** 选项销毁所有中间快照。

-rR 选项不以递归方式销毁递归快照的子快照。这其中的任一选项仅销毁顶层递归快照。要完全回滚递归快照，必须回滚各个子快照。

-r
递归销毁比指定的快照更新的快照。

-R
递归销毁任何最新快照以及这些快照的任何克隆。

-f
与 **-R** 选项配合使用以强制卸载要销毁的任何克隆文件系统。

zfs clone [-p] [-K] [-o *property=value*] ... *snapshot filesystem*|*volume*
创建给定快照的克隆。有关详细信息，请参见“克隆”一节。目标数据集可以位于 ZFS 分层结构中的任意位置，并且创建为与原始数据集相同的数据类型。

-p
创建所有非现有的父数据集。以这种方式创建的数据集会根据从父项继承的 **mountpoint** 属性自动挂载。如果目标文件系统或卷已存在，则操作将成功完成。

-o *property=value*
设置指定的属性；有关详细信息，请参见 **zfs create**。

-K
有关信息，请参见 **zfs_encrypt(1M)**。

zfs promote clone-filesystem

将克隆文件系统提升为不再依赖于其“原始”快照。这使得可以销毁创建克隆所基于的文件系统。克隆父项-子项相关性关系将被颠倒，以便原始文件系统变为指定文件系统的克隆。

所克隆的快照和此快照之前的任何快照现在都归提升的克隆所有。这些快照使用的空间将从原始文件系统转移至提升的克隆，因此必须提供足够的空间容纳这些快照。此操作不占用任何新空间，但会调整空间记帐。提升的克隆不能有与其自身的名称冲突的任何快照名称。**rename** 子命令可用于重命名任何冲突的快照。

zfs rename filesystem|volume|snapshot

filesystem|volume|snapshot

zfs rename [-p] filesystem|volume filesystem|volume

重命名给定的数据集。新目标可以位于 ZFS 分层结构中的任意位置（快照除外）。只能在父文件系统或卷中重命名快照。重命名快照时，无需将快照的父文件系统指定为第二个参数的一部分。重命名的文件系统可以继承新挂载点，在这种情况下，将卸载这些文件系统并在新挂载点重新挂载。

-p

创建所有不存在的父数据集。以这种方式创建的数据集会根据从父项继承的 **mountpoint** 属性自动挂载。

zfs rename -r snapshot snapshot

递归重命名所有后代数据集的快照。快照是唯一可以递归重命名的数据集。

zfs list [-rH|-d depth] [-o property[,...]] [-t type[,...]] [-s property] ... [-S property] ...
[filesystem|volume|snapshot|path] ...

以表格格式列出给定数据集的属性信息。如果已指定，可以通过绝对路径名或相对路径名列出属性信息。缺省情况下，将显示所有文件系统和卷。如果 **listsnapshots** 属性为 **on**（缺省值为 **off**），将显示快照。将显示以下字段：**name**、**used**、**available**、**referenced**、**mountpoint**。

-H

用于脚本模式。不输出标题并通过单个制表符而不是任意空格分隔各个字段。

-r

在命令行上递归显示数据集的任何子项。

-d depth

递归显示数据集的任何子项，并将递归限制为最大 **depth**。深度为 **1** 时将仅显示数据集及其直接子项。

-o property

要显示的属性的逗号分隔列表。属性必须为：

- “本机属性”一节中介绍的属性之一
- 用户属性
- 值 **name**，将显示数据集名称

- 值 `space`，将显示文件系统和卷上的空间使用情况属性。这是指定
`-o name,avail,used,usedsnap,usedds,usedrefreserv,usedchild`
`-t filesystem,volume` 语法的快捷方式。

`-s property`

基于属性的值按列的升序对输出排序的属性。该属性必须为“属性”一节中介绍的属性之一，或者为按数据集名称排序的特殊值 `name`。可以使用多个 `-s` 属性选项一次指定多个属性。多个 `-s` 选项按重要性降序从左至右进行计算。

下面列出了排序条件：

- 数字类型按数字顺序排序。
- 字符串类型按字母顺序排序。
- 不适合某行的类型会将该行排序到字面值底部，而不管指定的排序方式为何。
- 如果未指定任何排序选项，将保留 `zfs list` 的现有行为。

`-S property`

与 `-s` 选项相同，但按属性的降序排序。

`-t type`

将显示逗号分隔的类型列表，其中 `type` 是 `filesystem`、`snapshot`、`volume` 或 `all` 之一。例如，指定 `-t snapshot` 时将仅显示快照。以下别名可用于替换类型说明符：`fs`（文件系统）、`snap`（快照）和 `vol`（卷）。

`zfs set [-r] property=value filesystem|volume|snapshot ...`

对每个数据集将该属性设置为给定值。只能编辑某些属性。有关可设置哪些属性和可接受的值的更多信息，请参见“属性”一节。指定数值作为确切值，或者使用后缀 `B`、`K`、`M`、`G`、`T`、`P`、`E`、`Z`（分别代表字节、KB、MB、GB、TB、PB、EB 和 ZB）以可读的格式指定数值。可以对快照设置用户属性。有关更多信息，请参见“用户属性”一节。

`-r`

在子数据集的整个子树中递归应用设置的有效值。根据属性，可以设置或继承有效值。

`zfs get [-rHp] [-d depth] [-o all | field[,...]] [-s source[,...]] all | property[,...] filesystem|volume|snapshot ...`

显示给定数据集的属性。如果未指定任何数据集，则该命令将显示系统中所有数据集的属性。对于每个属性，将显示以下列：

<code>name</code>	Dataset name
<code>property</code>	Property name
<code>value</code>	Property value
<code>source</code>	Property source. Can either be local, default, temporary, inherited, or none (-).

缺省情况下，将显示除 `RECEIVED` 列之外的所有列；使用 `-o` 选项指定特定列或所有列。此命令接受“本机属性”和“用户属性”一节中介绍的逗号分隔的属性列表。

特殊值 `all` 可用于显示应用于给定数据集的类型（文件系统、卷或快照）的所有属性。

`-r`

递归显示任何子项的属性。

`-d depth`

递归显示数据集的任何子项，并将递归限制为 *depth*。深度为 1 时将仅显示数据集及其直接子项。

`-H`

以更容易被脚本解析的格式显示输出。将省略所有标题，并通过单个制表符而不是任意数量的空格显式分隔各个字段。

`-o field`

将显示字段集。包括一个或多个：

`name,property,value,received,source`

以逗号分隔的列表形式显示多个字段。缺省值为：

`name,property,value,source`

关键字 `all` 指定所有资源。

`-s source`

将显示逗号分隔的源列表。将忽略来自此列表中的源之外的某个源的属性。每个源必须是以下源之一：

`local,default,inherited,temporary,received,none`

缺省值是所有源。

`-p`

显示可解析（确切）值中的数字。

```
zfs set share=name=sharename,path=pathname,[desc=description],prot=smb |
nfs[,property=value[,...]] filesystem
```

有关 `zfs set share` 语法的完整描述，请参见 [zfs_share\(1M\)](#)。

```
zfs inherit [-rS] property filesystem|volume|snapshot ...
```

清除指定的属性，以使从祖先继承该属性。如果没有祖先设置该属性，则将使用缺省值。有关缺省值的列表和可以继承的属性的详细信息，请参见“属性”一节。

`-r`

递归继承所有子项的给定属性。

`-S`

恢复为接收的属性值（如果有）。如果属性没有接收值，则 `zfs inherit -S` 的行为与不带 `-S` 的 `zfs inherit` 相同。如果属性有接收值，则 `zfs inherit` 会用继承的值覆盖接收的值，直到 `zfs inherit -S` 恢复为接收的值。

zfs upgrade

标识文件系统版本，这将确定当前运行的软件发行版中可用的文件系统功能。您可以继续使用旧文件系统版本，但某些功能可能不可用。文件系统可以使用 **zfs upgrade -a** 命令进行升级。无法在运行早期软件版本的系统上访问更高版本的文件系统。

zfs upgrade [-v]

显示当前软件支持的 ZFS 文件系统版本。将显示当前 ZFS 文件系统版本和所有先前支持的版本，同时显示每一版本所提供功能的说明。

zfs upgrade [-r] [-V version] [-a | *filesystem*]

将文件系统升级到新的盘上版本。如果升级文件系统，则意味着在运行旧软件版本的系统上将不再可以访问该文件系统。无法在运行旧软件版本的系统上访问从新文件系统快照生成的 **zfs send** 流。

通常，文件系统版本与池版本无关。有关 **zpool upgrade** 命令的信息，请参见 [zpool\(1M\)](#)。

在某些情况下，文件系统版本和池版本是相互关联的，在升级文件系统版本之前，必须升级池版本。

-a

升级所有已导入池中的所有文件系统。

filesystem

升级指定的文件系统。

-r

升级指定的文件系统和所有后代文件系统。

-v version

升级到指定的 *version*。如果未指定 **-v** 标志，此命令将升级到最新版本。此选项只能用于增加版本号，最高只能增加到此软件支持的最新版本。

zfs userspace [-hniHp] [-o field[,...]] [-sS field]... [-t type [,...]] *filesystem* | *snapshot*

显示指定文件系统或快照中每个用户占用的空间和每个用户的配额。这对应于 **userused@user** 和 **userquota@user** 属性。

-h

显示语法帮助消息并退出。

-n

输出数字 ID 而不是用户/组名称。

-H

不输出标题，使用制表符分隔的输出。

-p

使用确切的（可解析的）数字输出。

-o *field*[,...]

仅显示以下集中的指定字段：`type`, `name`, `used`, `quota`。缺省设置是显示所有字段。

-s *field*

按此字段对输出排序。可以多次指定 `s` 和 `S` 标志，以先按一个字段排序，然后按另一个字段排序。缺省值为 `-s type -s name`。

-S *field*

按此字段的相反顺序排序。请参见 `-s`。

-t *type*[,...]

仅输出以下集中的指定类型：`all`, `posixuser`, `smbuser`, `posixgroup`, `smbgroup`。

缺省值为 `-t posixuser, smbuser`

缺省值可以更改为包含组类型。

-i

将 SID 转换为 POSIX ID。如果不存在映射，POSIX ID 可能是临时的。一般 POSIX 接口（例如，`stat(2)`、`ls -l`）执行此转换，因此通过 `-i` 选项，可以将 `zfs userspace` 的输出直接与这些实用程序进行比较。但是，如果在建立 SMB 至 POSIX 名称映射之前，SMB 用户已创建了某些文件，`-i` 可能会引起混淆。在这种情况下，一些文件归 SMB 实体所有，一些归 POSIX 实体所有。但是，`-i` 选项将报告，POSIX 实体具有这两种实体的总计使用量和配额。

zfs groupspace [-hniHp] [-o *field*[,...]] [-sS *field*]... [-t *type* [...]] *filesystem* | *snapshot*
显示指定文件系统或快照中每个组占用的空间和每个组的配额。此子命令等同于 `zfs userspace`，但不同的是将显示的缺省类型为 `-t posixgroup, smbgroup`。

-

zfs mount

显示当前挂载的所有 ZFS 文件系统。

zfs mount [-vOS] [-o *options*] -a | *filesystem*

挂载 ZFS 文件系统。自动在引导过程中调用。

-o *options*

将临时在挂载的持续时间中使用的可选、逗号分隔的挂载选项列表。有关详细信息，请参见“临时挂载点属性”一节。

-O

执行覆盖挂载。有关更多信息，请参见 `mount(1M)`。

-v

报告挂载进度。

-a

挂载所有可用的 ZFS 文件系统。自动在引导过程中调用。

filesystem

挂载指定的文件系统。

针对已加密数据集的 `zfs mount` 操作可能会提示您输入密钥，具体取决于 `keysource` 属性值。例如，如果 `keysource` 定位器设置为 `prompt`，可能会发生这种情况。

`zfs unmount [-f] -a | filesystem|mountpoint`

卸载当前挂载的 ZFS 文件系统。自动在关闭过程中调用。

-f

强制卸载文件系统，即使当前正在使用该文件系统。

-a

卸载所有可用的 ZFS 文件系统。自动在引导过程中调用。

filesystem|*mountpoint*

卸载指定的文件系统。也可以为该命令指定系统上 ZFS 文件系统挂载点的路径。

对于已加密的数据集，卸载文件系统时，不会卸载密钥。要卸载密钥，请参见 `zfs key`。

`zfs share -a | dataset`

有关 `zfs share` 语法和示例的完整描述，请参见 [zfs_share\(1M\)](#)。

`zfs unshare -a | dataset|mountpoint`

有关 `zfs unshare` 语法和示例的完整描述，请参见 [zfs_share\(1M\)](#)。

`zfs send [-DvRbp] [-iI] snapshot snapshot`

`zfs send -r [-Dvbcp] [-i snapshot] snapshot`

创建第二个 *snapshot* 的流表示，该快照将写入标准输出。输出可以重定向到文件或不同的系统（例如，使用 [ssh\(1\)](#)）。缺省情况下，生成完整的流。

-b

仅发送接收的属性值，无论这些值是否将被本地设置覆盖，但仅当已接收到数据集时才发送。当希望 `zfs receive` 恢复在所发送数据集中备份的接收属性，并避免发送可能与源数据集没有任何关系的本地设置，但仅通过备份数据的方式时，可使用此选项。

-c

创建自包含流。自包含流不依赖于任何不包含在流软件包中的数据集。仅在与 `-r` 选项一起使用时有效。如果与 `-i` 选项一起使用，则流将依赖于指定为该选项的参数的快照。

有关详细信息，请参见《Solaris ZFS 管理指南》中的“ZFS 流”部分。

-D

对流执行 `dedup` 处理。在不支持流重复数据删除功能的系统上，无法接收重复数据删除流。

要了解复制流软件包与递归流软件包的不同之处，请参见《Solaris ZFS 管理指南》中的“ZFS 流”部分。

-i snapshot

从第一个 *snapshot* 到第二个 *snapshot* 执行增量流。增量源（第一个 *snapshot*）可以指定为快照名称的最后一部分（例如，@ 之后的部分），假设其来自于与第二个 *snapshot* 相同的文件系统。

如果目标是克隆，则源可以是原始快照，必须完整指定该快照（例如 `pool/fs@origin`，而不仅是 `@origin`）。

-I snapshot

生成流软件包，该软件包将第一个快照中的所有中间快照发送到第二个快照。例如，`-I @a fs@d` 类似于 `-i @a fs@b; -i @b fs@c; -i @c fs@d`。可以与 `-i` 选项一样指定增量源快照。

-R

生成复制流软件包，这将复制指定的文件系统、所有后代文件系统以及命名快照。接收时，所有属性、快照、后代文件系统和克隆都将被保留。

如果将 `-i` 或 `-I` 标志与 `-R` 标志结合使用，则会生成增量复制流。属性的当前值、当前快照和文件系统名称是在接收流时设置的。如果接收该流时指定了 `-F` 标志，则会销毁发送端不存在的快照和文件系统。

-r

生成递归流软件包。递归流软件包包含一系列完整流和/或增量流。接收后，将保留所有属性和后代文件系统。不同于使用 `-R` 标识生成的复制流软件包，中间快照不被保留，除非中间快照是包含在流中的克隆的原始版本。

如果 `-i` 选项与 `-r` 选项结合使用，则生成增量式递归流。属性的当前值、当前快照以及文件系统名称是在接收流时设置的。如果接收该流时指定了 `-F` 选项，则会销毁发送端不存在的快照和文件系统。`-I` 选项不能与 `-r` 选项一起使用。

与 `-c` 选项结合时，将创建自包含递归流软件包。如果同时使用 `-c` 和 `-i` 选项，则不具有使用 `-i` 选项指定的快照的文件系统和卷将作为自包含流发送。

要了解递归流软件包与复制流软件包的不同之处，请参见《Solaris ZFS 管理指南》中的“ZFS 流”部分。

-p

发送属性。

-v

输出有关所生成流软件包的详细信息。

将提交流的格式。您将接收到有关将来版本的 ZFS 的流。

```
zfs receive [-vnFu] [[-o property=value] | [-x property]] ... filesystem|volume| snapshot
zfs receive [-vnFu] [[-o property=value] | [-x property]] ... [-d | -e] filesystem
```


创建一个快照，其内容与在标准输入中提供的流中指定的内容相同。如果接收了完整的流，那么同时会创建一个新文件系统。流使用 `zfs send` 子命令创建，缺省情况下通过该命令将创建完整的流。`zfs recv` 可以用作 `zfs receive` 的别名。

如果接收增量流，则目标文件系统必须已存在，并且其最新快照必须与增量流的源匹配。对于 ZFS 卷，将销毁并重新创建目标设备链接，这意味着在 `receive` 操作期间无法访问该卷。

当接收使用 `zfs send -R` 命令生成的快照复制软件包流时，将使用 `zfs destroy -d` 命令销毁发送位置中不存在的任何快照。如果指定了 `-o property=value` 或 `-x` 属性，则其将应用于复制数据集的整个子树中属性的有效值。可以设置或继承有效属性值，具体取决于属性以及数据集是否位于复制子树中的最顶层。接收的属性尽管将被覆盖仍会保留，可以使用 `zfs inherit -rS` 或 `zfs send -Rb` 恢复。

此子命令创建的快照的名称（和文件系统，如果接收完整的流）取决于参数类型和 `-d` 或 `-e` 选项。

如果参数为快照名称，则将创建指定的 *snapshot*。如果参数为文件系统或卷名称，则会在指定的 *filesystem* 或 *volume* 中创建与发送的快照名称相同的快照。如果指定了 `-d` 或 `-e` 选项，则快照的名称通过将所发送快照的名称附加到指定的文件系统来确定。如果指定了 `-d` 选项，则将附加除所发送快照路径的池名称之外的所有内容（例如，附加的 `b/c@1` 来自所发送的快照 `a/b/c@1`），如果指定了 `-e` 选项，则仅附加所发送快照路径的尾部（例如，附加的 `c@1` 来自所发送的快照 `a/b/c@1`）。在使用 `-d` 的情况下，将在指定的文件系统中创建复制所发送快照的路径需要的所有文件系统。

-d

使用除所发送快照路径的第一个元素之外的任何元素（除池名称之外的所有内容）确定上述段落中介绍的新快照的名称。

-e

使用所发送快照路径的最后一个元素确定上述段落中介绍的新快照的名称。

-F

执行接收操作之前，将文件系统强制回滚到最新快照。如果接收增量复制流（例如，`zfs send -R -[iI]` 生成的流），则将销毁发送端不存在的快照和文件系统。

-n

不实际接收该流。此选项可与 `-v` 选项结合使用来验证接收操作将使用的名称。

-o property=value

设置指定的属性，就好像在调用命令 `zfs set property=value` 的同时，从非增量 `send` 流创建了接收的数据集，或从增量 `send` 流更新了接收的数据集。也可以在接收时设置任何可编辑的 ZFS 属性。绑定到接收数据的一次性设置属性（例如 `normalization` 和 `casesensitivity`）无法在接收时编辑，甚至当通过 `zfs receive` 新创建数据集时也是如此。可以指定多个 `-o` 选项。如果在多个 `-o` 或 `-x` 选项中指定相同的属性，则会发生错误。

-u
未挂载与接收流关联的文件系统。

-v
输出有关执行接收操作所需的流和时间的详细信息。

-x *property*
确保 `receive` 后的指定属性的有效值不受 `send` 流（如果有）中该属性的值的影响，就好像属性已从 `send` 流中排除一样。如果指定的属性在 `send` 流中不存在，则此选项不起作用。如果需要覆盖接收的属性，则可以设置或继承有效值，具体取决于属性。在使用增量更新的情况下，`-x` 会保持任何现有的本地设置或显式继承不变（因为已覆盖接收属性）。所有 `-o` 限制同样适用于 `-x`。

`zfs allow filesystem | volume`

`zfs allow [-ldug] everyone|user|group[,...] perm |@setname[,...] filesystem| volume`

`zfs allow [-ld] -e perm|@setname[,...] filesystem| volume`

`zfs allow -c perm|@setname[,...] filesystem|volume`

`zfs allow -s @setname perm|@setname[,...] filesystem|volume`

有关 `zfs allow` 语法和示例的完整描述，请参见 [zfs_allow\(1M\)](#)。

`zfs unallow [-rldug] everyone|user|group[,...] [perm|@setname[, ...]] filesystem| volume`

`zfs unallow [-rld] -e [perm|@setname [,...]] filesystem|volume`

`zfs unallow [- r] -c [perm|@setname[,...]]`

`filesystem|volume`

`zfs unallow [-r] -s @setname [ perm|@setname[,...]]`

`filesystem|volume`

有关 `zfs unallow` 语法和示例的完整描述，请参见 [zfs_allow\(1M\)](#)。

`zfs hold [-r] tag snapshot...`

向指定快照中添加通过 `tag` 参数命名的单个引用。每个快照都有其自己的标志名称空间，标志在该空间内必须是唯一的。

如果一个快照上存在一个保持，尝试使用 `zfs destroy` 命令销毁该快照将返回 `EBUSY`。

-r
指定将具有特定标志的保持递归应用于所有后代文件系统的快照。

`zfs holds [-r] snapshot...`

列出指定快照的所有现有用户引用。

-r
除了列出命名快照上的保持外，还列出在命名后代快照上设置的保持。

`zfs release [-r] tag snapshot...`

从指定快照中删除通过 `tag` 参数命名的单个引用。每个快照的标志必须已经存在。

如果一个快照上存在一个保持，尝试使用 `zfs destroy` 命令销毁该快照将返回 `EBUSY`。

-r

在所有后代文件系统的快照上递归释放具有特定标志的保持。

```
zfs diff [-FHt] [-o field] ... snapshot [snapshot | filesystem]
```

```
zfs diff -E [- FHt] [-o field] ... snapshot | filesystem
```

提供快照与后代数据集之间的区别的高级说明。后代可以是数据集的快照或当前数据集。

如果指定单个快照，将提供该快照和当前数据集之间的区别。

对于已执行原始快照与后代之间的更改的每个文件，将与文件名一起说明更改的类型。如果重命名，旧名称和新名称都显示。对于在文件名中找到的空格字符、反斜杠字符及其他无法输出或非 7 位的 ASCII 字符，都显示为三位八进制表示的字节值前面的反斜杠字符。

如果已指定 -t 选项，该命令的第一列输出是文件的 st_ctim 值。对于已删除的文件，这是以前快照中的最终 st_ctim。

更改类型遵循显示的任何时间戳并通过单字符说明：

- + 指明文件是在后来的数据集内添加的。
- 指明文件是在后来的数据集内删除的。
- M 指明文件是在后来的数据集内修改的。
- R 指明文件是在后来的数据集内重命名的。

如果已指定 -F 选项，下一列输出将是说明文件类型的单字符。这些映射为：

- F 正规文件
- / 目录
- B 块设备
- > 门
- | FIFO
- @ 符号链接
- P 事件门户
- = 套接字

如果修改涉及非目录文件的链接计数中的更改，该更改将表示为修改行中的圆扩号内的增量。如果重命名该文件，则会使用字符串 "->" 将旧名称与新名称隔开。

如果选择 -H 选项，将生成更易于解析的输出。字段用单个制表符隔开，重命名的旧名称与新名称之间不放置箭头字符串 (->)。不能保证非 -H 输出的字段之间存在间距。

如果选中 **-e** 选项，则将逐一列出两个快照之间添加或修改的所有文件，不会显示删除的文件。无论修改是何种类型，更改类型将始终报告为 "+"。

如果选中 **-E** 选项，则提供的区别与空的快照和指定快照或数据集之间的区别一样。

如果选中 **-o** 字段选项，则只显示选中字段。每行以通过 **-F** 和 **-t** 选项请求的标准字段开始，后跟后续 **-o** 选项中请求的字段。与 **-H** 选项一样，所有字段都以单个制表符分隔。允许的字段名称包括：

object	ls-i 输出的文件编号
parent	ls-i 输出的文件所在目录的编号
size	ls-s 显示的文件大小
links	指向文件的链接数
linkschange	指向文件的链接数的更改
name	文件名称
oldname	重命名之前文件的名称，或者 -（连字符）（如果文件未重命名）
user	ls 显示的文件属主的用户名
group	ls 显示的文件组名
ctime	上一次修改文件元数据的时间戳
mtime	上一次修改文件的时间戳
atime	上一次访问文件的时间戳
crttime	创建文件的时间戳

除非用户已具有 {PRIV_SYS_CONFIG} 或 {PRIV_SYS_MOUNT} 特权，否则必须通过 **zfs allow** 向他们授予 **diff** 权限才能使用该子命令。

```
zfs key -l | {-a | [-r] filesystem|volume}
zfs key -u [-f] | {-a | [-r] filesystem|volume}
zfs key -c [-o keysource=value] {-a | [-r] filesystem| volume}
zfs key -K {-a | [-r] filesystem|volume}
有关 zfs key 语法和示例的完整描述，请参见 zfs\_encrypt\(1M\)。
```

示例

示例 1 创建 ZFS 文件系统分层结构

以下命令创建名为 **pool/home** 的文件系统和名为 **pool/home/bob** 的文件系统。对父文件系统设置挂载点 **/export/home**，子文件系统将自动继承该挂载点。

```
# zfs create pool/home
# zfs set mountpoint=/export/home pool/home
# zfs create pool/home/bob
```

示例2 创建 ZFS 快照

以下命令创建名为 `yesterday` 的快照。该快照按要求挂载在 `pool/home/bob` 文件系统根目录中的 `.zfs/snapshot` 目录下。

```
# zfs snapshot pool/home/bob@yesterday
```

示例3 创建和销毁多个快照

以下命令为 `pool/home` 及其所有后代文件系统创建名为 `yesterday` 的快照。每个快照都按要求挂载在其文件系统根目录中的 `.zfs/snapshot` 目录下。第二个命令销毁新建的快照。

```
# zfs snapshot -r pool/home@yesterday
# zfs destroy -r pool/home@yesterday
```

示例4 禁用和启用文件系统压缩

以下命令禁用 `pool/home` 下的所有文件系统的 `compression` 属性。下一个命令显式启用 `pool/home/anne` 的 `compression`。

```
# zfs set compression=off pool/home
# zfs set compression=on pool/home/anne
```

示例5 列出 ZFS 数据集

以下命令列出所有的活动文件系统和系统中的卷。如果 `listsnapshots` 属性为 `on`，将显示快照。缺省值为 `off`。有关池属性的更多信息，请参见 [zpool\(1M\)](#)。

```
# zfs list
```

NAME	USED	AVAIL	REFER	MOUNTPOINT
pool	450K	457G	18K	/pool
pool/home	315K	457G	21K	/export/home
pool/home/anne	18K	457G	18K	/export/home/anne
pool/home/bob	276K	457G	276K	/export/home/bob

示例6 在 ZFS 文件系统上设置配额

以下命令将 `pool/home/bob` 的配额设置为 30 GB。

```
# zfs set quota=30G pool/home/bob
```

示例7 列出 ZFS 属性

以下命令列出 `pool/home/bob` 的所有属性。

```
# zfs get all pool/home/bob
```

NAME	PROPERTY	VALUE	SOURCE
pool/home/bob	type	filesystem	-
pool/home/bob	creation	Fri Jul 15 11:53 2011	-
pool/home/bob	used	31K	-
pool/home/bob	available	30.0G	-
pool/home/bob	referenced	31K	-

示例7 列出 ZFS 属性 (续)

pool/home/bob	compressratio	1.00x	-
pool/home/bob	mounted	yes	-
pool/home/bob	quota	30G	local
pool/home/bob	reservation	none	default
pool/home/bob	recordsize	128K	default
pool/home/bob	mountpoint	/pool/home/bob	default
pool/home/bob	sharenfs	off	default
pool/home/bob	checksum	on	default
pool/home/bob	compression	on	default
pool/home/bob	atime	on	default
pool/home/bob	devices	on	default
pool/home/bob	exec	on	default
pool/home/bob	setuid	on	default
pool/home/bob	readonly	off	default
pool/home/bob	zoned	off	default
pool/home/bob	snapdir	hidden	default
pool/home/bob	aclmode	discard	default
pool/home/bob	aclinherit	restricted	default
pool/home/bob	canmount	on	default
pool/home/bob	xattr	on	default
pool/home/bob	copies	1	default
pool/home/bob	version	5	-
pool/home/bob	utf8only	off	-
pool/home/bob	normalization	none	-
pool/home/bob	casesensitivity	mixed	-
pool/home/bob	vscan	off	default
pool/home/bob	nbmand	off	default
pool/home/bob	sharesmb	off	default
pool/home/bob	refquota	none	default
pool/home/bob	refreservation	none	default
pool/home/bob	primarycache	all	default
pool/home/bob	secondarycache	all	default
pool/home/bob	usedbysnapshots	0	-
pool/home/bob	usedbydataset	31K	-
pool/home/bob	usedbychildren	0	-
pool/home/bob	usedbyrefreservation	0	-
pool/home/bob	logbias	latency	default
pool/home/bob	dedup	off	default
pool/home/bob	mlslabel	none	default
pool/home/bob	sync	standard	default
pool/home/bob	encryption	off	-
pool/home/bob	keysource	none	default
pool/home/bob	keystatus	none	-
pool/home/bob	rekeydate	-	default
pool/home/bob	rstchown	on	default
pool/home/bob	shadow	none	-

示例7 列出 ZFS 属性 (续)

以下命令获取单个属性值。

```
# zfs get -H -o value compression pool/home/bob
on
```

以下命令列出具有 pool/home/bob 的本地设置的所有属性。

```
# zfs get -r -s local -o name,property,value all pool/home/bob
```

NAME	PROPERTY	VALUE
pool/home/bob	quota	30G
pool/home/bob	compression	on

示例8 回滚 ZFS 文件系统

以下命令将 pool/home/anne 的内容恢复为名为 yesterday 的快照，并删除所有中间快照。

```
# zfs rollback -r pool/home/anne@yesterday
```

示例9 创建 ZFS 克隆

以下命令创建初始内容与 pool/home/bob@yesterday 相同的可写文件系统。

```
# zfs clone pool/home/bob@yesterday pool/clone
```

示例10 提升 ZFS 克隆

以下命令说明如何使用克隆、克隆提升及重命名检验文件系统的更改，然后将原始文件系统替换为更改的文件系统。

```
# zfs create pool/project/production
  populate /pool/project/production with data
# zfs snapshot pool/project/production@today
# zfs clone pool/project/production@today pool/project/beta
make changes to /pool/project/beta and test them
# zfs promote pool/project/beta
# zfs rename pool/project/production pool/project/legacy
# zfs rename pool/project/beta pool/project/production
once the legacy version is no longer needed, it can be destroyed
# zfs destroy pool/project/legacy
```

示例11 继承 ZFS 属性

以下命令使 pool/home/bob 和 pool/home/anne 继承其父项的 checksum 属性。

```
# zfs inherit checksum pool/home/bob pool/home/anne
```

示例 12 远程复制 ZFS 数据

以下命令将完整流和增量流依次发送到远程计算机中，以将它们分别恢复到 poolB/received/fs@a 和 poolB/received/fs@b 中。poolB 必须包含文件系统 poolB/received，最初不能包含 poolB/received/fs。

```
# zfs send pool/fs@a | \
    ssh host zfs receive poolB/received/fs@a
# zfs send -i a pool/fs@b | ssh host \
    zfs receive poolB/received/fs
```

示例 13 使用 zfs receive -d 选项

以下命令将 poolA/fsA/fsB@snap 的完整流发送到远程计算机中，以将其接收到 poolB/received/fsA/fsB@snap 中。已接收快照的名称中的 fsA/fsB@snap 部分是基于发送快照的名称决定的。poolB 必须包含文件系统 poolB/received。如果 poolB/received/fsA 不存在，将其作为空文件系统创建。

```
# zfs send poolA/fsA/fsB@snap | \
    ssh host zfs receive -d poolB/received
```

示例 14 设置用户属性

以下示例可以为数据集设置用户定义的 com.example:department 属性。

```
# zfs set com.example:department=12345 tank/accounting
```

示例 15 执行滚动快照

以下示例显示如何通过一致的命名方案维护快照的历史信息。为保留一周的快照，用户销毁最早的快照，重命名剩余的快照，然后创建新快照，如下所示：

```
# zfs destroy -r pool/users@7daysago
# zfs rename -r pool/users@6daysago @7daysago
# zfs rename -r pool/users@5daysago @6daysago
# zfs rename -r pool/users@4daysago @5daysago
# zfs rename -r pool/users@3daysago @4daysago
# zfs rename -r pool/users@2daysago @3daysago
# zfs rename -r pool/users@yesterday @2daysago
# zfs rename -r pool/users@today @yesterday
# zfs snapshot -r pool/users@today
```

示例 16 显示 ZFS 快照的差别

以下示例是指定的 zfs diff -F 和 -t 选项的输出：

```
# zfs diff -Ft myfiles@snap1
1269962501.206726811  M  /  /myfiles/
1269962444.207369955  M  F  /myfiles/link_to_me      (+1)
1269962499.207519034  R      /myfiles/rename_me -> /myfiles/renamed
1269962431.813566720  -  F  /myfiles/delete_me
1269962518.666905544  +  F  /myfiles/new_file
```


示例 16 显示 ZFS 快照的差别 (续)

```
1269962501.393099817 + | /myfiles/new_pipe
```

退出状态 将返回以下退出值：

- 0 成功完成。
- 1 出现错误。
- 2 指定的命令行选项无效。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/zfs
接口稳定性	Committed（已确定）

另请参见 [chmod\(1\)](#)、[chown\(1\)](#)、[pktool\(1\)](#)、[ssh\(1\)](#)、[mount\(1M\)](#)、[shadowd\(1M\)](#)、[share\(1M\)](#)、[share_nfs\(1M\)](#)

请参见 [gzip\(1\)](#) 手册页，该页不属于 SunOS 手册页集合。

有关其他 ZFS 功能的信息，请参见 [zfs_allow\(1M\)](#)、[zfs_encrypt\(1M\)](#)、[zfs_share\(1M\)](#) 和《Oracle Solaris 管理：ZFS 文件系统》。

附注 如通过 `diff` 子命令修改中所述的文件可能已通过多种方式进行修改。使 `st_ctim`（请参见 [stat\(2\)](#)）中发生更改的任何操作都是报告修改的基础。

引用名

zfs_allow – 将 ZFS 文件系统管理权限委托给非特权用户

用法概要

```
zfs allow filesystem|volume

zfs allow [-ldug] everyone|user|group[,...] perm|@setname[,...]
    filesystem|volume

zfs allow [-ld] -e perm|@setname[,...] filesystem|volume

zfs allow -c perm|@setname[,...] filesystem|volume

zfs allow -s @setname perm|@setname[,...] filesystem|volume

zfs unallow [-rldug] everyone|user|group[,...] [perm|@setname[,... ]]
    filesystem|volume

zfs unallow [-rld] -e [perm|@setname[,... ]] filesystem|volume

zfs unallow [-r] -c [perm|@setname[ ... ]] filesystem|volume

zfs unallow [-r] -s @setname [perm|@setname[,... ]] filesystem|volume
```

描述

zfs allow 命令可用于将权限委托给非特权用户，以便管理 ZFS 存储池中的 ZFS 文件系统，如 [zpool\(1M\)](#) 中所述。可以使用 zfs unallow 命令撤消管理权限。

权限通常是指使用 ZFS 子命令或更改 ZFS 属性的能力。以下列出了可用的权限：

NAME	TYPE	NOTES
allow	subcommand	Must also have the permission that is being allowed
clone	subcommand	Must also have the 'create' ability and 'mount' ability in the origin file system
create	subcommand	Must also have the 'mount' ability
destroy	subcommand	Must also have the 'mount' ability
diff	subcommand	Allows user to compare differences between snapshots of a ZFS dataset
hold	subcommand	Allows adding a user hold to a snapshot
key	subcommand	Allows key loading or unloading
keychange	other	Allows wrapping key change and re-key
mount	subcommand	Allows mount/umount of ZFS datasets
promote	subcommand	Must also have the 'mount' and 'promote' ability in the origin file system
receive	subcommand	Must also have the 'mount' and 'create' ability
release	subcommand	Allows releasing a user hold which might destroy the snapshot
rename	subcommand	Must also have the 'mount' and 'create' ability in the new parent
rollback	subcommand	
send	subcommand	
share	subcommand	Allows sharing file systems over NFS or SMB protocols
snapshot	subcommand	
groupquota	other	Allows accessing any groupquota@...

		property
groupused	other	Allows reading any groupused@... property
userprop	other	Allows changing any user property
userquota	other	Allows accessing any userquota@... property
userused	other	Allows reading any userused@... property
aclinherit	property	
aclmode	property	
atime	property	
canmount	property	
casesensitivity	property	
checksum	property	
compression	property	
copies	property	
dedup	property	
devices	property	
encryption	property	
exec	property	
keysource	property	
logbias	property	
mlslabel	property	
mountpoint	property	
nbmand	property	
normalization	property	
primarycache	property	
quota	property	
readonly	property	
recordsize	property	
refquota	property	
refreservation	property	
reservation	property	
rstchown	property	
secondarycache	property	
setuid	property	
sharenfs	property	
sharesmb	property	
snapdir	property	
sync	property	
utf8only	property	
version	property	
volblocksize	property	
volsize	property	
vscan	property	
xattr	property	
zoned	property	

子命令 修改状态的所有子命令将以原始形式永久记录到池中。

zfs ?
显示帮助消息。

zfs help *command* | help | *property property-name* | *permission*

显示 **zfs** 命令的用法信息。可以显示特定命令、属性或委托权限的帮助。如果显示特定命令或属性的帮助，将显示命令语法或属性值。在不带任何参数的情况下使用 **zfs help** 将显示完整的 **zfs** 命令列表。

zfs allow *filesystem* | *volume*

显示已在指定的文件系统或卷中委托的权限。有关更多信息，请参见 **zfs allow** 的其他形式。

zfs allow [-ldug] *everyone*|*user*|*group*[,...] *perm*|*@setname*[,...] *filesystem* | *volume*

zfs allow [-ld] -e *perm*|*@setname*[,...] *filesystem* | *volume*

将文件系统的 ZFS 管理权限委托给非特权用户。

[-ug] *everyone*|*user*|*group*[,...]

指定权限将要委托给的用户。可以按逗号分隔的列表形式指定多个实体。如果未指定两个 **-ug** 选项中的任何一个，则优先将该参数解释为关键字

everyone，然后解释为用户名，最后解释为组名。要指定名为 **"everyone"** 的用户或组，请使用 **-u** 或 **-g** 选项。要将同名的组指定为用户，请使用 **-g** 选项。

[-e] *perm*|*@setname*[,...]

指定权限将委托给 **everyone**。可以按逗号分隔的列表形式指定多个权限。权限名称与 ZFS 子命令和属性的名称相同。请参见下面的属性列表。可以指定以 **at** 符号 (**@**) 开头的属性集名称。有关详细信息，请参见下面的 **-s** 格式。

[-ld] *filesystem*|*volume*

指定将委托权限的位置。如果 **-ld** 两个选项均未指定或同时指定了这两个选项，则权限允许用于文件系统或卷及其所有后代。如果仅使用 **-l** 选项，则仅在本机允许用于指定的文件系统。如果仅使用 **-d** 选项，则仅允许用于后代文件系统。

zfs allow -c *perm*|*@setname*[,...] *filesystem*|*volume*

设置“创建时间”权限。这些权限在本地授予给任何新创建的后代文件系统的创建者。

zfs allow -s *@setname perm*|*@setname*[,...] *filesystem*|*volume*

定义权限或向权限集内添加权限。其他 **zfs allow** 命令可将权限集用于指定的文件系统及其后代。可以对权限集进行动态评估，因此对权限集的更改可立即反映。权限集与 ZFS 文件系统遵循相同的命名限制，但名称必须以 **"at 符号" (@)** 开头，且长度不能超过 64 个字符。

zfs unallow [-rldug] *everyone*|*user*|*group*[,...] [*perm*|*@setname*[,...]] *filesystem* | *volume*

zfs unallow [-rld] -e [*perm*|*@setname* [...]] *filesystem*|*volume*

zfs unallow [-r] -c [*perm*|*@setname*[,...]]

filesystem*|*volume

删除通过 **zfs allow** 命令授予的权限。未显式拒绝任何权限，因此授予的其他权限仍有效。例如，如果权限是由祖先授予的。如果未指定任何权限，指定

user、**group** 或 **everyone** 的所有权限均被删除。指定 **everyone**（或使用 **-e** 选项）仅删除授予 **everyone** 的权限，而不是每个用户和组的所有权限。有关 **-ldug** 选项的说明，请参见 **zfs allow** 命令。

-r

从该文件系统和所有后代中递归删除权限。

```
zfs unallow [-r] -s @setname [perm|@setname[,...]]
filesystem|volume
```

从权限集中删除权限。如果未指定权限，将删除所有权限，因此会删除整个权限集。

示例

示例1 委托对 ZFS 数据集的 ZFS 管理权限

以下示例说明如何设置权限，以使用户 `anne` 可以创建、销毁和挂载 `pool/home/anne`，并可以捕获 `pool/home/anne` 的快照。还会显示对 `pool/home/anne` 的权限。

```
# zfs allow anne create,destroy,mount,snapshot pool/home/anne
# zfs allow pool/home/anne
---- Permissions on pool/home/anne -----
Local+Descendent permissions:
    user anne create,destroy,mount,snapshot
```

因为缺省情况下 `pool/home/anne` 挂载点权限设置为 755，所以用户 `anne` 将无法在 `pool/home/anne` 下挂载文件系统。设置类似以下语法的 ACL 可提供挂载点访问权限：

```
# chmod A+user:anne:add_subdirectory:allow /pool/home/anne
```

示例2 委托对 ZFS 数据集的创建时间权限

以下示例显示如何授予 `staff` 组中的每个人在 `pool/home` 中创建文件系统的权限。该语法还允许员工成员销毁他们自己的文件系统，但不允许他们销毁其他人的文件系统。还会显示对 `pool/home` 的权限。

```
# zfs allow staff create,mount pool/home
# zfs allow -c destroy pool/home
# zfs allow pool/home
---- Permissions on pool/home -----
Create time permissions:
    destroy
Local+Descendent permissions:
    group staff create,mount
```

示例3 定义和授予对 ZFS 数据集的权限集

以下示例显示如何定义和授予对 `pool/home` 文件系统的权限集。还会显示对 `pool/home` 的权限。

```
# zfs allow -s @pset create,destroy,snapshot,mount pool/home
# zfs allow staff @pset pool/home
# zfs allow pool/home
---- Permissions on pool/home -----
Permission sets:
    @pset create,destroy,mount,snapshot
```

示例3 定义和授予对 ZFS 数据集的权限集 (续)

```
Create time permissions:
    destroy
Local+Descendent permissions:
    group staff @pset,create,mount
```

示例4 委托对 ZFS 数据集的属性权限

以下示例显示如何授予对 tank/users 文件系统的设置配额和预留空间的权限。还会显示对 tank/users 的权限。

```
# zfs allow mark quota,reservation tank/users
# zfs allow tank/users
---- Permissions on tank/users -----
Local+Descendent permissions:
    user mark quota,reservation
mark% zfs set quota=10G tank/users/tim
mark% zfs get quota tank/users/tim
NAME                PROPERTY  VALUE  SOURCE
tank/users/tim      quota     10G    local
```

示例5 删除对 ZFS 数据集的 ZFS 委托权限

以下示例显示如何从 @pset 权限集中删除 staff 组对 pool/home 文件系统的快照权限。还会显示对 pool/home 的权限。

```
# zfs unallow -s @pset snapshot pool/home
# zfs allow pool/home
---- Permissions on pool/home -----
Permission sets:
    @pset create,destroy,mount
Create time permissions:
    destroy
Local+Descendent permissions:
    group staff @pset,create,mount
```

退出状态

将返回以下退出值：

- 0
成功完成。
- 1
出现错误。
- 2
指定的命令行选项无效。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/zfs
接口稳定性	Committed（已确定）

另请参见 [zfs\(1M\)](#)、[zpool\(1M\)](#)、[chmod\(2\)](#)、[chown\(2\)](#)、[attributes\(5\)](#)

有关使用其他 ZFS 功能的信息，请参见 [zfs_encrypt\(1M\)](#)、[zfs_share\(1M\)](#)、[zfs\(1M\)](#) 和 《[Oracle Solaris 管理：ZFS 文件系统](#)》。

引用名	zfs_encrypt – 对 ZFS 文件系统进行加密
用法概要	<pre>zfs [-?] zfs help subcommand help property property-name permission zfs create -o encryption=on [-o keysource=raw hex passphrase,prompt file:// pkcs11: https://] ... zfs clone [-p] [-K] [-o property=value] ... snapshot filesystem volume zfs get [-r -d depth][-Hp][-o all field[,...]] [-s source[,...]] all property[,...] filesystem volume snapshot ... zfs key-l {-a [-r] filesystem volume} zfs key -u [-f] {-a [-r] filesystem volume} zfs key -c [-o keysource=value] {-a [-r] filesystem volume} zfs key -K {-a [-r] filesystem volume} zfs mount zfs mount [-v0] [-o options] -a filesystem zfs unmount [-f] -a filesystem mountpoint</pre>
描述	<code>zfs create -o encryption</code> 命令可对 ZFS 存储池中新创建的 ZFS 数据集进行加密，如 zpool(1M) 中所述。
加密	加密是对数据进行编码以实现保密性的过程，数据所有者需要使用密钥才能访问编码的数据。您可以在创建 ZFS 数据集时设置加密策略，但是该策略不可进行更改。有关详细信息，请参见“本机属性”中的 <code>encryption</code> 和 <code>keysource</code> 属性说明。 数据集加密将被永久继承，在数据集克隆期间无法删除。接收重复的数据集流时，如果希望加密，则目标数据集必须启用加密。否则，数据将以明文形式存储。完全复制加密的数据集流可以使数据集得以加密，但会采用新生成的密钥，并且流本身不会被加密。
本机 ZFS 加密属性	以下与 ZFS 加密相关的本机属性包含有关数据集的只读统计信息。无法设置或继承这些属性。除非另有说明，否则本机属性适用于所有数据集类型。有关 ZFS 本机属性的完整描述和列表，请参见 zfs(1M)。 keystatus 标识数据集的加密密钥状态。通过显示 <code>available</code> 或 <code>unavailable</code> 的状态，可指明数据集密钥的可用性。对于没有启用加密的数据集，则显示 <code>none</code>。 mounted 对于文件系统，指示当前是否已挂载文件系统。此属性可以是 <code>yes</code> 或 <code>no</code>。 rekeydate 上次因对此数据集执行 <code>zfs key -K</code> 或 <code>zfs clone -K</code> 操作而导致数据加密密钥发生更改的日期。如果未执行密钥更新操作，则 <code>rekeydate</code> 与 <code>creation</code> 日期相同。

在创建文件系统后无法更改以下属性，因此在创建文件系统时应设置这些属性。如果未使用 `zfs create` 或 `zpool create` 命令设置属性，则从父数据集继承这些属性。如果在这些功能获得支持以前已创建的父数据集缺少这些属性，则新文件系统将为这些属性采用缺省值。

`encryption=off | on | aes-128-ccm | aes-192-ccm | aes-256-ccm | aes-128-gcm | aes-192-gcm | aes-256-gcm`

定义用于加密数据集的加密算法和密钥长度。on 值等于 `aes-128-ccm`。缺省值为 `off`。将加密设置为 `off` 以外的值时，`checksum` 属性被设置为 `sha256+mac` 并且变为 `readonly`。

在创建时必须指定以下属性，并可以使用特殊命令修改这些属性：

`keysource=raw | hex | passphrase,prompt | file:/// | pkcs11: | https://`

定义包含数据集密钥的密钥格式和位置。使用 `zfs key -l` 命令创建、挂载或装入数据集时，必须提供密钥。

`keysource` 属性接受两个值：`format`，确定提供密钥的方式；`locator`，标识密钥的来源位置。

`format` 接受三个值：

- `raw`：原始的密钥字节
- `hex`：十六进制密钥字符串
- `passphrase`：生成密钥的字符串

`locator` 接受两个值：

- `prompt`：创建或挂载数据集时提示您输入密钥
- `file:///filename`：文件系统中的密钥文件位置
- `pkcs11`：描述 PKCS#11 标记中密钥位置的 URI
- `https://location`：安全的服务器上的密钥文件位置

有关使用 `https://` 定位器创建密钥的示例，请参见“示例”。

要更改包装密钥值或密钥，必须运行 `zfs key -c` 命令。如果只需要更改密钥位置（例如，文件名更改），则使用 `zfs set` 命令及 `keysource` 属性。请注意，当只使用 `zfs set` 命令更改密钥位置时，ZFS 不会执行检查，如新位置是否具有有效的包装密钥。

如果未指定和未继承 `keysource`，则对于已启用加密的数据集，缺省的 `keysource` 将设置为 `passphrase,prompt`；对于已禁用加密的数据集，该属性将设置为 `none`。

子命令

修改状态的所有子命令将以原始形式永久记录到池中。

`zfs ?`

显示帮助消息。

zfs help *command* | help | *property property-name* | *permission*

显示 **zfs** 命令的用法信息。可以显示特定命令、属性或委托权限的帮助。如果显示特定命令或属性的帮助，将显示命令语法或属性值。在不带任何参数的情况下使用 **zfs help** 将显示完整的 **zfs** 命令列表。

zfs create [-p] [-o *encryption=on*] [-o *keysource=raw* | *hex* | *passphrase, prompt* | *file://|pkcs11:|https://*] ... *filesystem*

创建新的 ZFS 文件系统并启用加密，它使用 **aes-128-ccm**。有关支持的加密算法的列表，请参见加密属性描述。

-p

创建所有非现有的父数据集。以这种方式创建的数据集会根据从父项继承的 **mountpoint** 属性自动挂载。在命令行上使用 **-o** 选项指定的任何属性将被忽略。如果目标文件系统已经存在，则会成功完成此操作。

-o *encryption=value*

将加密属性设置为 *value*。可以指定多个 **-o** 选项。如果在多个 **-o** 选项中指定了相同的属性，则会导致错误。

zfs clone [-p] [-K] [-o *property=value*] ... *snapshot filesystem* | *volume*

创建给定快照的克隆。有关详细信息，请参见“克隆”一节。目标数据集可以位于 ZFS 分层结构中的任意位置，并且创建为与原始数据集相同的数据类型。

-p

创建所有非现有的父数据集。以这种方式创建的数据集会根据从父项继承的 **mountpoint** 属性自动挂载。如果目标文件系统或卷已存在，则操作将成功完成。

-o *property=value*

设置指定的属性；有关详细信息，请参见 **zfs create**。

-K

在此数据集的钥匙串中创建一个新的数据加密密钥。写入克隆中的数据将使用新的数据加密密钥，该密钥与其原始快照不同。

zfs set *keysource=value filesystem* | *volume* | ...

对每个数据集将 **keysource** 属性设置为给定值。只能更改 **keysource** 位置。如果要更改包装密钥值，请使用 **zfs key -c** 命令。

-r

在子数据集的整个子树中递归应用设置的有效值。根据属性，可以设置或继承有效值。

zfs get *encryption* | *keysource* | *keystatus* | *rekeydate filesystem* | *volume* | ...

显示给定数据集的属性。

-r

递归显示任何子项的属性。

-d *depth*

递归显示数据集的任何子项，并将递归限制为 *depth*。深度为 1 时将仅显示数据集及其直接子项。

-H

以更容易被脚本解析的格式显示输出。将省略所有标题，并通过单个制表符而不是任意数量的空格显式分隔各个字段。

zfs key-l [{-a} | [-r] *filesystem* [*volume*]

装入一个数据集以及继承该密钥的任何数据集的加密密钥。通过该命令提供的密钥不是用于加密数据集的实际密钥。它是数据集的数据加密密钥集的包装密钥。

-l

装入包装密钥以取消加密数据集和继承该密钥的数据集的锁定。该命令基于数据集的 *keysource* 属性定义的内容装入密钥。

在池导入期间，在挂载数据集时执行密钥装入操作。在引导期间，如果包装密钥可用并且 *keysource* 未设置为 *prompt*，将执行密钥装入操作。

-a

应用于系统上所有池中的所有数据集。

-r

将该操作递归应用于命名文件系统或卷下的所有数据集。

zfs key-u [-f] [{-a} | [-r] *filesystem* [*volume*]

卸载一个数据集以及继承该密钥的任何数据集的加密密钥。

-u

先取消挂载数据集，然后尝试卸载加密数据集以及继承密钥的数据集的包装密钥。如果卸载成功，数据集将无法访问并且已取消挂载。

-f

在尝试卸载密钥前尝试强制取消挂载数据集。如果未指定，将尝试正常取消挂载。

-a

应用于系统上所有池中的所有数据集。

-r

将该操作递归应用于命名文件系统或卷下的所有数据集。

zfs key-c [-o *keysource*=*value*] [{-a} | [-r] *filesystem* [*volume*]

更改包装密钥。如果新密钥具有不同的格式或定位器，*keysource* 属性必须包含在该命令中。仅 **zfs key -c** 命令中包含的 *keysource* 属性可以更改。

-c

更改加密数据集和继承包装密钥的数据集的包装密钥。现有密钥必须装入之后，密钥更改操作才能发生。ZFS 不会提示您输入现有口令短语。

-o *property=value*

密钥更改操作中包含的要更改的属性。**keysource** 属性是唯一一个可以在密钥更改操作中更改的选项。

您必须具有更改 **keysource** 属性的权限。

-a

应用于系统上所有池中的所有数据集。

-r

将该操作递归应用于命名文件系统或卷下的所有数据集。

zfs key -K { -a | [-r] *filesystem* | *volume* }

创建新的数据加密密钥。新的数据加密密钥由与该数据集的任何现有数据加密密钥相同的包装密钥包装。

-K

为该数据集创建新的数据加密密钥。在该操作之后写入的数据将使用新的数据加密密钥。

-a

应用于系统上所有池中的所有数据集。

-r

将该操作递归应用于命名文件系统或卷下的所有数据集。

zfs mount

zfs mount [-v0] [-o *options*] -a | *filesystem*

挂载 ZFS 文件系统。自动在引导过程中调用。有关 **zfs mount** 语法的完整描述，请参见 [zfs\(1M\)](#)。

filesystem

挂载指定的文件系统。

已加密数据集的 **zfs mount** 操作可能会提示您输入密钥，具体取决于 **keysource** 属性值。例如，如果 **keysource** 定位器设置为 **prompt**，可能会发生这种情况。

zfs unmount [-f] -a | *filesystem* | *mountpoint*

卸载当前挂载的 ZFS 文件系统。自动在关闭过程中调用。有关 **zfs unmount** 语法的完整描述，请参见 [zfs\(1M\)](#)。

filesystem | *mountpoint*

卸载指定的文件系统。也可以为该命令指定系统上 ZFS 文件系统挂载点的路径。

对于已加密的数据集，卸载文件系统时，不会卸载密钥。要卸载密钥，请参见 **zfs key**。

示例

示例1 创建加密数据集

以下示例显示如何使用 passphrase 提示符（即 keysource 属性的缺省值）创建加密数据集。该示例假定 tank/home 数据集未加密。

```
# zfs create -o encryption=on tank/home/bob
Enter passphrase for 'tank/home/bob/': *****
Enter again: *****
```

在以下示例中，`pktool(1)` 命令用于生成文件的原始密钥。接下来，通过 aes-256-ccm 算法和 pktool 生成的原始密钥文件创建加密数据集 (tank/home/anne)。

```
# pktool genkey keystore=file outkey=/media/stick/mykey \
keytype=aes keylen=256
# zfs create encryption=aes-256-ccm \
-o keysource=raw,file:///rmdisk/stick/mykey tank/home/anne
```

该示例显示如何创建提示您密钥存储在 http 位置的 ZFS 加密文件系统。

```
# zfs create -o encryption=on \
-o keysource=passphrase,https://keys.example.com/keys/42 tank/home/fs1
```

该示例显示如何生成 PKCS#11 令牌中的原始密钥。然后，使用从 pktool 生成的原始 PKCS#11 密钥创建加密数据集。

```
# pktool genkey keystore=pkcs11 keytype=aes keylen=128 label=fs2
Enter PIN for Sun Software PKCS#11 softtoken: xxxxx
# zfs create -o encryption=on -o keysource=raw,pkcs11:object=fs2 \
tank/home/fs2
Enter PKCS#11 token PIN for 'tank/home/fs2': xxxxx
```

该示例显示如何生成 KMS 令牌中的原始密钥。然后，使用从 pktool 生成的原始 KMS 密钥创建加密数据集。

```
# pktool genkey keystore=pkcs11 keytype=aes keylen=256 token=KMS \
label=fs3
Enter PIN for KMS: xxxxx
# zfs create -o encryption=aes-256-ccm \
-o keysource="raw,pkcs11:token=KMS;object=fs3" tank/home/fs3
Enter 'KMS' PKCS#11 token PIN for 'tank/home/fs3': xxxxx
```

示例2 通过不同的加密算法创建加密数据集

在该示例中，所有 tank/home 数据集都继承 keysource 属性，但使用不同的加密算法创建 tank/home/bob 数据集。

```
# zpool create tank ....
# zfs create -o encryption=on tank/home
# zfs get keystatus tank/home
NAME          PROPERTY  VALUE          SOURCE
tank/home     keystatus  available      -
```

示例2 通过不同的加密算法创建加密数据集 (续)

```
# zfs create -o encryption=aes-256-ccm tank/home/bob
```

示例3 继承加密和密钥源属性

在该示例中，所有的 tank/home 数据集都继承 encryption 和 keysource 属性。

```
# zpool create -o encryption=on -o keysource=raw,file:///... tank ...
# zfs create tank/home
```

示例4 更改加密数据集的包装密钥和密钥源

该示例显示如何将数据集的包装密钥更改为数据集的 keysource 属性定义的新密钥。

```
# zfs get keysource tank/home/bob
NAME PROPERTY          VALUE                               SOURCE
tank  keysource          raw,file:///etc/keyfile default

# zfs key -c -o keysource=passphrase,prompt tank/home/bob
Enter passphrase for 'tank/home/bob/': *****
Enter again: *****
```

以下示例显示如何更改数据集的包装密钥的 http 位置。

```
# zfs get keysource tank/home/bob
NAME          PROPERTY  VALUE                               SOURCE
tank/home/bob keysource  passphrase,prompt local

# zfs set keysource=passphrase,https://internal.example.com/keys/bob/zfs \
tank/home/bob
```

您必须具有委托的 key 和 keychange 权限才能更改 keysource 属性。

示例5 重设数据集的加密密钥

该示例显示如何更改您或管理员既看不到也未管理的数据集加密密钥。数据集的加密密钥由 keysource 属性中指定的密钥包装（加密）。

```
# zfs key -K tank/project42
# zfs get rekeydate,creation tank/project42
```

您必须具有委托的 keychange 权限才能执行密钥更改操作。

退出状态

将返回以下退出值：

- 0 成功完成。
- 1 出现错误。

2
指定的命令行选项无效。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/zfs
接口稳定性	Committed（已确定）

另请参见 [chmod\(1\)](#)、[chown\(1\)](#)、[pktool\(1\)](#)、[ssh\(1\)](#)、[mount\(1M\)](#)、[zfs\(1M\)](#)、[zpool\(1M\)](#)、[chmod\(2\)](#)、[chown](#)
有关使用其他 ZFS 功能的信息，请参见 [zfs_allow\(1M\)](#)、[zfs_share\(1M\)](#)、[zfs\(1M\)](#)
和 《Oracle Solaris 管理：ZFS 文件系统》。

引用名	<code>zfs_share</code> – 共享和取消共享 ZFS 文件系统
用法概要	<pre> zfs [-?] zfs help subcommand help property property-name permission zfs create [-p] [-o sharenfs=on sharesmb=on] [-o property ... filesystem zfs set share=name=sharename,path=pathname,[desc=description], prot=smb nfs[,property=value[,...]] filesystem zfs set [-r] sharenfs=on off filesystem zfs set [-r] sharesmb=on off filesystem zfs set -c share=name=sharename filesystem zfs get [-r] sharenfs filesystem ... zfs get [-r] sharesmb filesystem ... zfs get [-r] share filesystem ... zfs share -a filesystem zfs unshare -a filesystem mountpoint </pre>
描述	<code>zfs set share</code> 命令用于在 ZFS 存储池中创建 NFS 或 SMB 共享，如 zpool(1M) 中所述。将 <code>sharenfs</code> 或 <code>sharesmb</code> 属性设置为 <code>on</code> 可发布共享。您也可以通过使用 <code>zfs share</code> 命令发布 ZFS 数据集的 NFS 共享，通过使用 <code>zfs unshare</code> 命令撤销发布 ZFS 数据集的 NFS 共享。
本机共享文件系统属性	文件系统属性分为两种类型：本机属性和用户定义的（或“用户”）属性。本机属性用于显示信息或控制 ZFS 行为。此外，本机属性分为可编辑属性和只读属性。 属性继承自父项，但子项所覆盖的除外。有些属性仅适用于某些类型的数据集（文件系统、卷或快照）。 以下本机属性可用于更改 ZFS 文件系统的行为，共享某个文件系统时通常会使用这些本机属性。 使用以下本机 NFS 和 SMB 属性可以更改 ZFS 共享的行为。 nbmand=on off 控制是否应该用 <code>nbmand</code>（非分块强制性锁定）挂载文件系统。此选项用于 SMB 客户机。对此属性所做的更改只有在卸载文件系统并重新挂载后才有效。有关 <code>nbmand</code> 挂载的更多信息，请参见 mount(1M)。 readonly=on off 控制是否可以修改此数据集。缺省值为 <code>off</code>（关闭）。 该属性也可通过其简短列名 <code>rdonly</code> 来引用。

sharenfs=on | off

控制是否作为 NFS 共享发布 ZFS 数据集。您也可以通过使用 `zfs share` 命令发布 ZFS 数据集的 NFS 共享，通过使用 `zfs unshare` 命令撤销发布 ZFS 数据集的 NFS 共享。两种发布 NFS 共享的方法均需要事先设置 NFS 共享属性。有关设置 NFS 共享属性的方法，请参见 `zfs set share` 命令。

在更改 `sharenfs` 属性时，只有在此属性以前设置为 `off` 或在更改此属性前已发布共享的情况下，才会使用已通过 `zfs set share` 命令设置的任何新选项重新发布文件系统共享及继承此属性的任何子项。如果新属性值是 `off`，则撤销发布文件系统共享。

sharesmb=on | off

控制是否作为 SMB 共享发布 ZFS 数据集。您也可以通过使用 `zfs share` 命令发布 ZFS 数据集的 SMB 共享，通过使用 `zfs unshare` 命令撤销发布 ZFS 数据集的 SMB 共享。两种发布 SMB 共享的方法均需要事先设置 SMB 共享属性。有关设置 SMB 共享属性的信息，请参见 `zfs set share` 命令。

由于 SMB 共享需要资源名称，因此根据数据集名称构建唯一的资源名称。构建的名称是数据集名称的一个副本，区别在于资源名称中不合法的数据集名称字符被替换为下划线字符 (`_`)。此外，还支持伪属性 `"name"`，以便您用指定的名称替换数据集名称。然后使用指定的名称替换继承情况下的数据集前缀。例如，如果将数据集 `data/home/john` 设置为 `name=john`，则 `data/home/john` 的资源名称是 `john`。如果 `john` 具有一个子数据集 `data/home/john/backups`，则其资源名称是 `john_backups`。

在创建 SMB 共享时，以 `.zfs/shares` 目录中的条目显示 SMB 共享名称。您可以使用 `ls` 或 `chmod` 命令在此目录的条目上显示共享级别的 ACL。

在更改 `sharesmb` 属性时，如果此属性以前设置为 `off` 或在更改此属性前已发布共享，则会使用已通过 `zfs set share` 命令设置的任何新选项重新发布数据集共享及继承此属性的任何子项。如果新属性值是 `off`，则撤销发布数据集共享。

vscan=on | off

控制打开和关闭常规文件时是否应为其扫描病毒。除了启用此属性以外，为了进行病毒扫描也必须启用病毒扫描服务。缺省值为 `off`（关闭）。

在创建文件系统后无法更改以下属性，因此在创建文件系统时应设置这些属性。如果未使用 `zfs create` 或 `zpool create` 命令设置属性，则从父数据集继承这些属性。如果在这些功能获得支持以前已创建的父数据集缺少这些属性，则新文件系统将为这些属性采用缺省值。

casesensitivity=sensitive | insensitive | mixed

指示文件系统使用的文件名匹配算法是否应区分大小写、不区分大小写或允许混用大小写的组合匹配方式。`casesensitivity` 属性的缺省值为 `mixed`。传统上，UNIX 和 POSIX 文件系统的文件名区分大小写。

`casesensitivity` 属性的值 `mixed` 指示文件系统对区分大小写和不区分大小写的匹配行为要求均可支持。当前，在支持混合行为的文件系统上，不区分大小写的匹配

行为仅限于 Solaris SMB 服务器产品。有关 `mixed` 值的行为的更多信息，请参见《[Oracle Solaris 管理：ZFS 文件系统](#)》。

`normalization = none | formC | formD | formKC | formKD`

指示每当比较两个文件名时，文件系统是否应为文件名执行 `unicode` 标准化，以及应使用哪种标准化算法。文件名在存储时始终保持未修改状态，并将其标准化作为任何比较进程的一部分。如果将此属性设置为 `none` 以外的合法值，并且不指定 `utf8only` 属性，则 `utf8only` 属性自动设置为 `on`。`normalization` 属性的缺省值为 `none`。在创建文件系统后无法更改此属性。

`utf8only=on | off`

指示在文件名含有 UTF-8 字符代码集中不存在的字符情况下，文件系统是否应拒绝此类文件名。如果已将此属性显式设置为 `off`，则不允许将规范化属性显式设置或常规设置为 `none`。`utf8only` 属性的缺省值为 `off`。在创建文件系统后无法更改此属性。

子命令

修改状态的所有子命令将以原始形式永久记录到池中。

`zfs ?`

显示帮助消息。

`zfs help command | help | property property-name | permission`

显示 `zfs` 命令的用法信息。可以显示特定命令、属性或委托权限的帮助。如果显示特定命令或属性的帮助，将显示命令语法或属性值。在不带任何参数的情况下使用 `zfs help` 将显示完整的 `zfs` 命令列表。

`zfs create [-p] [-o sharenfs=on | sharesmb=on -o property=value] ... filesystem`

创建新的 ZFS 文件系统。该文件系统会根据从父项继承的 `mountpoint` 属性自动挂载。

`-o property=value`

与在创建数据集时调用命令 `zfs set property=value` 一样设置指定的属性。在创建时也可设置任何可编辑的 ZFS 属性。可以指定多个 `-o` 选项。如果在多个 `-o` 选项中指定了相同的属性，则会导致错误。

`zfs get [-r] share [filesystem]`

显示所有定义的共享，或显示指定文件系统或后代文件系统的定义共享。

`zfs set share=name=sharename,path=pathname,[desc=description],prot=smb | nfs[,property=value[,...]] filesystem`

定义 ZFS 数据集的 NFS 或 SMB 文件共享属性。建立 NFS 或 SMB 文件共享属性后，可以通过将 `sharenfs` 或 `sharesmb` 属性设置为 `on` 来发布 ZFS 文件系统的 NFS 或 SMB 共享。

需要以下共享属性：

`share=name=sharename`

标识 NFS 或 SMB 共享名称。共享名称的最大长度为 80 个字符。

path=pathname

标识要共享的数据集的物理路径或要共享的数据集中的子目录。

prot=smb|nfs

标识 NFS 或 SMB 文件共享协议。

以下共享属性是可选的：

desc=description

标识用于描述共享资源的文本字符串。描述中的空格或逗号必须括在引号中 ("")。

ro=|rw=

标识共享将以读/写还是只读形式提供给所有客户机。您也可以指定包含主机名、IP 地址或网络组的逗号分隔列表。

root=

标识指定主机的超级用户或拥有根访问权限的主机列表。缺省情况下，没有主机拥有根访问权限。

sec=

标识 NFS 服务器安全模式，例如 `sys`、`dh`、`krb5` 等。有关支持的安全模式信息，请参见 [nfssec\(5\)](#)。

必须在 `prot=nfs` 之后但在任何 `sec` 属性之前指定以下可选的 NFS 属性。

```
anon=TYPE_USER
nosub=[true|false]
nosuid=[true|false]
aclok=[true|false]
public=[true|false]
index=filename
log=TYPE_LOGTAG
cksum=TYPE_STRINGSET
```

必须在 `prot=smb` 之后指定以下可选的 SMB 属性：

```
ad-container=string
abe=[true|false]
csc=[disabled|manual|auto|vdo]
catia=[true|false]
guestok=[true|false]
ro=access-list
rw=access-list
none=access-list
```

有关 NFS 和 SMB 共享属性的详细说明，请参见 [share_nfs\(1M\)](#) 和 [share_smb\(1M\)](#)。

zfs set [-r] property=value filesystem|volume|snapshot ...

对每个数据集将该属性设置为给定值。只能编辑某些属性。有关可设置哪些属性和可接受的值的更多信息，请参见“属性”一节。

-r

在子数据集的整个子树中递归应用设置的有效值。根据属性，可以设置或继承有效值。

zfs set -c share=name=sharename filesystem

删除文件系统共享。-c 选项将此子命令与上面介绍的 **zfs set share** 命令区别开来。

zfs share -a | dataset

根据共享属性值并且如果已设置 **sharenfs** 或 **sharesmb** 属性，发布 ZFS 数据集的 NFS 或 SMB 共享。

通过 NFS 或 SMB 协议共享数据集意味着，数据集数据将在网络上可用。设置了 **sharenfs** 或 **sharesmb** 属性的 ZFS 数据集将在引导系统时自动共享。

-a

根据共享属性值以及 **sharenfs** 和 **sharesmb** 属性的设置，共享所有 ZFS 数据集。

dataset

根据共享属性值以及 **sharenfs** 和 **sharesmb** 属性的设置，共享指定的数据集。

zfs unshare -a | dataset|mountpoint

取消共享设置了 **sharenfs** 或 **sharesmb** 属性的所有 ZFS 数据集。

-a

取消共享所有共享的 ZFS 数据集。自动在引导过程中调用。

dataset|mountpoint

取消共享指定的数据集。也可以为该命令指定系统上共享的 ZFS 数据集的路径。

示例

示例 1 创建 ZFS 文件系统的 NFS 或 SMB 共享

创建 ZFS 文件系统的 NFS 或 SMB 共享的过程包含以下两个步骤：

1. 使用 **zfs set share** 命令创建共享以识别 NFS 或 SMB 共享属性。

要创建 NFS 共享，请执行以下命令：

```
# zfs create -o mountpoint=/data tank/data
# zfs set share=name=data,path=/data,prot=nfs tank/data
```

要创建 SMB 共享，请执行以下命令：

```
# zfs set share=name=fs,path=/fs,prot=smb,ro=@10-132-60-234 \
tank/fs
```

2. 通过设置 **sharenfs** 或 **sharesmb** 属性发布共享。使用诸如以下的命令：

```
# zfs set sharenfs=on tank/data
# zfs set sharesmb=on tank/fs
```

示例2 显示 NFS 或 SMB 共享信息

显示 NFS 或 SMB 共享信息：

```
# zfs get share
NAME                                PROPERTY  VALUE  SOURCE
tank/data                          share     name=data,path=/data,prot=nfs  local
tank/fs                             share     name=fs,path=/fs,prot=smb,ro... local
```

示例3 添加或更改 ZFS 文件系统上的共享属性

如下所示，通过添加新选项（如 hostA 的只读）更改共享：

```
# zfs set share=name=data,prot=nfs,ro=hostA tank/data
```

示例4 删除 ZFS 文件系统的 NFS 或 SMB 共享

使用 zfs set -c 命令删除共享：

```
# zfs set -c share=name=fs tank/fs
share 'fs' was removed.
```

退出状态

将返回以下退出值：

- 0
成功完成。
- 1
出现错误。
- 2
指定的命令行选项无效。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/zfs
接口稳定性	Committed（已确定）

另请参见

[chmod\(1\)](#)、[chown\(1\)](#)、[pktool\(1\)](#)、[ssh\(1\)](#)、[mount\(1M\)](#)、[shadowd\(1M\)](#)、[share\(1M\)](#)、[share_nfs\(1M\)](#)

请参见 [gzip\(1\)](#) 手册页，该页不属于 SunOS 手册页集合。

有关使用其他 ZFS 功能的信息，请参见《[Oracle Solaris 管理：ZFS 文件系统](#)》。

引用名

zic – time zone compiler

用法概要

zic [--version] [-s] [-v] [-d *directory* | -l *localtime*]
[-p *posixrules*] [-y *yearistype*] [*filename*]...

描述

zic reads text from the file(s) named on the command line and creates the time conversion information files specified in this input. If a *filename* is '-', the standard input is read.

Input lines are made up of fields. Fields are separated by any number of white space characters. Leading and trailing white space on input lines is ignored. A pound sign (#) indicates a comment that extends to the end of the line. White space characters and pound signs can be enclosed within double quotes (") if they are to be used as part of a field. Any line that is blank (after comment stripping) is ignored. Non-blank lines are expected to be of one of three types: rule lines, zone lines, or link lines.

Rule

A rule line has the form:

For example:

Rule NAME FROM TO TYPE IN ON AT SAVE LETTER/S

The fields that make up a rule line are:

Rule USA 1969 1973 - Apr lastSun 2:00 1:00 D

- NAME Gives the (arbitrary) name of the set of rules this rule is part of.
- FROM Gives the first year in which the rule applies. The word *minimum* (or an abbreviation) means the minimum year with a representable time value. The word *maximum* (or an abbreviation) means the maximum year with a representable time value.
- TO Gives the final year in which the rule applies. In addition to *minimum* and *maximum* (as above), the word *only* (or an abbreviation) can be used to repeat the value of the FROM field.
- TYPE Gives the type of year in which the rule applies. If TYPE is:
- '-' The rule applies in all years between FROM and TO, inclusive.
 - uspres The rule applies in U.S. Presidential election years.
 - nonpres The rule applies in years other than U.S. Presidential election years.
 - even The rule applies to even-numbered years.
 - odd The rule applies to odd-numbered years.
- If TYPE is something else, then zic will attempt to execute the command *yearistype year type*

to check the type of a year: an exit status of 0 means that the year is of the given type; an exit status of 1 means that the year is not of the given type. The yearistype command is not currently provided in the Solaris environment.

IN	Names the month in which the rule takes effect. Month names can be abbreviated.										
ON	Gives the day on which the rule takes effect. Recognized forms include: <table><tr><td>5</td><td>the fifth day of the month</td></tr><tr><td>lastSun</td><td>The last Sunday in the month</td></tr><tr><td>lastMon</td><td>The last Monday in the month</td></tr><tr><td>Sun>=8</td><td>First Sunday on or after the eighth</td></tr><tr><td>Sun<=25</td><td>Last Sunday on or before the 25th</td></tr></table> Names of days of the week can be abbreviated or spelled out in full. Note: There cannot be spaces within the ON field.	5	the fifth day of the month	lastSun	The last Sunday in the month	lastMon	The last Monday in the month	Sun>=8	First Sunday on or after the eighth	Sun<=25	Last Sunday on or before the 25th
5	the fifth day of the month										
lastSun	The last Sunday in the month										
lastMon	The last Monday in the month										
Sun>=8	First Sunday on or after the eighth										
Sun<=25	Last Sunday on or before the 25th										
AT	Gives the time of day at which the rule takes effect. Recognized forms include: <table><tr><td>2</td><td>Time in hours</td></tr><tr><td>2:00</td><td>Time in hours and minutes</td></tr><tr><td>15:00</td><td>24-hour format time (for times after noon)</td></tr><tr><td>1:28:14</td><td>Time in hours, minutes, and seconds, where hour 0 is midnight at the start of the day and hour 24 is midnight at the end of the day.</td></tr></table> Any of these forms can be followed by the letter w if the given time is local “wall clock” time; s if the given time is local “standard” time; or u (or g or z) if the given time is universal time. In the absence of an indicator, wall clock time is assumed.	2	Time in hours	2:00	Time in hours and minutes	15:00	24-hour format time (for times after noon)	1:28:14	Time in hours, minutes, and seconds, where hour 0 is midnight at the start of the day and hour 24 is midnight at the end of the day.		
2	Time in hours										
2:00	Time in hours and minutes										
15:00	24-hour format time (for times after noon)										
1:28:14	Time in hours, minutes, and seconds, where hour 0 is midnight at the start of the day and hour 24 is midnight at the end of the day.										
SAVE	Gives the amount of time to be added to local standard time when the rule is in effect. This field has the same format as the AT field (without the w and s suffixes).										
LETTER/S	Gives the “variable part” (for example, the “S” or “D” in “EST” or “EDT” of time zone abbreviations to be used when this rule is in effect. If this field is '-', the variable part is null.										

Zone A zone line has the form:

Zone NAME GMTOFF RULES/SAVE FORMAT [UNTIL]

For example:

```
Zone Australia/SouthWest 9:30      -      CST    1992 Mar 15 12:00
                        8:30      Aus    CST
```

The fields that make up a zone line are:

- NAME** The name of the time zone. This is the name used in creating the time conversion information file for the zone.
- GMTOFF** The amount of time to add to UTC to get standard time in this zone. This field has the same format as the AT and SAVE fields of rule lines; begin the field with a minus sign to subtract time from UTC.
- RULES/SAVE** The name of the rule(s) that apply in the time zone or, alternately, an amount of time to add to local standard time. If this field is '-', then standard time always applies in the time zone.
- FORMAT** The format for time zone abbreviations in this time zone. The pair of characters %s is used to show where the "variable part" of the time zone abbreviation goes. Alternately, a slash (/) separates standard and daylight abbreviations.
- UNTIL** The time at which the UTC offset or the rule(s) change for a location. It is specified as a year, a month, a day, and a time of day. The time of day has the same format as the AT field of rule lines. If this is specified, the time zone information is generated from the given UTC offset and rule change until the time specified.

The month, day, and time of day have the same format as the IN, ON, and AT columns of a rule; trailing columns can be omitted, and default to the earliest possible value for the missing columns.

The next line must be a "continuation" line. This line has the same form as a zone line except that the string "Zone" and the name are omitted. The continuation line places information starting at the time specified as the UNTIL field in the previous line in the file used by the previous line. Continuation lines can contain an UNTIL field, just as zone lines do, indicating that the next line is a further continuation.

Link A link line has the form:

```
Link    LINK-FROM    LINK-TO
```

For example:

```
Link    Europe/Istanbul Asia/Istanbul
```

The LINK-FROM field should appear as the NAME field in some zone line; the LINK-TO field is used as an alternate name for that zone.

	Except for continuation lines, lines can appear in any order in the input.	
选项	- -version	Outputs version information and exits.
	-d directory	Creates time conversion information files in the directory <i>directory</i> rather than in the standard directory <code>/usr/share/lib/zoneinfo</code> .
		The -d and -l options are mutually exclusive.
	-l localtime	Uses the given time zone as local time <i>localtime</i> . zic acts as if the file contained a link line of the form: Link <i>localtime</i> localtime
		The -d and -l options are mutually exclusive.
	-p posixrules	Uses the rules of the given time zone <i>posixrules</i> when handling POSIX-format time zone environment variables. zic acts as if the input contained a link line of the form: Link <i>posixrules</i> posixrules
		This option is not used by <code>ctime(3C)</code> and <code>mktime(3C)</code> in the Solaris environment.
	-s	Limits time values stored in output files to values that are the same whether they are taken to be signed or unsigned. You can use this option to generate SVVS-compatible files. This option is obsolete and may be removed in a future release.
操作数	-v	Complains if a year that appears in a data file is outside the range of years representable by system time values (0:00:00 a.m. UTC, January 1, 1970, to 3:14:07 a.m. UTC, January 19, 2038). This option also complains if a time of 24:00 (which cannot be handled by pre-1998 versions of zic) appears in the input.
	-y yearistype	Uses the given command <i>yearistype</i> rather than <code>yearistype</code> when checking year types (see Rules under DESCRIPTION).
文件	filename	A file containing input lines that specify the time conversion information files to be created. If a <i>filename</i> is '-', the standard input is read.
属性	/usr/share/lib/zoneinfo	Standard directory used for created files
	/usr/share/lib/zoneinfo/src	Directory containing source files
	See attributes(5) for descriptions of the following attributes:	

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/core-os
Interface Stability	Committed*

* The -s option is obsolete.

另请参见

[time\(1\)](#), [zdump\(1M\)](#), [ctime\(3C\)](#), [mktime\(3C\)](#), [attributes\(5\)](#)

附注

For areas with more than two types of local time, you might need to use local standard time in the AT field of the earliest transition time's rule to ensure that the earliest transition time recorded in the compiled file is correct.

If the current *timezone* file is edited and compiled using the “zic” command, the changes will only be reflected in any new processes that are running. The most accurate way to reflect the changes for the whole system would be a reboot.

引用名	zoneadm – 管理区域
用法概要	<pre>zoneadm -z zonename [-u uuid-match] subcommand [subcommand_options] zoneadm [-R root] [-z zonename] [-u uuid-match] list [list_options] zoneadm [-R root] -z zonename [-u uuid-match] mark incomplete</pre>
描述	zoneadm 实用程序可用于管理系统区域。区域是指由操作系统运行时维护的应用程序容器。
安全	一旦将进程放入 0 区域之外的区域，此进程或其任何子进程便不能更改区域。
安全	除了简单的列表和帮助功能，zoneadm 只能由在全局系统区域中进行操作的用户使用，且必须使用有效的根目录用户 ID 来执行。此外，用户要执行特定的子命令还必须经过授权。 zoneadm 会检查授权字符串，这些字符串可能会选择性地包括用作后缀的指定 <i>zonename</i>，以及用作前缀的斜杠字符。如果省略区域名称，用户将有权进入任何区域。 只提供信息的子命令（例如 <i>help</i> 或 <i>list</i>）不需任何授权。所有其他子命令都需要授权 <i>solaris.zone.manage/zonename</i>。 一旦将进程放入 0 区域之外的区域，该进程或其任何子进程便不能更改区域。
选项	所支持的选项如下： -R root 指定备用根（引导环境）。此选项只能与 <i>list</i> 和 <i>mark</i> 子命令配合使用。 -u uuid-match 区域的唯一标识符，由 libuuid(3LIB) 分配。如果该选项存在且参数为非空字符串，则会选择与 UUID 相匹配的区域（如果存在这样的区域），而非 <i>-z</i> 选项指定的区域。 -z zonename 区域的字符串标识符。
子命令	会导致破坏性操作或丢失工作的子命令有一个强制执行操作的 -F 标志。如果输入来自某个终端设备，则当此类命令不带 -F 标志时，系统会提示用户；否则，如果此类命令不带 -F 标志，则系统将禁止操作，并将诊断消息写入标准错误中。如果区域安装或卸载中断，该区域将处于未完成状态。使用 <i>uninstall</i> 可将此类区域重置为已配置状态。 支持以下子命令：

`attach [-u] [-F] [-n path] [brand-specific options]`

`attach` 子命令获取从某个系统分离的区域，然后将该区域连接到新的系统。因此，我们建议（而非要求）您在进行“连接”操作之前，先运行 `detach` 子命令。将新区域设置为已配置状态后，可使用 `attach` 子命令来设置区域根目录，不必以新区域的形式来安装该区域。

可使用 `-F` 选项在未经验证的情况下将区域强行设置为 "installed"（已安装）状态。此选项应慎用，因为在将区域从源系统移至无法正常托管该区域的目标系统时，使用此选项会令区域处于不受支持的状态。`-n` 选项可用于执行 `attach` 子命令的“预运行”。此选项将 `detach -n` 子命令的输出用作输入，可用于确认是否存在冲突问题（如网络设备不兼容），以及主机能否支持该区域。路径可以是 `.`，用于读取来自标准输入设备的输入。

区域的标记可能包括其他用于管理区域连接方式的选项。具体标记信息请参见 [brands\(5\)](#)。

要连接的区域必须首先使用 `zonecfg`（请参见 [zonecfg\(1M\)](#)）命令进行配置。运行 `attach -n` 时，上述规定不适用。

请使用以下命令来连接区域：

```
# zoneadm -z my-zone attach
```

请使用以下命令来连接和更新区域：

```
# zoneadm -z my-zone attach -u
```

缺少 `-n` 选项时（如上所示），必须先停止源区域，然后才能使用该子命令。

`-n path`

读取区域清单并在不实际执行连接操作的情况下，验证目标机器是否拥有托管该区域所需的正确配置。在执行试验性附加之前，不必在新主机上配置目标系统中的区域。

`-u`

更新连接的区域。

`boot [-w|-W] [-- boot_options]`

引导（或激活）指定区域。

`boot` 子命令具有以下互斥选项：

`-w`

引导包含可写根的区域，有效地覆盖区域配置中的 `file-mac-profile` 设置。此选项仅对本次引导周期有效：后续重新引导将引导区域，使 `file-mac-profile` 再次有效。

`-W`

以瞬态读取/写入模式引导区域；当区域完成自组合时，该区域将以只读模式重新引导。对非只读根区域没有任何影响。

支持以下 *boot_options* :

-i altinit

选择一个备用可执行文件作为初始进程。*altinit* 是指向可执行文件的有效路径。缺省的初始进程为 *init(1M)*。

-m smf_options

smf_options 包括两类控制服务管理工具引导行为的选项：恢复选项和消息选项。

消息选项决定了 *smf(5)* 在引导过程中显示的消息的类型和数量。服务选项决定了用于引导系统的服务。有关 *-m* 子选项的列表，请参见 *kernel(1M)*。

-s

仅引导至里程碑 *svc:/milestone/single-user:default*。此里程碑相当于 *init* 级别 *s*。请参见 *svc.startd(1M)* 和 *init(1M)*。

clone [-m copy] [-s zfs_snapshot] [brand-specific options] source_zone

安装某个区域时，可直接复制现有的已安装区域。该子命令是一种备选的区域安装方法。

-m copy

强制克隆成为副本，即使可以进行“ZFS clone”也是如此。

-s zfs_snapshot

指定用作克隆源的 ZFS 快照的名称。此处的 *snapshot* 必须是从之前的 *zoneadm clone* 安装获取的源区域的 *snapshot*。

必须先停止源区域，然后才能使用该子命令。

detach [-F] -n]

分离指定的区域。将区域从一个系统移至另一个系统时，第一步是分离该区域。区域迁移的完整过程如下：分离区域，将 *zonepath* 目录移至新主机，然后再将该区域连接到新主机。区域分离后，将一直处于已配置状态。如果尝试安装某个已分离的配置区域或者要克隆到该区域，您将接收到一条错误消息，且 *install* 或 *clone* 子命令无法继续。*-n* 选项可用于执行 *detach* 子命令的“预运行”。在这一过程中，会生成运行 *attach -n* 子命令所需的信息，该信息可用于确认是否存在冲突问题（如网络设备不兼容）或主机能否支持该区域。该信息将发送至标准输出，并可保存到文件中或传输至 *attach -n* 子命令。*-F* 选项可用于强制分离区域，而无需对现有的 *zonepath* 执行验证检查。

请使用以下命令来分离区域：

```
# zoneadm -z my-zone detach
```

除非使用 *-n* 选项，否则必须停止源区域，然后才能使用该子命令。

-F

强制分离区域而无需对区域的存储执行验证检查。如果该主机无法再访问 *zonepath* 的存储，则通常使用该选项。当区域的存储已故障转移到备用主机（手动执行或作为群集的一部分）时，通常出现此情况。

-n

在不实际分离某区域的情况下，生成关于该运行区域的区域清单。源系统中区域的状态不会改变。区域清单会被发送到 `stdout`。全局管理员可以将此输出定向到某一文件，或将此输出传输到远程命令以便立即在目标主机上进行验证。

halt

停止指定的区域。`halt` 跳过在区域内运行关闭脚本这一步骤，并删除区域的运行时资源。

另请参见下面的 `shutdown` 子命令。

help [*subcommand*]

显示一般帮助。如果您指定 *subcommand*，则会显示有关 *subcommand* 的帮助。

install [*brand-specific options*]

将指定区域安装到系统中。该子命令会自动首先尝试进行验证。如果验证步骤失败，则会拒绝安装。请参见 `verify` 子命令。

区域的标记可能包括其他用于管理区域中软件安装方式的选项。具体标记信息请参见 [brands\(5\)](#)。

list [*list_options*]

显示当前区域或指定区域（如果指定）的名称。

缺省情况下，列出所有运行区域。与 `zoneadm -z zonename` 选项结合使用时，此子命令只列出指定区域（不管其状态如何）。这种情况下，将禁用 `-i` 和 `-c` 选项。

如果既未给出 `-i` 选项，也未给出 `-c` 选项，则会列出所有运行区域。

支持以下 *list_options*：

-c

显示所有已配置区域。此选项将覆盖 `-i` 选项。

-i

将显示扩展到所有已安装区域。

-p

请求可供机器解析的输出。输出格式是一系列行，每个区域一行，其中包含以冒号分隔的字段。这些字段包括：

```
zoneid:zonename:state:zonepath:uuid:brand:ip-type:\
r/w:file-mac-profile
```

如果 *zonepath* 包含嵌入的冒号，则可通过反斜杠“\”对这些冒号进行转义，并可通过使用 `shell` 函数 `read(1)` 及环境变量 `IFS` 来解析该反斜杠。*uuid* 值是在安装区域时由 `libuuid(3LIB)` 分配的，可用于标识存在于备用引导环境中或已在该环境中重命名的同一区域。任何可解析 `zoneadm list -p` 命令输出的软件都必须能够处理将来可能会添加的任何字段。

`-a` 和 `-d` 选项相互排斥。如果 `-v` 和 `-p` 均未使用，则只列出区域名称。

-v

显示详细信息，包括区域名称、ID、当前状态、根目录、标记类型、ip-type 和选项。

-a 和 **-d** 选项相互排斥。如果 **-v** 和 **-p** 均未使用，则只列出区域名称。

mark incomplete

将某个已安装区域的状态更改为 "incomplete"（不完整）。对系统的管理性更改导致区域不可用或不一致时，此命令可能会很有用。此更改无法撤消（除非采用卸载整个区域的方式）。

move new_zonepath

将 *zonepath* 移至 *new_zonepath*。必须先停止区域，然后才能使用该子命令。*new_zonepath* 必须是本地文件系统，适用针对 *zonepath* 的一般限制。

ready

让区域做好运行应用程序的准备，但不在区域中启动任何用户进程。

reboot [--boot_options]

重新启动区域。这等效于 **haltboot** 序列（关机脚本未运行）。

有关支持的引导选项，请参见 **boot** 子命令。

shutdown [-r [--boot_options]]

完全关闭区域（等效于在区域中运行 **/usr/sbin/init 0**）。**shutdown** 子命令会一直等到区域成功关闭；如果关机过程需要很长时间，可以使用 **zoneadmhalt** 强制停止该区域。

如果指定了 **-r**，则重新引导区域。有关支持的引导选项，请参见 **boot** 子命令。

uninstall [-F]

从系统中卸载指定区域。请慎用此子命令。它会删除指定区域的 *zonepath* 下的所有文件。您可以使用 **-F** 标志强制执行操作。

verify

检查确认指定区域的配置可以安全地安装在计算机上。下面是按 **resource/property**（资源/属性）类型分类的检查：

zonepath

zonepath 及其父目录均存在，并属于相应模式的根目录。相应模式是指当 *zonepath* 为 **700** 时，其父目录既不是 **group**，也不是 **world-writable** 之类的目录。*zonepath* 不通过 NFS 挂载。不存在名为 **root** 的 *zonepath* 子目录。

如果 *zonepath* 不存在，**verify** 子命令不会失败，而只会警告用户：后续安装会尝试创建该目录并设置相应的权限。如果出现错误，后续的 **verify** 子命令可能会失败。

zonepath 不能是符号链接。

fs

任何 **fs** 资源都会对其 *type* 值进行检查。如果值是 **proc**、**mntfs**、**autofs** 或 **nfs** 中的一个，或者文件系统在 **/usr/lib/fs/<fstype>/mount** 下没有关联的挂载二进制数据，则会报告一个错误。

这是一种将 *directory* 设置为相对路径时会出现的错误。

这是一种将 **raw** 所指定的路径设置为相对路径或者在 **/usr/lib/fs/<fstype>/fsck** 下没有给定文件系统类型的 **fsck** 二进制数据时会出现的错误。如果存在相应的 **fsck** 二进制数据但未指定 **raw** 路径，则也会出现这种错误。

net

所有物理网络接口都存在。所有网络地址资源都属于下面某种类型：

- 有效的 IPv4 地址，可以后接 / 和前缀长度；
- 有效的 IPv6 地址，必须后接 / 和前缀长度；
- 可解析为 IPv4 地址的主机名。

请注意，不支持解析为 IPv6 地址的主机名。

物理接口名称即网络接口名称。

可将区域配置为专用 IP 区域或共享 IP 区域。共享 IP 区域必须同时设置物理属性和地址属性。专用 IP 区域必须设置物理属性，不能设置地址属性。

anet

它验证用来自动创建 VNIC 的下游链路 (lower-link) 是否存在。

rctl

它也可以用于验证任何定义的资源控制值在当前计算机上是否有效。这意味着特权级别为 **privileged**（特权），限制低于当前定义的系统值，且定义的操作与给定资源控制的有效操作一致。

示例**示例 1 使用 -m 选项**

以下命令说明了如何使用 **-m** 选项。

```
# zoneadm boot -- -m verbose
```

示例 2 使用 -i 选项

以下命令说明了如何使用 **-i** 选项。

```
# zoneadm boot -- -i /usr/sbin/init
```

示例 3 使用 -s 选项

以下命令说明了如何使用 **-s** 选项。

```
# zoneadm boot -- -s
```


引用名 zoneadmd – 区域管理守护进程

用法概要 /usr/lib/zones/zoneadmd

描述 zoneadmd 是一个系统守护进程，在系统需要管理特定区域时启动。由于 zoneadmd 守护进程的每个实例都管理一个特定区域，因此看到多个 zoneadmd 守护进程同时运行是很正常的。

此守护进程由区域管理软件自动启动，而不应直接调用。守护进程在不再使用时将自动关闭。它不属于编程接口，而应归类为专用接口。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/zones
接口稳定性	Private（专用）

另请参见 [svcs\(1\)](#)、[zlogin\(1\)](#)、[svcadm\(1M\)](#)、[zoneadm\(1M\)](#)、[attributes\(5\)](#)、[smf\(5\)](#)、[zones\(5\)](#)

附注 [zones\(5\)](#) 服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

svc:/system/zones:default

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作，如启用、禁用或请求重新启动。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

引用名	zonecfg – 设置区域配置
用法概要	<pre>zonecfg -z zonename zonecfg -z zonename subcommand zonecfg -z zonename -f command_file zonecfg help</pre>
描述	zonecfg 实用程序可以创建、修改和列出区域的配置。创建和修改功能只可供授权的用户使用，并且要求使用 root 的有效用户 ID 来执行该进程。否则，它将以只读模式运行。 区域的配置由若干个资源和属性组成。 为了简化用户界面，zonecfg 采用了范围的概念。缺省范围为全局。 zonecfg 命令的以下用法概要适用于交互式使用模式： <pre>zonecfg -z zonename subcommand</pre> 通过 zonecfg 更改的参数不会影响正在运行的区域。必须重新引导区域才能使更改生效。 除了创建和修改区域，zonecfg 实用程序还可以用来持久性地指定全局区域的资源管理设置。 在下面的文本中，rctl 表示 resource control（资源控制）的缩写。请参见 resource_controls(5)。 每个区域都配有一个关联标记。该标记决定了在区域中使用的用户级环境，以及在安装、引导或关闭区域时区域的各种行为。在安装区域后，无法更改该标记。缺省标记由全局区域中安装的分发软件决定。有些标记不能支持某些 zonecfg 属性和资源。有关每个标记的更多详细信息，请参见特定于标记的手册页。有关标记的概述，请参见 brands(5) 手册页。
资源	支持下列资源类型： <pre>attr 通用属性。 capped-cpu CPU 的使用限制。 capped-memory 物理内存、交换内存和锁定内存的限制。 dataset ZFS 数据集。 dedicated-cpu 当此区域运行时，专用于此区域的系统处理器的子集。</pre>

device
设备。

fs
文件系统。

net
网络接口。

anet
自动网络接口。

admin
受委托的管理员。

rctl
资源控制。

稀疏根和完全根非全局区域

以前版本的 Solaris 提供了**稀疏根区域**概念。该功能与 SVr4 包管理系统密切相关，旨在节省磁盘空间并减少管理工作。

在选择要在区域中安装哪些包方面，新的包管理系统 IPS 提供了更大的灵活性。这种灵活性以及文件系统技术方面的优势（其中值得一提的是 ZFS **重复数据删除**）意味着删除稀疏根区域最为明智。通过组合利用 IPS 包管理和文件系统的优点，可以使所有区域都具备稀疏根区域的优点。

属性

每种资源类型有一个或多个属性。另外还有一些全局属性，即整个配置的属性，而不是某个特定资源的属性。

支持下列属性：

(全局)
zonename

(全局)
zonepath

(全局)
autoboot

(全局)
bootargs

(全局)
pool

(全局)
limitpriv

(全局)
brand

```
(全局)
  cpu-shares

(全局)
  hostid

(全局)
  max-lwps

(全局)
  max-msg-ids
(全局)
  max-processes

(全局)
  max-sem-ids

(全局)
  max-shm-ids

(全局)
  max-shm-memory

(全局)
  scheduling-class

(全局)
  fs-allowed

(全局)
  file-mac-profile

fs
  dir、special、raw、type、options

net
  address、allowed-address、configure-allowed-address、physical、defrouter

anet
  linkname、lower-link、allowed-address、configure-allowed-address、defrouter、mac-

device
  match

rctl
  name、value

attr
  name、type、value

dataset
  name、alias
```

```

dedicated-cpu
    ncpus、 importance

capped-memory
    physical、 swap、 locked

capped-cpu
    ncpus

admin
    user、 auths

```

对于与这些名称配对的属性值，它们可能是简单值、复杂值，也可能是列表。允许的类型是特定于属性的。简单值是字符串，可根据需要用引号引起来。复杂值具有语法：

```
(<name>=<value>,<name>=<value>,...)
```

其中每个 *<value>* 都是简单值，*<name>* 字符串在给定的属性中是唯一的。列表具有语法：

```
[<value>,...]
```

其中每个 *<value>* 可以是简单值或复杂值。只包含单个值（简单值或复杂值）的列表等效于在不使用列表语法的情况下指定该值。例如，"foo" 等效于 "[foo]"。列表可以为空（用 "[]" 表示）。

在解释属性值时，zonecfg 接受如 [fnmatch\(5\)](#) 中所指定的正则表达式。请参见“示例”部分。

属性类型如下所述：

全局：zonename
区域的名称。

全局：zonepath
区域的文件系统的路径。

全局：autoboot
布尔值，指示在系统引导时是否应自动引导某个区域。请注意，如果禁用了区域服务，则无论此属性的设置如何，都不会自动引导区域。可以使用 `svcadm` 命令启用区域服务，例如：

```
# svcadm enable svc:/system/zones:default
```

将 `enable`（启用）替换为 `disable`（禁用）可禁用区域服务。请参见 [svcadm\(1M\)](#)。

全局：bootargs
要传递到区域引导的参数（选项）；但如果向 "zoneadm boot" 命令提供了选项，则将优先采用提供给该命令的选项。有效的参数在 [zoneadm\(1M\)](#) 中进行了描述。

全局: pool

此区域在引导时必须绑定到的资源池的名称。此属性与 `dedicated-cpu` 资源不兼容。

全局: limitpriv

此区域中的任何进程可以获取的最大特权集。此属性应包含以逗号分隔的特权集指定，如 `priv_str_to_set(3C)` 中所述。可以从结果集中排除特权集，方法是在其名称前加上连字符 (-) 或叹号 (!)。在此上下文中，不支持特殊特权字符串 "zone"。如果特殊字符串 "default" 作为第一个标记出现在该属性中，它将扩展为一个安全的特权集，其中保留了 `zones(5)` 中所述的资源和安全隔离。缺少该属性或该属性为空等效于该安全特权集。

系统管理员在为区域配置特权时必须特别注意。无法通过此机制排除某些特权，因为它们是引导区域时所必需的。此外，不能将某些特权授予给区域，因为如果授予这些特权可能会导致一个区域中的进程不当地影响其他区域中的进程。在尝试“引导”或“准备”区域时，`zoneadm(1M)` 指示何时已在区域的特权集中添加或删除了某个无效的特权。

有关特权的描述，请参见 `privileges(5)`。命令 `ppriv -l`（请参见 `ppriv(1)`）可以生成所有 Solaris 特权的列表。您可以按 `ppriv` 显示的格式指定特权。在 `privileges(5)` 中，特权以 `PRIV_privilege_name` 格式列出。例如，在此属性中指定的特权 `sys_time` 在 `privileges(5)` 中将显示为 `PRIV_SYS_TIME`。

全局: brand

区域的标记类型。

全局: ip-type

区域既可以拥有自己的独占 IP 实例（缺省），也可以与全局区域共享 IP 实例。在缺省区域模板 `SYSdefault` 中，`ip-type` 设置为 `exclusive`。在提供的另外一个 `SYSdefault-shared-ip` 模板中，`ip-type` 设置为 `shared`。

此属性可以接受的值为 `exclusive` 和 `shared`。

全局: hostid

区域可以模仿 32 位主机标识符以简化系统整合。缺省情况下，区域的 `hostid` 属性是空的，这意味着区域不模拟主机标识符。区域主机标识符必须是 0 到 FFFFFFFE 之间的十六进制值。0x 或 0X 前缀是可选的。大写和小写十六进制数字都可以接受。

全局: fs-allowed

可以在区域中挂载的其他文件系统的逗号分隔列表；例如，`ufs,pcfs`。缺省情况下，只能挂载 `hsfs(7FS)` 和网络文件系统。

此属性不适用于通过 `add fs` 或 `add dataset` 挂载到区域中的文件系统。

注意—如果允许挂载非缺省文件系统，将使得区域管理员能够使用伪造的文件系统映像来危害系统安全，因此不支持此类挂载。

全局: `file-mac-profile`

定义文件系统的哪些部分不使用只读策略，即，允许区域写入文件系统的哪些部分。

该属性当前支持 4 个值: `none`、`strict`、`fixed-configuration` 和 `flexible-configuration`。

`none` 使该区域与常规读取/写入区域完全相同。`strict` 禁止使用只读策略以外的策略。`fixed-configuration` 允许区域写入 `/var` 中及其子目录下的文件，但包含配置文件的目录除外：

```
/var/ld
/var/lib/portsnap
/var/pkg
/var/spool/cron,
/var/spool/portsnap
/var/svc/manifest
/var/svc/profiles
```

`flexible-configuration` 与 `fixed-configuration` 相当，但还允许写入 `/etc` 中的文件。

`fs`: `dir`、`special`、`raw`、`type`、`options`

确定文件系统的挂载方式、挂载位置等信息时所需的值。请参见 [mount\(1M\)](#)、[mount\(2\)](#)、[fsck\(1M\)](#) 和 [vfstab\(4\)](#)。

`net`: `address`、`allowed-address`、`configure-allowed-address`、`physical`、`defrouter`

`net` 资源表示为一个区域分配的物理网络资源。在分配之前，该资源必须在全局区域中存在。

网络接口的网络地址和物理接口名称。网络地址为以下地址之一：

- 有效的 IPv4 地址，后面可以跟有可选的 `/` 和前缀长度；
- 有效的 IPv6 地址，必须后接 `/` 和前缀长度；
- 可解析为 IPv4 地址的主机名。

请注意，不支持解析为 IPv6 地址的主机名称。

物理接口名称即网络接口名称。

用来指定可选的缺省路由器值的方式与指定网络地址的方式类似，但它不能后跟 `/`（斜杠）和网络前缀长度。要正确地使用 `defrouter` 功能，使用该属性的区域所在的子网必须不同于全局区域所在的子网。此外，使用不同的 `defrouter` 设置的每个区域（或区域集）必须位于不同的子网上。

可将区域配置为专用 IP 区域或共享 IP 区域。对于共享 IP 区域，必须同时设置 `physical` 和 `address` 属性；是否设置缺省路由器是可选的。在引导非全局区域之前，必须在全局区域中激活在 `physical` 属性中指定的接口。不过，如果全局区域不使用该接口，应在全局区域中将其配置为 `down`，并应在此处为该接口指定缺省路由器。不能为共享 IP 区域设置 `allowed-address` 属性。

对于独占 IP 区域，必须设置 `physical` 属性，不能设置 `address` 属性。此外，还可以通过指定 `allowed-address` 属性来限制独占 IP 区域可以使用的 IP 地址集。如果未指定 `allowed-address`，则独占 IP 区域可以使用 `net` 资源的关联 `physical` 接口上的任何 IP 地址。而当指定了 `allowed-address` 时，独占 IP 区域将无法使用不在 `physical` 地址的 `allowed-address` 列表中的 IP 地址。如果 `configure-allowed-address` 设置为 `true`，在安装后初始引导时总是对接口配置 `allowed-address`。如果它设置为 `false`，在区域引导时将不会配置 `allowed-address`。缺省情况下，当指定 `allowed-address` 时，`configure-allowed-address` 将设置为 `true`。此外，如果填充了 `allowed-address` 列表，还可以选择指定 `defrouter` 属性。但是，如果指定了 `defrouter` 值且 `configure-allowed-address` 设置为 `false`，则将忽略 `defrouter` 值，并显示相应的警告消息。为 `physical` 属性指定的接口必须没有在全局区域中使用。如果通过 `zonecfg` 指定了 `allowed-address` 和缺省路由器，则这些属性将在通过 `ipadm(1M)` 在非全局独占 IP 区域中启用接口时（通常是在区域引导期间）应用于接口。非全局独占 IP 区域将无法向该接口应用任何其他地址，也无法传输具有指定 IP 版本的不同源地址的数据包。通过 `zonecfg` 设置的缺省路由器无法使用 `-p` 标志与 `route(1M)` 从非全局独占 IP 区域内永久删除。

请注意，不能在多个独占 IP 区域间共享单个数据链路。

`anet`: `linkname`、`lower-link`、`allowed-address`、`configure-allowed-address`、`defrouter`、`mac-address`
`anet` 资源表示为独占 IP 区域自动创建的网络资源。`zonecfg` 使用缺省 `SYSdefault` 模板创建区域时，将在区域配置中自动包括具有以下属性的 `anet` 资源：

```
linkname=net0
lower-link=auto
mac-address=random
link-protection=mac-nospoof
```

当此类区域引导时，将自动确定 `lower-link`（请参见下文）并通过该链路为区域自动创建临时 VNIC。当区域停止时，将删除 `lower-link` 和 VNIC。

下面介绍了支持的属性。所有这些属性都是可选的。只允许全局区域修改自动创建的 VNIC 或其属性。如果 `zonecfg` 中设置的属性不能在创建 VNIC 时分配给 VNIC，则该区域将无法引导。

`linkname`

为自动创建的 VNIC 指定名称。缺省情况下，该属性将自动设置为 `netN` 格式（其中 `N` 为非负整数）的第一个可用名称（对于区域）。例如：`net0`、`net1` 等。`info` 子命令会显示自动选择的 `linkname`。

多个区域（包括全局区域）可同时使用同名链路。

`lower-link`

指定用来创建 VNIC 的链路。该属性的缺省值为 `auto`。在添加 `anet` 资源时，管理员可以显式指定一个值。链路可以是接受为 `dladm create-vnic-l` 选项的参数 的任何链路（请参见 `dladm(1M)`）。如果该属性设置为 `linkname`（而非 `auto`）且

该链路不存在，则区域将无法引导。如果设置为 `auto`，每次区域引导时，`zoneadmd(1M)` 守护进程都会自动选择用来创建 VNIC 的链路。将会使用以下试探性操作选择链路：

1. 第一个物理数据链路的链路状态是否为 `up`。
2. 如果不存在 `up` 状态，则使用名为 `net0` 的数据链路。

如果以上两点都不满足，区域将无法引导。

`allowed-address`

请参见 `net` 资源中对独占 IP 区域 `allowed-address` 属性的描述。

`configure-allowed-address`

请参见 `net` 资源中对独占 IP 区域 `configure-allowed-address` 属性的描述。

`defrouter`

请参见 `net` 资源中对独占 IP 区域 `defrouter` 属性的描述。

`mac-address`

根据指定值或关键字设置 VNIC 的 MAC 地址。如果该值不是关键字，则将其解释为单点传送 MAC 地址。支持的关键字包括：

- `factory`：为 VNIC 分配出厂 MAC 地址。当请求出厂 MAC 地址时，可以使用 `mac-slot` 属性指定 MAC 地址插槽标识符。否则，将使用下一个可用的出厂 MAC 地址。
- `random`：为 VNIC 分配随机 MAC 地址。使用 `mac-prefix` 属性指定一个前缀。否则，将使用缺省前缀，该前缀包含有效 IEEE OUI，并设置了本地位。这是缺省值。
- `auto`：首先尝试使用出厂 MAC 地址。如果该地址不可用，则分配随机 MAC 地址。

如果选择了随机 MAC 地址，则将在区域引导和区域分离/连接后保留生成的地址。这将允许区域通过维护稳定的客户机 ID 来保留它们的 DHCP 租用，还会充分利用具有稳定 MAC 地址的其他优势。

`mac-prefix`

指定 MAC 地址前缀（如果请求随机 MAC 地址）。否则，将忽略该属性。

`mac-slot`

指定 MAC 地址插槽标识符（如果请求出厂 MAC 地址）。否则，将忽略该属性。

`mtu`

VNIC 的最大传输单元（以字节为单位）。请参见 `dladm(1M)` 中的 `mtu` 属性。

`maxbw`

指定 VNIC 的全双工带宽。请参见 `dladm(1M)` 中的 `maxbw` 属性。缺省情况下，VNIC 将使用在 `lower-link` 上设置的 `maxbw`，如果未进行设置，则不存在任何带宽限制。

priority

指定 VNIC 的相对优先级。有关支持的值和缺省值，请参见 [dladm\(1M\)](#) 中的 **priority** 属性。

vlan-id

对该 VNIC 启用 VLAN 标记并指定 VLAN 标记的 id。没有指定缺省值，这意味着如果未设置该属性，VNIC 将不会参与任何 VLAN。

rxfanout

指定接收端扇出线程数。请参见 [dladm\(1M\)](#) 中的描述。

rxrings

指定 VNIC 的接收环。有关支持的值和缺省值，请参见 [dladm\(1M\)](#) 中的 **rxrings** 属性。

txrings

指定 VNIC 的传送环。有关支持的值和缺省值，请参见 [dladm\(1M\)](#) 中的 **txrings** 属性。

link-protection

启用一个或多个链路保护类型（使用逗号分隔值）。有关支持的值，请参见 [dladm\(1M\)](#) 中的 **protection** 属性。它的缺省值为 **mac-nospoof**。

请注意，向该属性添加 **ip-nospoof** 不会有任何效果，除非同时设置了 **allowed-address**。设置 **allowed-address** 会将 **ip-nospoof** 隐式添加到 **link-protection** 集，而清除 **allowed-address** 会将其删除。

allowed-dhcp-cids

设置该属性将对 VNIC 启用 **dhcp-nospoof**。有关详细信息，请参见 [dladm\(1M\)](#)。

device: match

要匹配的设备名称。这可以是要匹配的模式，也可以是绝对路径名。请注意，设备资源和别名数据集可能会在 **/dev/zvol** 中存在名称空间冲突。请参见 [dev\(7FS\)](#)。

注—一般情况下，向区域中添加设备可能会危害系统的安全性；请参见“附注”部分。

rctl: name \ value

资源控制的名称和 *priv/limit/action* 三要素。请参见 [prctl\(1\)](#) 和 [rctladm\(1M\)](#)。设置 **rctl** 值的首选方式是使用与特定 **rctl** 关联的全局属性名称。

attr: name \ type \ value

通用属性的名称、类型和值。**type** 必须是 **int**、**uint**、**boolean** 或 **string** 中的一个，**value** 必须属于该类型。**uint** 的意思是不带符号的整数，即非负整数。

dataset: name \ alias

要从区域内访问的 ZFS 数据集的名称。请参见 [zfs\(1M\)](#)。为每个数据集设置了别名，以便它在区域中显示为虚拟 ZFS 池。**alias** 是该虚拟池的名称。有关应用于 ZFS 池名称，并进而还应用于数据集别名值的名称限制，请参见 [zpool\(1M\)](#)。别名

rpool 是从区域的 rpool 数据集中保留的。请注意，别名数据集和设备资源可能会在 `/dev/zvol` 中存在名称空间冲突。请参见 [dev\(7FS\)](#)。

全局: `cpu-shares`

要分配给此区域的公平份额调度器 (Fair Share Scheduler, FSS) 份额的数量。此属性与 `dedicated-cpu` 资源不兼容。此属性是设置 `zone.cpu-shares rctl` 的首选方式。

全局: `max-lwps`

此区域可同时使用的最大 LWP 数。此属性是设置 `zone.max-lwps rctl` 的首选方式。

全局: `max-msg-ids`

此区域允许的最大消息队列 ID 数。此属性是设置 `zone.max-msg-ids rctl` 的首选方式。

全局: `max-processes`

此区域可同时使用的最大进程表槽数。此属性是设置 `zone.max-processes rctl` 的首选方式。除非已显式设置了 `max-lwps` 属性，否则设置此属性会隐式地将 `max-lwps` 属性的值设置为进程槽数的 10 倍。

全局: `max-sem-ids`

此区域允许的最大信号量 ID 数。此属性是设置 `zone.max-sem-ids rctl` 的首选方式。

全局: `max-shm-ids`

此区域允许的最大共享内存 ID 数。此属性是设置 `zone.max-shm-ids rctl` 的首选方式。

全局: `max-shm-memory`

此区域允许的最大共享内存数量。此属性是设置 `zone.max-shm-memory rctl` 的首选方式。可以向此数量值应用一个比例 (K、M、G、T)，例如，1M 代表 100 万字节。

全局: `scheduling-class`

指定用于在某个区域中运行的进程的调度类。未指定此属性时，将按如下方式建立调度类：

- 如果设置了 `cpu-shares` 属性或等效的 `rctl`，将使用调度类 FSS。
- 如果既未设置 `cpu-shares` 也未设置等效的 `rctl`，且区域的 `pool` 属性引用了一个具有缺省调度类的池，则将使用该类。
- 在任何其他情况下，将使用系统缺省调度类。

`dedicated-cpu: ncpu、importance`

应分配给该区域供其独占使用的 CPU 数量。该区域在引导时将创建一个池和处理器集。有关资源池的更多信息，请参见 [pooladm\(1M\)](#) 和 [poolcfg\(1M\)](#)。`ncpu` 属性可以指定单个值，也可以指定一个处理器数量范围（例如 1-4）。`importance` 属性是可选的；如果设置了该属性，则它指定 `pset.importance` 值供 [poolld\(1M\)](#) 使用。如果使用了此资源，则在该区域引导时必须有足够的空闲处理器分配给该区域，否则

该区域将不会引导。分配给该区域的处理器将不可供全局区域或其他区域使用。此资源与 pool 和 cpu-shares 属性都不兼容。只能将此资源的单个实例添加到区域中。

capped-memory: physical、swap、locked

该区域可以使用的内存量上限。可以向其中的每个数量值应用一个比例（K、M、G、T），例如，1M代表100万字节。其中每个属性都是可选的，但在添加此资源时必须至少设置一个属性。只能将此资源的单个实例添加到区域中。physical 属性设置该区域的 max-rss。这将由在全局区域中运行的 rcapd(1M) 强制实施。swap 属性是设置 zone.max-swap rctl 的首选方式。locked 属性是设置 zone.max-locked-memory rctl 的首选方式。

capped-cpu: ncpus

设置可供某个区域使用的 CPU 时间量的限制。使用的单元将换算为可供区域中的所有用户线程使用的对单个 CPU 的百分比，以小数表示（例如，.75），或以混合数字表示（整数和小数，例如，1.25）。ncpu 值为 1 表示一个 CPU 的 100%，值为 1.25 表示 125%，.75 表示 75%，依此类推。当设置有上限的区域中的项目有其自己的上限时，优先采用最小值。

capped-cpu 属性是 zone.cpu-cap 资源控制的别名，与 zone.cpu-cap 资源控制相关。请参见 resource_controls(5)。

admin: user、auths

将区域管理授权委托给指定的用户或角色。该用户必须对应于一个有效的本地帐户。auths 的允许值包括：

login

允许在经过验证的情况下 zlogin(1) 到该区域。

manage

允许对所配置的区域进行正常管理。

copyfrom

允许将指定的区域作为用来克隆新区域的来源。

下表汇总了资源、属性名称和类型：

resource	property-name	type
(global)	zonename	simple
(global)	zonepath	simple
(global)	autoboot	simple
(global)	bootargs	simple
(global)	pool	simple
(global)	limitpriv	simple
(global)	brand	simple
(global)	ip-type	simple
(global)	hostid	simple
(global)	cpu-shares	simple
(global)	max-lwps	simple

(global)	max-msg-ids	simple
(global)	max-sem-ids	simple
(global)	max-shm-ids	simple
(global)	max-shm-memory	simple
(global)	scheduling-class	simple
(global)	fs-allowed	list of simple
fs	dir	simple
	special	simple
	raw	simple
	type	simple
	options	list of simple
net	address	simple
	allowed-address	list of simple
	configure-allowed-address	simple
	defrouter	list of simple
	physical	simple
anet	linkname	simple
	lower-link	simple
	allowed-address	list of simple
	configure-allowed-address	simple
	defrouter	list of simple
	mac-address	simple
	mac-slot	simple
	mac-prefix	simple
	mtu	simple
	maxbw	simple
	priority	simple
	vlan-id	simple
	rxfanout	simple
	rxrings	simple
	txrings	simple
	link-protection	list of simple
	allowed-dhcp-cids	list of simple
device	match	simple
rctl	name	simple
	value	list of complex
attr	name	simple
	type	simple
	value	simple
dataset	name	simple
	alias	simple
dedicated-cpu	ncpus	simple or range
	importance	simple
capped-memory	physical	simple with scale
	swap	simple with scale
	locked	simple with scale

capped-cpu	ncpus	simple
admin	user	simple
	auths	list of simple

为进一步指定各个事项 ("rctl" 资源类型的复杂属性 "value" 的细目)，它包含了三个名称/值对，名称为 `priv`、`limit` 和 `action`，其中每个名称都接受一个简单值。"attr" 资源的 "name" 属性在语法方面的限制类似于区域名称，但不完全相同：它必须以字母数字开头，可以包含字母数字以及连字符 (-)、下划线 () 和点 (.) 字符。以 "zone" 开头的属性名称保留供系统使用。最后，"autoboot" 全局属性的值必须为 `true` 或 `false`。

使用内核统计信息
来监视 CPU 上限

系统使用内核统计信息 (`kstat(3KSTAT)`) 模块 `caps` 为所有设有上限的项目和区域维护该信息。您可以通过读取内核统计信息 (`kstat(3KSTAT)`) (指定 `caps` 作为 `kstat` 模块名称) 来访问该信息。以下命令显示所有活动 CPU 上限的内核统计信息：

```
# kstat caps::'/cpucaps/'
```

在某个区域中运行的 `kstat(1M)` 命令只显示与该区域以及该区域中的项目相关的 CUP 上限。请参见“示例”。

下面是可与 `kstat(1M)` 一起使用的与上限相关的参数：

- `caps`
 `kstat` 模块。
- `project_caps` 或 `zone_caps`
 `kstat` 类，可与 `kstat -c` 选项一起使用。
- `cpucaps_project_id` 或 `cpucaps_zone_id`
 `kstat` 名称，可与 `kstat -n` 选项一起使用。*id* 是项目或区域的标识符。

响应请求所有 CPU 上限统计信息的 `kstat(1M)` 命令时将显示以下字段。

- `module`
 在 `kstat` 的此使用情况中，此字段的值将是 `caps`。
- `name`
 如上所述，`cpucaps_project_id` 或 `cpucaps_zone_id`
- `above_sec`
 在上限之上度过的总时间（以秒为单位）。
- `below_sec`
 在上限之下度过的总时间（以秒为单位）。
- `maxusage`
 观察到的最大 CPU 使用量。
- `nwait`
 在上限等待队列中的线程数。

usage

属于某个设有上限的项目或区域的所有线程当前使用的 CPU 总量，换算为对单个 CPU 的百分比。

value

上限值，换算为对单个 CPU 的百分比。

zonename

为其显示统计信息的区域的名称。

有 `kstat` 命令的输出的样例，请参见“示例”部分。

选项

所支持的选项如下：

-f *command_file*

指定 `zonecfg` 命令文件的名称。*command_file* 是包含 `zonecfg` 子命令的文本文件，每个子命令占一行。

-z *zonename*

指定区域的名称。区域名称区分大小写。区域名称必须以字母数字字符开头，并可包含字母数字字符、下划线(_)、连字符(-)和点(.)。名称 `global` 和所有以 `SYS` 开头的名称均保留，不能使用。

子命令

可以使用 `add` 和 `select` 子命令来选择特定的资源，此时作用域将更改为该资源。`end` 和 `cancel` 子命令用于结束资源指定，此时作用域将恢复为全局。某些子命令（例如 `add`、`remove` 和 `set`）在每个作用域中有不同的语义。

`zonecfg` 支持以分号分隔的子命令列表。例如：

```
# zonecfg -z myzone "add net; set physical=myvnic; end"
```

可导致破坏性操作或导致工作丢失的子命令具有一个强制执行操作的 `-F` 选项。当输入来自终端设备时，如果在给定此类命令时未指定 `-F` 选项，系统会在适当的时候提示用户；其他情况下，如果在给定此类命令时未指定 `-F` 选项，将不允许操作，并向标准错误写入一条诊断消息。

支持以下子命令：

add *resource-type*（全局范围）

add *property-name property-value*（资源范围）

在全局范围中，开始给定资源类型的指定。将范围更改为此资源类型。

在资源范围中，添加具有给定名称和给定值的属性。用于属性值的语法随属性类型不同而不同。一般情况下，它是一个简单值或者是用方括号括起来并以逗号分隔的简单值列表 (`[foo,bar,baz]`)。请参见“属性”部分。

cancel

结束资源指定并将作用域重置为全局。放弃任何部分指定的资源。`cancel` 仅在资源范围内适用。

clear *property-name*

清除属性的值。

commit

将当前配置从内存提交到稳定存储器。必须提交配置以供 `zoneadm` 使用。在提交内存中的配置之前，可以使用 `revert` 子命令删除更改。在 `zonecfg` 会话完成时，会自动尝试 `commit` 操作。因为配置必须是正确的才能提交，因此该操作将自动执行验证。

create [-F] [-a *path* | -b | -t *template*]

为指定的区域创建内存中配置。可以使用 `create` 开始配置一个新区域。有关将此保存到稳定存储器的信息，请参见 `commit`。

如果要覆盖现有的配置，请指定 `-F` 选项以强制执行操作。指定 `-t template` 选项可创建一个与 *template* 完全相同的配置，其中 *template* 是所配置的某个区域的名称。

`create` 使用缺省模板 `SYSdefault`。可以使用 `svc:/system/zones:default` 服务的 `default_template` SMF 属性在系统范围内更改缺省模板。管理员可以使用 `-t`（使用特定模板）或 `-b`（使用空白模板）覆盖该区域的缺省值。

使用 `-a path` 选项可方便地在新主机上匹配一个分离区域。*path* 参数是已经移动到此新主机上的分离区域的 `zonepath` 位置。配置了分离区域后，应使用 `zoneadm attach` 命令安装它（请参见 [zoneadm\(1M\)](#)）。对新区域的所有验证都发生在 `attach` 进程期间，而不是在区域配置期间。

使用 `-b` 选项可以创建一个空白配置。如果未指定参数，`create` 将应用 Oracle Sun 缺省设置。

delete [-F]

从内存和稳定存储器中删除指定的配置。此操作是即时的，不需要提交。删除的配置无法恢复。

使用 `-F` 选项强制执行操作：

end

结束资源指定。此子命令仅在资源范围内适用。`zonecfg` 进行检查以确保完整地指定了当前资源。如果完整地指定了当前资源，则会将其添加到内存中配置（有关将此配置保存到稳定存储器的信息，请参见 `commit`），作用域将恢复为全局。如果指定不完整，它将发出相应的错误消息。

export [-f *output-file*]

将配置输出至标准输出。使用 `-f` 选项可将配置输出到 *output-file*。此选项以适合在命令文件中使用的格式生成输出。

help [usage] [subcommand] [syntax] [command-name]

输出常规帮助或有关给定主题的帮助。

```
info zonename | zonepath | autoboot | brand | pool | limitpriv
info [resource-type [property-name=property-value ]*]
```

显示有关当前配置的信息。如果指定了 *resource-type*，将仅显示关于相关类型的资源的信息。如果指定了任何 *property-name* 值对，将仅显示符合指定条件的资源的相关信息。在资源范围中，所有参数都将被忽略，*info* 将仅显示当前正在添加或修改的资源的相关信息。

remove resource-type{property-name=property-value} (全局范围)

在全局范围中，删除指定的资源。[] 语法表示在方括号内有 0 项或多项内容。如果您希望仅删除资源的单个实例，必须为要唯一标识的资源指定足够的属性名称-值对。如果没有指定属性名称-值对，则会删除所有实例。如果指定了多个名称-值对，除非使用了 *-F* 选项，否则将需要确认。

select resource-type {property-name=property-value}

选择与给定的 *property-name* *property-value* 名称/值对条件匹配的给定类型的资源进行修改。此子命令仅在全局范围内适用。将范围更改为此资源类型。{} 语法表示在花括号内有 1 项或多项内容。必须为要唯一标识的资源指定足够的 *property-name* *property-value* 对。

set property-name=property-value

将给定属性名称设置为给定值。某些属性（例如，*zonename* 和 *zonepath*）是全局的，另外一些属性是特定于资源的。此子命令同时适用于全局范围和资源范围。

verify [-v]

检验当前配置是否正确：

- 所有资源都指定了其所有必需属性。
- 指定了 *zonepath*。

如果指定了 *-v* 选项，当设备资源中指定的设备可能存在冲突并隐藏别名数据集内创建的 ZFS 卷时，将会发出警告。请参见 [dev\(7FS\)](#)。

revert [-F]

将配置恢复到上次提交时的状态。可使用 *-F* 选项强制执行操作。

exit [-F]

退出 *zonecfg* 会话。如果需要，会自动尝试 *commit*。还可以使用 EOF 字符来退出 *zonecfg*。可使用 *-F* 选项强制执行操作。

示例

示例 1 创建新区域的环境

在下面的示例中，*zonecfg* 创建新区域的环境。*/usr/local* 从全局区域回送挂载到 */opt/local*。*/opt/sfw* 从全局区域回送挂载，*nxge0* 上的 VNIC 添加到具有三个 IP 地址的区域，使用 *rctl* 资源类型设置对区域的公平份额调度程序 (*fair-share scheduler*, *FSS*) CPU 份额数的限制。该示例还说明了如何选择用于进行修改的给定资源，本例中的方法是选择由 *zonecfg* 自动创建的 *anet* 资源。

```
example# zonecfg -z myzone3
my-zone3: No such zone configured
Use 'create' to begin configuring a new zone.
zonecfg:myzone3> create
zonecfg:myzone3> set zonepath=/export/home/my-zone3
```

示例1 创建新区域的环境 (续)

```

zonecfg:myzone3> set autoboot=true
zonecfg:myzone3> add fs
zonecfg:myzone3:fs> set dir=/usr/local
zonecfg:myzone3:fs> set special=/opt/local
zonecfg:myzone3:fs> set type=lofs
zonecfg:myzone3:fs> add options [ro,nodevices]
zonecfg:myzone3:fs> end
zonecfg:myzone3> add fs
zonecfg:myzone3:fs> set dir=/mnt
zonecfg:myzone3:fs> set special=/dev/dsk/c0t0d0s7
zonecfg:myzone3:fs> set raw=/dev/rdisk/c0t0d0s7
zonecfg:myzone3:fs> set type=ufs
zonecfg:myzone3:fs> end
zonecfg:myzone3> add fs
zonecfg:myzone3:fs> set dir=/opt/sfw
zonecfg:myzone3:fs> set special=/opt/sfw
zonecfg:myzone3:fs> set type=lofs
zonecfg:myzone3:fs> add options [ro,nodevices]
zonecfg:myzone3:fs> end
zonecfg:myzone3> select anet linkname=net0
zonecfg:myzone3:anet> set lower-link=nxge0
zonecfg:myzone3:anet> set allowed-address="192.168.0.1/24,192.168.1.2/\
24,192.168.2.3/24"
zonecfg:myzone3:anet> end
zonecfg:my-zone3> set cpu-shares=5
zonecfg:my-zone3> add capped-memory
zonecfg:my-zone3:capped-memory> set physical=50m
zonecfg:my-zone3:capped-memory> set swap=100m
zonecfg:my-zone3:capped-memory> end
zonecfg:myzone3> exit

```

示例2 创建独占IP区域

以下示例将创建一个分配了名为 net0 的 VNIC 的区域。用来创建 VNIC 的链路是自动确定的。IP 地址和路由是在新区域内部使用 [ipadm\(1M\)](#) 配置的。

```

example# zonecfg -z excl
zonecfg:excl> create
zonecfg:excl> set zonepath=/export/zones/excl
zonecfg:excl> exit

```

示例3 创建共享IP区域

以下示例创建的区域将与全局区域共享 IP 栈，并被分配一个 IP 地址和缺省路由器。

```

example# zonecfg -b -z shared
zonecfg:shared> create
zonecfg:shared> set zonepath=/export/zones/shared

```

示例3 创建共享 IP 区域 (续)

```

zonecfg:shared> set ip-type=shared
zonecfg:shared> add net
zonecfg:shared:net> set physical=nge0
zonecfg:shared:net> set address=192.168.0.3/24
zonecfg:shared:net> set defrouter=192.168.0.1
zonecfg:shared:net> end
zonecfg:shared> exit

```

示例4 将区域与资源池相关联

以下示例展示了如何将现有区域与现有的资源池相关联：

```

example# zonecfg -z myzone
zonecfg:myzone> set pool=mypool
zonecfg:myzone> exit

```

有关资源池的更多信息，请参见 [pooladm\(1M\)](#) 和 [poolcfg\(1M\)](#)。

示例5 更改区域的名称

以下示例展示了如何更改现有区域的名称：

```

example# zonecfg -z myzone
zonecfg:myzone> set zonename=myzone2
zonecfg:myzone2> exit

```

示例6 更改区域的特权集

以下示例展示了如何更改现有区域的特权集，该区域的进程在区域下次引导时将受限于该特权集。在这个特定的案例中，特权集将是区域正常情况下拥有的标准安全特权集，加上更改系统日期和时间的特权。

```

example# zonecfg -z myzone
zonecfg:myzone> set limitpriv="default,sys_time"
zonecfg:myzone2> exit

```

示例7 设置全局区域的 zone.cpu-shares 属性

以下命令可设置全局区域的 zone.cpu-shares 属性：

```

example# zonecfg -z global
zonecfg:global> set cpu-shares=5
zonecfg:global> exit

```

示例8 使用模式匹配

以下命令展示了 zonecfg 对模式匹配的支持。在区域 flexlm 中，输入：

```

zonecfg:flexlm> add device
zonecfg:flexlm:device> set match="/dev/cua/a00[2-5]"
zonecfg:flexlm:device> end

```

示例 8 使用模式匹配 (续)

在全局区域中，输入：

```
global# ls /dev/cua
a      a000  a001  a002  a003  a004  a005  a006  a007  b
```

在区域 flexlm 中，输入：

```
flexlm# ls /dev/cua
a002  a003  a004  a005
```

示例 9 将区域的上限设置为三个 CPU

以下序列使用 zonecfg 命令将区域的 CPU 上限设置为三个 CPU。

```
zonecfg:myzone> add capped-cpu
zonecfg:myzone>capped-cpu> set ncpus=3
zonecfg:myzone>capped-cpu>capped-cpu> end
```

上面的序列使用了 capped-cpu 属性，它等效于使用了 zone.cpu-cap 资源控制的以下序列。

```
zonecfg:myzone> add rctl
zonecfg:myzone:rctl> set name=zone.cpu-cap
zonecfg:myzone:rctl> add value (priv=privileged,limit=300,action=none)
zonecfg:myzone:rctl> end
```

示例 10 使用 kstat 监视 CPU 上限

以下命令显示有关所有 CPU 上限的信息。

```
# kstat -n /cpucaps/
module: caps                               instance: 0
name:    cpucaps_project_0                class:    project_caps
         above_sec                          0
         below_sec                         2157
         crtime                             821.048183159
         maxusage                           2
         nwait                              0
         snaptime                          235885.637253027
         usage                              0
         value                             18446743151372347932
         zonename                          global

module: caps                               instance: 0
name:    cpucaps_project_1                class:    project_caps
         above_sec                          0
         below_sec                         0
         crtime                             225339.192787265
         maxusage                           5
```

示例 10 使用 kstat 监视 CPU 上限 (续)

```

nwait      0
snaptime   235885.637591677
usage      5
value      18446743151372347932
zonename    global

module: caps      instance: 0
name:  cpucaps_project_201      class:  project_caps
      above_sec      0
      below_sec      235105
      crtime         780.37961782
      maxusage       100
      nwait          0
      snaptime       235885.637789687
      usage          43
      value          100
      zonename       global

module: caps      instance: 0
name:  cpucaps_project_202      class:  project_caps
      above_sec      0
      below_sec      235094
      crtime         791.72983782
      maxusage       100
      nwait          0
      snaptime       235885.637967512
      usage          48
      value          100
      zonename       global

module: caps      instance: 0
name:  cpucaps_project_203      class:  project_caps
      above_sec      0
      below_sec      235034
      crtime         852.104401481
      maxusage       75
      nwait          0
      snaptime       235885.638144304
      usage          47
      value          100
      zonename       global

module: caps      instance: 0
name:  cpucaps_project_86710     class:  project_caps
      above_sec      22
      below_sec      235166
```

示例 10 使用 kstat 监视 CPU 上限 (续)

```

        crtime                698.441717859
        maxusage              101
        nwait                 0
        snaptime              235885.638319871
        usage                 54
        value                 100
        zonename              global

module: caps                  instance: 0
name:   cpucaps_zone_0       class:   zone_caps
        above_sec            100733
        below_sec            134332
        crtime               821.048177123
        maxusage             207
        nwait                2
        snaptime             235885.638497731
        usage                199
        value                200
        zonename             global

module: caps                  instance: 1
name:   cpucaps_project_0    class:   project_caps
        above_sec            0
        below_sec            0
        crtime               225360.256448422
        maxusage             7
        nwait                0
        snaptime             235885.638714404
        usage                7
        value                18446743151372347932
        zonename             test_001

module: caps                  instance: 1
name:   cpucaps_zone_1       class:   zone_caps
        above_sec            2
        below_sec            10524
        crtime               225360.256440278
        maxusage             106
        nwait                0
        snaptime             235885.638896443
        usage                7
        value                100
        zonename             test_001
```

示例 11 显示特定区域或项目的 CPU 上限

使用 `kstat -c` 和 `-i` 选项，可以显示特定区域或项目的 CPU 上限，如下所示。第一个命令生成针对某个特定项目的显示，第二个命令生成针对区域 1 中同一个项目的显示。

```
# kstat -c project_caps
```

```
# kstat -c project_caps -i 1
```

示例 12 委托区域管理权限

以下示例展示了如何将当前区域的管理权限分配给某个角色。

```
example# zonecfg -z myzone
zonecfg:myzone> add admin
zonecfg:myzone:admin> set user=zadmin
zonecfg:myzone:admin> set auths=login,manage
zonecfg:myzone:admin> end
zonecfg:myzone> commit
```

这些命令的执行结果将是 RBAC `user_attr(4)` 数据库中的更新条目，类似于以下条目：

```
zadmin:::type=role;\
auths=solaris.zone.login/myzone,solaris.zone.manage/myzone;profiles=\
Zone Management
```

示例 13 创建具有非缺省属性的独占 IP 区域

以下示例创建的区域具有通过 `mylink0` 自动创建的 VNIC，并指定了给定的 MAC 地址、100 Mbps 的最大带宽、高优先级、RX 端的专用硬件环、TX 端的非专用硬件环（即，基于软件）以及 VLAN id 2。

```
example# zonecfg -b -z excl
excl: No such zone configured
Use 'create' to begin configuring a new zone
zonecfg:excl> create
zonecfg:excl> set zonepath=/export/zones/excl
zonecfg:excl> add anet
zonecfg:excl:anet> set linkname=mynic0
zonecfg:excl:anet> set lower-link=mylink0
zonecfg:excl:anet> set mac-address=8:0:20:fe:4e:b8
zonecfg:excl:anet> set maxbw=100M
zonecfg:excl:anet> set priority=high
zonecfg:excl:anet> set vlan-id=2
zonecfg:excl:anet> set rxrings=hw
zonecfg:excl:anet> set txrings=sw
zonecfg:excl:anet> end
zonecfg:excl> exit
```


示例14 创建只读区域

以下示例将创建一个可防止区域对其根文件系统进行修改的新区域。由于应用了 `fixed-configuration` 配置文件，`/var` 中的文件是可写的。

```
example# zonecfg -z rozone
rozone: No such zone configured
Use 'create' to begin configuring a new zone
zonecfg:rozone> create
zonecfg:rozone> set brand=solaris
zonecfg:rozone> set zonepath=/export/zones/rozone
zonecfg:rozone> set autoboot=true
zonecfg:rozone> set file-mac-profile=fixed-configuration
zonecfg:rozone> set ip-type=exclusive
zonecfg:rozone> add net
zonecfg:rozone:net> set physical=vnic0
zonecfg:rozone:net> end
zonecfg:rozone> exit
```

退出状态

将返回以下退出值：

- 0 成功完成。
- 1 出现错误。
- 2 用法无效。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/zones
接口稳定性	Volatile（可变）

另请参见

[ppriv\(1\)](#)、[prctl\(1\)](#)、[zlogin\(1\)](#)、[dladm\(1M\)](#)、[ipadm\(1M\)](#)、[kstat\(1M\)](#)、[mount\(1M\)](#)、[pooladm\(1M\)](#)
《[Oracle Solaris 管理：Oracle Solaris Zones、Oracle Solaris 10 Zones 和资源管理](#)》

附注

`zonecfg` 使用的所有字符都必须采用 US-ASCII 编码。

一般情况下，向区域中添加设备可能会使区域对系统的安全性和稳定性产生不利影响，因为不是所有设备都已经过审核并确认可以在区域内安全地使用。

可以将使用 `sd` 或 `ssd` 目标驱动程序存储设备（例如，这可以使用 `prtconf -D /dev/dsk/c2t40d3` 来检查）安全地委托给区域。这将使得区域管理员可以为此类设备设置标签和分区。

只委托一个分片（例如，`match=/dev/dsk/c2t40d3s0`）也将使得区域管理员能够为包含该分片的整个磁盘设置标签和分区。要阻止这种情况发生，请通过以下方式取消从区域运行 `format` 的功能：

```
set limitpriv=default,-storage_config
```

通过将 `raw_storage` 特权添加到区域的 `limitpriv` 中，可以启用对存储设备的原始访问。这是不安全的，因为它会允许区域进程执行原始 SCSI 命令（请参见 [uscsi\(7I\)](#)）。

引用名	zonep2vchk – check a global zone's configuration for physical to virtual migration into non-global zone						
用法概要	<pre> zonep2vchk -V zonep2vchk [-T release] -c zonep2vchk [-T release] [-P] [-b] [-s path[,path...]] [-S file] [-r {time}(h m s)] [-x] [-e execname[,execname...]] [-E file] </pre>						
描述	The zonep2vchk utility is used to evaluate a global zone's configuration before the process of physical-to-virtual (p2v) migration into a non-global zone. The p2v process involves archiving a global zone (source), and then installing a non-global zone (target) using that archive. See the <code>install -a</code> documentation in the solaris(5) and solaris10(5) man pages. zonep2vchk serves two functions. First, it can be used to report issues on the source which might prevent a successful p2v migration. Second, it can output a template <code>zonecfg</code>, which can be used to assist in configuring the non-global zone target. zonep2vchk can be executed on a Solaris 10 or later global zone. To execute on Solaris 10, copy the zonep2vchk utility to the Solaris 10 source global zone.						
安全	The zonep2vchk utility must be run with an effective user id of zero. It interrogates the configuration state of a variety of Solaris subsystems.						
选项	The following options are supported: -V Display the command version and exit. -T release Specify the target release. The defaults are: <table> <tr> <th>Global Zone</th><th>Default Target</th></tr> <tr> <td>Solaris 10</td><td>S10</td></tr> <tr> <td>Solaris 11</td><td>S11</td></tr> </table> Any configuration files generated by zonep2chk will be applicable to the target release. See -c below. When run on Solaris 10, a target release of S11 can be specified, which will check for p2v into a Solaris 10 Branded zone. When the target is S10, it is assumed that a shared stack will be used. Any issues that will require an exclusive IP stack will be reported. When the target is S11, it is assumed that an exclusive IP stack will be used.	Global Zone	Default Target	Solaris 10	S10	Solaris 11	S11
Global Zone	Default Target						
Solaris 10	S10						
Solaris 11	S11						

If a particular feature in use by the global zone requires a particular patch/update level of the target to function, this information will be printed in the zonep2vchk output.

-P

Generate machine-parseable output. See the section “Parseable Output Format” below.

-c

Display a template zone configuration on stdout in the form of [zonecfg\(1M\)](#) export output. This configuration will contain resource limits and network configuration based on the source host's physical resources and networking configuration.

-b

Perform basic checks. This will check the global zone for issues that could prevent a successful p2v. This is the default behavior if none of -b, -c, -s, -S, -r, -x are specified.

-r{*time*} (h|m|s)

Perform runtime checks for the specified duration. This will analyze the currently executing processes in the global zone, and report issues that could prevent successful execution inside a non-global zone. Issues reported reflect actions made by the processes during the time in which zonep2vchk was executing.

-x

Perform runtime checks (as with -r) until SIGINT is received, such as is delivered by Ctrl -c from most shells.

-e *execname*[,*execname*...]

When performing runtime analysis (-r, -x), limit inspected programs to those matching the specified list of *execnames*. The *execname* is the name of process, as returned by ps -o comm. It is not necessary for named processes to exist when zonep2vchk is invoked. Any matching processes created while zonep2vchk is running will be inspected.

-E *file*

Similar to -e, but reads the list of *execnames* from *file*, one per line.

-s *path*[,*path*...]

Perform static binary analysis on the files or directories specified. This will inspect ELF binaries for system and library calls that might affect function inside a zone. Directories will be recursed, and non-ELF files will be ignored.

-S *file*

Similar to -s, but reads the path list from *file*, one per line.

Parseable Output Format

zonep2vchk will output a single line of parseable output for each issue detected. The line format is:

```
category:issue:field1:[field2:...]
```

Each field is delimited by a colon (:). Colon characters escaped with a backslash (\:) should not be treated as field delimiters.

Multiple instances of the same issue can be reported, each with fields describing the particular instance of the issue.

Below the existing categories and issues are defined. Future versions of zonep2vchk might include additional categories and issues. Existing issues might have new fields added after the existing fields for existing issues.

header Category The header category lists information about the source, target, and zonep2vchk version. The issues in this category are:

version

The version of the zonep2vchk command.

Field1: The version of the zonep2vchk command.

source

Information about the source system.

Field1: The nodename of the source system.

Field2: The /etc/release version of the source system.

Field3: The kernel version of the source system.

Field4: The platform of the source system.

target

Information about the specified target of the p2v check.

Field1: The Solaris version of the target.

Field2: The brand type that would be used on the target.

Field3: The ip-type of the expected zone on the target.

footer Category The footer category lists final summary information. The issues in this category are:

issues

A summary of the number of issues found.

Field1: The number of issues detected.

incompatible Category The incompatible category represents issues that will not function in a non-global zone. The issues in this category are:

etcsystem

An /etc/system tunable exists. These tunables do not function inside a zone. The /etc/system tunable can be transferred to the target global zone, but it will affect the entire system, including all zones and the global zone. If there is an alternate tunable that can be configured from within the zone, this tunable is described.

Field1:

The /etc/system tunable setting.

Field2:

One of:

noalternate	There is no alternate tunable from within a non-global zone.
obsolete	The tunable is obsolete on the target. It no longer serves any function.
replaced	The tunable has been replaced on the target. The replacement is configured in the global zone, and described by fields 3 and 4.
alternate	An alternate tunable exists. This tunable can be configured from within a non-global zone. The tunable is described by fields 3 and 4.
noinfo	zonep2vchk is not knowledgeable of the tunable. Tunable likely has no alternate inside a zone.

Field3:

Type of alternate/replacement tunable.

Field4:

Description of alternate/replacement tunable.

be

More than one boot environment exists. Only the active boot environment will be transferable to the non-global zone.

Field1: The name of the non-active boot environment.

unsupported

A feature is enabled that will not function in a zone.

Field1: mobileip The mobile IP agent, which does not function in a zone, is configured. See [mipagent\(1M\)](#) for details.

nfs (S10 sources only)

The system is sharing a filesystem by means of NFS. Native zones on Solaris 10 and Solaris 10 zones on Solaris 11 cannot share by means of NFS.

Field1: Path of file system being shared.

smb

The system is sharing a filesystem by means of in-kernel smb/cifs. Zones cannot share filesystems by means of SMB.

Field1: Path of file system being shared.

pkg

A package delivering software known not to work in a zone is installed.

Field1: Name of the package.

iscsi-target

The system is exporting an iSCSI target. Zones cannot export iSCSI targets.

Field1: Name of the iSCSI target.

fcoe-target

The system has configured an FCOE target. Zones cannot configure FCOE targets.

Field1: Ethernet device used.

Field2: WWN of the FCOE target.

fc-target

The system has configured an Fiberchannel target. Zones cannot configure Fiberchannel targets.

Field1: WWN of the Fiberchannel target.

npiv

The system has configured a virtual NPIV HBA. Zones cannot configure virtual HBAs.

Field1: Physical WWN hosting the virtual HBA.

Field2: Virtual WWN.

scsi

The system has configured an SCSI block device. Zones cannot configure scsi block devices.

Field1: Object configured as a SCSI device.

svcnotallowed

A service is enabled that will not function in a zone.

Field1: Name of the service.

resourcepool

A Solaris resource pool is configured. Zones cannot configure resource pools.

Field1: Name of the pool.

pset

A processor set is configured. Zones cannot configure processor sets.

Field1: Processor set ID.

Field2: List of CPU IDs in the processor set.

zones

Zones are configured. A zone cannot host zones. Any zones will not exist in the target non-global zone after p2v. Zones can be migrated separately using the detach/attach features in [zoneadm\(1M\)](#).

Field1: Name of the zone.

Field2: State of the zone.

lofi (Solaris 10 targets only)

A lofi device is configured. A zone cannot configure lofi devices.

Field1: Name of the lofi device.

Field2: Path of the file backing the device.

syscall (generated by -s and -f)

A binary makes a system or library call that cannot be made from a zone.

Field1: Name of the the binary file.

Field2: Name of the system or library call.

syscallargs (generated by -s and -f)

A binary makes a system or library call that cannot be made from a zone if called with certain arguments.

Field1: Name of the system or library call.

See regular output (no -P) for details on disallowed arguments.

lib (generated by -s and -f)

A binary links with a library that cannot be used inside a zone.

Field1: Name of the binary file.

Field2: Name of the disallowed library.

privnotallowed (generated by -r and -x)

A privilege is used by a process that cannot be added to a zone.

Field1: Name of the process.

Field2: Name of the privilege.

devnotallowed (generated by -r and -x)

A device is opened by a process that cannot be added to a zone.

Field1: Name of the process.

Field2: Name of the device.

configuration Category The configuration category represents issues that will require a configuration setting to allow the issue to function inside the non-global zone. This could be a [zonecfg\(1M\)](#) configuration setting, a configuration change in the global zone, or both.

The issues in this category are:

datalink

A datalink feature is configured that cannot be configured from within a zone. The datalink feature must be configured in the global zone, and if necessary, delegated to the zone using `zonecfg add anet` (Solaris 11 only) or `zonecfg add net`.

Field1: Name of the datalink feature. One of:

aggr	Aggregation.
ibiface	Infiniband interface.
ibpart	Infiniband partition.
vnic	Virtual NIC.
etherstub	Ethernet stub.
bridge	A bridge instance.
secobj	A wireless WPA or WEB security object.

Field2: Datalink object name.

dhcp-server (Solaris 10 targets only)

The host is a DHCP server. To provide DHCP service, a zone must have `ip-type=exclusive`, or have the the privilege `net_rawaccess` and the device `/dev/ip`. Note that this will allow a shared stack zone to read and write raw IP packets on the network, similar to an exclusive stack zone or global zone.

Field1: FMRI of the DHCP server service.

ntp-client

An NTP client service is enabled. This service updates the system clock. Since all zones share the same system clock, this service is disabled automatically during p2v. If it is desired that the zone update the system clock on the target host, the zone will need the privilege `sys_time`, and the service will need to be enabled inside the zone after p2v.

Field1: FMRI of the client service.

driverconf

A networking device contains configuration settings in its `.conf` file. Zones cannot configure drivers. The driver must be configured in the global zone. Some network driver settings might be configurable using [dladm\(1M\)](#) instead of editing a driver configuration file.

Field1: Path of the configuration file.

ifname (Solaris 10 targets only)

An existing configuration file will be impacted by the change of a network device name. For example, an `/etc/hostname.bge0` file will be impacted if the network device given to the target non-global zone is not `bge0`.

Field1: Path of the impacted file.

iscsi-initiator

The system is accessing an iSCSI target as a client. Zones cannot access iSCSI targets. The global zone must be the iSCSI initiator. The device can then be added to the zone using `zonecfg add device`.

Field1: iSCSI target being accessed.

fcoe-initiator

The system has an FCOE initiator configured. A zone cannot configure an FCOE initiator. The global zone must configure the FCOE initiator, and make the SCSI target devices available to the zone using `zonecfg add fs` or `zonecfg add device`.

Field1: Ethernet network device.

Field2: WWN of the initiator.

fc-initiator

The system has an HBA Fiberchannel port online. A zone cannot access a Fiberchannel target. The target must be accessed from the global zone and made available to the zone.

Field1: Fiberchannel HBA port WWN.

linkprop

Datalink properties are configured. A zone cannot configure datalink properties. They must be configured from the global zone.

Field1: Name of the datalink.

Field2: Property name

Field3: Property value.

ndd

Tunables that cannot be configured by a zone have been configured using `ndd`. These tunables must be configured from the global zone.

Field1: File or script setting the tunable.

Field2: Driver being tuned.

Field3: Tunable parameter.

dynaddr

One or more dynamically assigned IP addresses are configured on a network interface. These addresses are not supported with shared-IP zones. These IP addresses could change as a result of MAC address changes. You may need to modify this system's address information on the DHCP server and on the DNS, LDAP, or NIS name servers.

Field1 can be one of:

<code>dhcp</code>	Configured DHCP address. In this case, <i>Field2</i> is the name of the interface configured for DHCP.
<code>v6autoconf</code>	IPv6 stateless address configuration is enabled. In this case, <i>Field2</i> is the name of the interface with IPv6 auto configuration.

`rarp` (Solaris 10 source only) Reverse ARP assigned address is enabled. In this case, *Field2* is the name of the interface with reverse ARP enabled.

`patch` (Solaris 10 source with Solaris 11+ target only)
A patch is required before p2v into a non-global zone.

Field1: The patch required.

`physif` (Solaris 10 targets only)
A physical interface exists on the source system that will have to be replaced with a dedicated physical or VLAN interface on the destination system if migrating to an exclusive-IP zone.

Field1: Name of the interface on the source system.

`sched`
The system is configured with a default scheduling class. The default scheduling class of a non-global zone can be configured using the `zonecfg set scheduler` property. This will be provided in the `-c` output.

Field1: The configured default scheduling class.

`sharedip` (Solaris 10 targets only)
If migrating to a shared-IP zone, the following networking features will need to be configured from the global zone on behalf of the zone.

Field1 can be one of:

`ipmpgroup` An IPMP group is configured. If IPMP is required, it must be configured from the global zone. In this case, *Field2* is the IPMP group name.

`vni` A virtual network interface is configured. These must be configured from the global zone. In this case, *Field2* is the VNI interface name.

`v4forwarding`

`v6forwarding` IP forwarding (v4 or v6) is configured on an interface. In this case, *Field2* is the interface with IP forwarding configured.

`staticroute` Static routes are configured. Static routes must be configured from the global zone.

`exclusiveonly` (Solaris 10 targets only)
A networking feature is configured that is not supported for use with shared-IP zones. The feature will work without modification in exclusive-IP zones.

Field1:

`iptun` A IPv4, IPv6, or 6to4 tunnel interface has been plumbed.

Field2: Name of the tunnel interface.

sharedonly

A networking feature is configured that is not supported in an exclusive-IP zone. When migrating to a shared-IP zone, the feature must be configured in the global zone to support communication.

Field1:

cgtp A Carrier Grade Transport Protocol interface has been plumbed.

Field2: Name of the CGTP interface.

netdevalloc

A networking feature requires its underlying device be allocated to the zone with the `zonecfg(1M) add device` command. This feature is not supported with shared-IP zones.

Field1: Can be :

ppp Point-to-Point Protocol (PPP). PPP configuration files exist under `/etc/ppp`. The underlying device that needs to be allocated to the zone is either a serial port or, in the case of pppoe, an Ethernet physical or VNIC interface.

svcexlip (Solaris 10 targets only)

A service is enabled that will require an exclusive-IP zone.

Field1: Name of the service FMRI.

svcpriv

A service is enabled that will require additional privileges be added to the zone using the `zonecfg(1M) limitpriv` property.

Field1: FMRI of the service.

Field2: List of the privileges required by the service.

svm

A Solaris Volume Manager metadvice is configured. Metadevices must be configured in the global zone, and made available to the non-global zone using `zonecfg(1M) add device, add fs, or add dataset`.

Field1: Name of the metadvice.

ramdisk

A ramdisk device is configured. A zone cannot configure ramdisk devices.

Field1: Ramdisk device path.

vfstab

A filesystem mount is configured by means of `/etc/vfstab`. The filesystem must be migrated to the target global zone and made available to the non-global zone.

Field1: Device being mounted.

Field2: Mountpoint.

zpool

The system has additional zpools configured. These zpools must be migrated to the target global zone, and made available to the zone using `zonecfg add dataset` or `zonecfg add fs`.

Field1: Name of the pool.

privexclip (Solaris 10 targets only)

A process used a privilege that requires and exclusive-IP stack. See [zonecfg\(1M\)](#) for a description of the `ip- type` property.

Field1: Name of the process.

Field2: Privilege used.

devexclip (Solaris 10 targets only)

A process opened a device that requires an exclusive IP stack. See [zonecfg\(1M\)](#) for a description of the `ip- type` property.

Field1: Name of the process.

Field2: Name of the device.

privoptional

A process used a privilege that requires additional privilege be added to the target non-global zone. See [zonecfg\(1M\)](#) for a description of the `limitpriv` property.

Field1: Name of the process.

Field2: Privilege used.

devoptional

A process opened a device that is not available in a zone by default. See [zonecfg\(1M\)](#) for a description of the `add device` resource.

Field1: Name of the process.

Field2: Path of the device.

syscallpriv (generated by -s and -f)

A binary makes a system or library call that might require additional privilege be added to the target non-global zone. See [zonecfg\(1M\)](#) for a description of the `limitpriv` property. See the non-parseable output for details concerning the system or library call.

Field1: Path of the binary

Field2: Name of the system call.

`syscallexclip` (generated by `-s` and `-f`)

A binary makes a system or library call that might require an exclusive-ip stack. See [zonecfg\(1M\)](#) for a description of the `ip-type` property. See the non-parseable output for details concerning the system or library call.

Field1: Path of the binary

Field2: Name of the system call.

示例

示例 1 Performing Static Binary Analysis

The following command performs static analysis on all ELF binaries in two application directory trees:

```
# zonep2vchk -s /opt/myapplication,/usr/local
```

示例 2 Generating a Template for the Target Zone

The following command will generate a template zone configuration for Solaris 11 when run on a Solaris 10 global zone.

```
# zonep2vchk T S11 -c
```

示例 3 Analyzing Running Applications for a Period

The following command will analyze the process named `myapplication` for one hour and report any activity that might not function in a zone.

```
# zonep2vchk -s 1h -e myapplication
```

示例 4 Performing Basic Checks

The following command will analyze the global zone for configuration and Solaris features in use that might not function in a zone. Each discovered issue will be reported as a single line of parseable output.

```
# zonep2vchk -bP
```

退出状态

The following exit values are returned:

- 0
Successful completion, no issues detected.
- 1
An internal error occurred.
- 2
Invalid usage.
- 3
One or more issues were detected.

属性 See [attributes\(5\)](#) for descriptions of the following attributes:

ATTRIBUTE TYPE	ATTRIBUTE VALUE
Availability	system/zones
Interface Stability	See below

Command invocation and parseable output is Committed. Human readable output (default output) is Uncommitted.

另请参见 [dladm\(1M\)](#), [mipagent\(1M\)](#), [zoneadm\(1M\)](#), [zonecfg\(1M\)](#), [attributes\(5\)](#), [solaris\(5\)](#), [solaris10\(5\)](#), [zones\(5\)](#)

附注 The static (-s and -f) checks make use of the [elfdump\(1\)](#) utility, which is delivered by the following package:

Solaris 11 developer/base-developer-utilities

Solaris 10 SUNWbtool

The runtime (-r) checks make use of the [dtrace\(1M\)](#) utility, which is delivered by the following package:

Solaris 11 system/dtrace

Solaris 10 SUNWdtrc

引用名 `zonestatd` – 区域监视守护进程

用法概要 `/usr/lib/zones/zonestatd`

描述 `zonestatd` 是一个在系统引导期间启动的系统守护进程。它监视区域使用系统资源的情况，以及区域和系统的配置信息，如 `psrset psets`、池 `psets` 和资源控制设置。

此守护进程由区域管理软件自动启动，而不应直接调用。它不属于编程接口，而应归类为专用接口。

安全 要想让每个非全局区域 (NGZ) 中的 `zonestat` 服务正常运行，全局区域中的 `zonestat` 服务必须处于联机状态。每个 NGZ 中的 `zonestat` 服务并不直接读取系统配置和使用数据，而是从全局区域上的 `zonestat` 服务中读取。

属性 有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	<code>system/zones</code>
接口稳定性	<code>Private</code> （专用）

另请参见 [prctl\(1\)](#)、[svcs\(1\)](#)、[zonestat\(1\)](#)、[acctadm\(1M\)](#)、[pooladm\(1M\)](#)、[poolcfg\(1M\)](#)、[rcapadm\(1M\)](#)、[svcadm\(1M\)](#)

附注 `zonestat` 服务由服务管理工具 [smf\(5\)](#) 管理，其服务标识符为：

`svc:/system/zones-monitoring:default`

可以使用 [svcadm\(1M\)](#) 来执行对此服务的管理操作（如启用、禁用或请求重新启动）。可以使用 [svcs\(1\)](#) 命令来查询服务的状态。

`zonestat` 服务具有以下 SMF 配置属性：

`config/sample_interval`

该属性设置 `zonestatd` 采样间隔。该间隔是区域监视守护进程 `zonestatd` 在对资源使用情况进行采样时所使用的间隔。该间隔也是用于确定配置更改（如处理器集更改、资源控制更改和区域状态更改）的间隔。

缺省间隔为 5 秒。

`zonestat` 服务利用了扩展记帐功能。该服务启用进程记帐资源跟踪功能（如果尚未启用），并配置进程记帐文件。`zonestat` 服务将按配置的间隔滚动进程记帐日志（请参见 [zonestat\(1\)](#)）。

如果在外部启用了扩展进程记帐，`zonestat` 服务将按照配置使用进程记帐日志。它不会滚动记帐日志，但如果记帐日志在外部滚动，它仍会正确工作。

引用名

zpool – 配置 ZFS 存储池

用法概要

```

zpool [-?]

zpool help command | help | property property-name

zpool add [-f] [-n [-l]] pool vdev ...

zpool attach [-f] pool device new_device

zpool clear [-F [-n]] pool [device]

zpool create [-f] [-n [-l]] [-o property=value] ...
    [-O file-system-property=value] ... [-m mountpoint]
    [-R root] pool vdev ...

zpool destroy [-f] pool

zpool detach pool device

zpool export [-f] pool ...

zpool get "all" | property[,...] pool ...

zpool history [-il] [pool] ...

zpool import [-d path ...] [-D]

zpool import [-o mntopts] [-o property=value] ... [-d path ... |
    -c cacheFile] [-D] [-f] [-m] [-N] [-R root] [-F [-n [-l]]] -a

zpool import [-o mntopts] [-o property=value] ... [-d path ... |
    -c cacheFile] [-D] [-f] [-m] [-R root] [-F [-n [-l]]]
    pool | id [newpool]

zpool iostat [-T d|u ] [-v [-l]] [pool] ... [interval [count]]

zpool list [-H] [-o property[,...]] [-T d|u ] [pool] ... [interval [count]]

zpool offline [-t] pool device ...

zpool online [-e] pool device ...

zpool remove pool device ...

zpool replace [-f] pool device [new_device]

zpool scrub [-s] pool ...

zpool set property=value pool

zpool split [-R altroot] [-n [-l]] [-o mntopts] [-o property=value] pool
    newpool [device ...]

zpool status [-l] [-v] [-x] [-T d|u ] [pool] ... [interval [count]]

zpool upgrade

zpool upgrade -v

zpool upgrade [-V version] -a | pool ...

```

描述 `zpool` 命令配置 ZFS 存储池。存储池是一种设备集合，提供 ZFS 数据集的物理存储和数据复制。

存储池内的所有数据集共享同一个空间。有关管理数据集的信息，请参见 [zfs\(1M\)](#)。

虚拟设备 (vdev) “虚拟设备”描述的是单个设备或者按特定的性能和故障特征组织的设备集合。支持的虚拟设备如下：

disk

一种块设备，通常位于 `/dev/dsk` 下。ZFS 可以使用单个分片或分区，但推荐的操作模式是使用整个磁盘。可以通过全路径指定磁盘，也可以使用缩略名称（`/dev/dsk` 下的相对路径部分）。通过省略分片或分区部分，可以指定整个磁盘。或者，可以使用描述磁盘当前位置的 `/dev/chassis/.../disk` 路径指定整个磁盘。如果指定整个磁盘，ZFS 会在必要时自动为磁盘添加标签。

file

一种常规文件。强烈建议不要使用文件作为后备存储。它主要是为试验目的而设计的，因为文件的容错能力与其所属的文件系统相同。必须用全路径指定文件。

mirror

两个或更多设备的镜像。数据以完全相同的方式在一个镜像的所有组件之间进行复制。包含 N 个大小为 X 的磁盘的镜像可以存放 X 字节，并可承受 $(N-1)$ 个设备故障而不会损害数据完整性。

raidz

raidz1

raidz2

raidz3

RAID-5 的一种变体，可以更好地分配奇偶校验，消除“RAID-5 写入漏洞”（在这种漏洞中，数据和奇偶校验会在断电后变得不一致）。数据和奇偶校验在 `raidz` 组内的所有磁盘间呈条带分布。

`raidz` 组可以有单一、双重或三重奇偶校验，即 `raidz` 组分别可以承受一次、两次或三次故障而不会丢失任何数据。`raidz1 vdev` 类型指定单一奇偶校验 `raidz` 组，`raidz2 vdev` 类型指定双重奇偶校验 `raidz` 组，`raidz3 vdev` 类型指定三重奇偶校验 `raidz` 组。`raidz vdev` 类型是 `raidz1` 的一个别名。

一个包含 N 个大小为 X 的 `raidz` 组（其中有 P 个奇偶校验磁盘）可以存放大约 $(N-P)*X$ 字节，并可承受 P 个设备故障而不会损害数据完整性。一个 `raidz` 组中的设备数量至少比奇偶校验磁盘数多一个。推荐的数量在 3 到 9 个之间，这有助于提高性能。

spare

一种特殊的假 `vdev`，可以跟踪池的可用热备件。有关更多信息，请参见“热备件”部分。

log

独立意图日志设备。如果指定了多个日志设备，则将在设备之间进行写负载均衡。可以镜像日志设备。但是，意图日志不支持 **raidz vdev** 类型。有关更多信息，请参见“意图日志”部分。

cache

用于高速缓存存储池数据的设备。不能将高速缓存设备配置为镜像或 **raidz** 组。有关更多信息，请参见“高速缓存设备”部分。

虚拟设备不能嵌套，因此镜像或 **raidz** 虚拟设备只能包含文件或磁盘。不允许有镜像的镜像（或其他组合）。

一个池可以在配置的顶层具有任意数目的虚拟设备，称为根 **vdev**。数据在所有顶层设备间动态分布，以平衡设备间的数据。添加新的虚拟设备时，ZFS 自动将数据放在新近可用的设备上。

虚拟设备在命令行上是逐次指定的，多个虚拟设备间用空格分隔。关键字 **mirror** 和 **raidz** 用于区分一个组的结束和另一个组的开始。例如，以下命令将创建两个根 **vdev**，其中每个根 **vdev** 都是两个磁盘的镜像。

```
# zpool create mypool mirror c0t0d0 c0t1d0 mirror c1t0d0 c1t1d0
```

或者，可以使用以下命令：

```
# zpool create tank \
mirror \
    /dev/chassis/RACK29.U01-04/DISK_00/disk \
    /dev/chassis/RACK29.U05-08/DISK_00/disk \
mirror \
    /dev/chassis/RACK29.U01-04/DISK_01/disk \
    /dev/chassis/RACK29.U05-08/DISK_01/disk
```

设备故障和恢复

ZFS 支持一套丰富的机制，用于处理设备故障和数据损坏。会对所有元数据和数据进行校验和计算，在检测到数据损坏时，ZFS 会自动根据好的副本修复坏数据。

为了利用这些功能，池必须采用某种形式的冗余（使用镜像组或 **raidz** 组）。虽然 ZFS 支持在非冗余配置下运行，但在每个根 **vdev** 都只是一个磁盘或文件的情况下，强烈建议不要这样做。只要出现一个数据位损坏，就会导致部分或所有数据不可用。

池的正常运行状况用三种状态之一进行描述：联机、降级或故障。联机池中的所有设备都正常运行。降级池中的一个或多个设备发生了故障，但由于有冗余配置，数据仍然可用。故障池中有损坏的元数据，或者有一个或多个故障设备，并且没有足够的副本支持继续运行。

顶层 **vdev**（如镜像或 **raidz** 设备）的运行状况可能会受与其关联的 **vdev** 或组件设备状态的影响。顶层 **vdev** 或组件设备处于以下状态之一：

DEGRADED（降级）

一个或多个顶层 **vdev** 处于降级状态，因为有一个或多个组件设备脱机。有足够的副本支持继续运行。

一个或多个组件设备处于降级或故障状态，但有足够的副本支持继续运行。底层条件如下：

- 校验和错误的数量超出了可接受的程度，设备降级，这表明某个组件可能出现问题。ZFS 在必要时会继续使用该设备。
- I/O 错误的数量超出了可接受的程度。不能将该设备标记为故障，因为还有足够的副本支持继续运行。

FAULTED (故障)

一个或多个顶层 vdev 处于故障状态，因为有一个或多个组件设备脱机。没有足够的副本支持继续运行。

一个或多个组件设备处于故障状态，并且没有足够的副本支持继续运行。底层条件如下：

- 可以打开设备，但内容与预期值不匹配。
- I/O 错误的数量超过了可接受的程度，设备处于故障状态，以防止继续使用。

OFFLINE (脱机)

已通过 `zpool offline` 命令明确将设备置于脱机状态。

ONLINE (联机)

设备联机并正常运行。

REMOVED (移除)

系统正在运行时已物理移除了该设备。设备移除检测依赖于硬件，而且并非在所有平台上都受支持。

UNAVAIL (不可用)

无法打开设备。如果在设备不可用时导入了某个池，则将用唯一标识符（而不是路径）标识该设备，因为最初的路径绝对不会正确。

如果移除设备后重新将其连接到系统，ZFS 会自动尝试将设备置于联机状态。设备连接检测依赖于硬件，而且并非在所有平台上都受支持。

热备件

ZFS 允许设备以“热备件”的方式与池关联。这些设备在池中的使用情况并不活跃，但在某个活动设备发生故障时，热备件会自动替换它。要创建带热备件的池，请使用任意数量的设备指定“备件”vdev。例如，

```
# zpool create pool mirror c0d0 c1d0 spare c2d0 c3d0
```

可以使用 `zpool add` 命令添加备件，使用 `zpool remove` 命令删除备件。一旦启动了备件替换，将在配置中创建一个新的“备件”vdev，并且该备件将保留到原始设备被替换为止。此时，如果又有另外的设备发生故障，热备件将重新变得可用。

通过分离热备件，可以取消正在进行的备件替换。如果分离了原始故障设备，则热备件将取代故障备件在配置中的位置，并会从所有活动池的备件列表中移除。

如果以物理方式替换了原始故障设备、使其重新联机或者清除了错误（通过 FMA 事件或者通过使用 `zpool online` 或 `zpool clear` 命令），且原始设备状态变为正常运行，则 INUSE 备件设备将再次变为 AVAIL。

备件不能替换日志设备。

意图日志

ZFS 意图日志 (ZIL) 可满足同步事务的 POSIX 要求。例如，数据库通常要求其事务在从系统调用返回时在稳定的存储设备上。NFS 和其他应用程序也可以使用 `fsync()` 来确保数据的稳定性。缺省情况下，意图日志是从主池内的块中分配的。但是，使用单独的意图日志设备（如 NVRAM 或专用磁盘）可能会获得更好的性能。例如：

```
# zpool create pool c0d0 c1d0 log c2d0
```

此外，还可以指定多个日志设备，并对它们进行镜像。有关镜像多个日志设备的示例，请参见“示例”部分。

可以将日志设备作为较大池的一部分添加、替换、附加、分离以及导入和导出。通过为日志指定顶层镜像，可以移除镜像日志设备。

高速缓存设备

可以将设备作为“高速缓存设备”添加到存储池中。这些设备在主内存和磁盘之间提供了一个进行高速缓存的附加层。对于读取操作量很大的工作负荷，当工作集大小比主内存中可高速缓存的内容大得多时，使用高速缓存设备可以通过低延迟介质处理此工作集的更大部分。使用高速缓存设备，可以最大程度地提高大多数静态内容的随机读取工作量的性能。

要创建带高速缓存设备的池，请使用任意数量的设备指定“高速缓存”`vdev`。例如：

```
# zpool create pool c0d0 c1d0 cache c2d0 c3d0
```

高速缓存设备不能镜像，也不能成为 `raidz` 配置的一部分。如果高速缓存设备发生读取错误，则该读取 I/O 会重新发送给原始存储池设备（它可能是镜像配置或 `raidz` 配置的一部分）。

高速缓存设备的内容被认为具有易失性，其他系统的高速缓存也具有这个特点。

进程

每个导入的池都有一个关联的进程 `zpool-poolname`。此进程中的线程是该池的 I/O 处理线程，该线程处理压缩、检验和操作以及与该池关联的所有 I/O 的其他任务。此进程的目的在于呈现系统存储池的 CPU 使用率信息。此进程的存在形成了一个不稳定的接口。

属性

每个池都有多个关联的属性。有些属性是只读的统计信息，而其他属性是可配置的，可以改变池的行为。以下是只读属性：

allocated

池中已实际分配的存储空间量。该属性也可通过其简短列名 `alloc` 来引用。

capacity

已用的池空间百分比。此属性也可通过其简短列名 `cap` 来引用。

dedupratio

为池指定的重复数据删除比率，用乘数因子表示。该值用一个十进制数表示。例如，`dedupratio` 值 1.76 表示存储了 1.76 个单位的数据，但实际消耗的磁盘空间只有 1 个单位。该属性也可通过其简短列名 `dedup` 来引用。

可按如下方式启用重复数据删除：

```
# zfs set dedup=on pool/dataset
```

缺省值为 `off`（关闭）。

有关重复数据删除功能的说明，请参见 [zfs\(1M\)](#)。

free

池中未分配的块的数量。

guid

池的唯一标识符。

health

池的当前运行状况。运行状况可以是 `"ONLINE"`、`"DEGRADED"` 或 `"FAULTED"`。

size

存储池的总大小。

这些空间使用属性报告存储池的实际可用物理空间。物理空间可能与空间总量不同，后者是所有包含的数据集可以实际使用的空间量。`raidz` 配置中使用的空间量取决于写入的数据的特征。此外，ZFS 会保留一些空间用于内部记帐，[zfs\(1M\)](#) 命令会考虑这一点，但 `zpool` 命令则不会。对于合理大小的未满载池，应该看不出这些影响。对于小型池，或者接近全满的池，这些差异可能会变得更加明显。

以下属性可在创建或导入时设置：

altroot

备用根目录。如果进行了设置，则该目录会被前置到池中的任何挂载点。检查未知池（无法信任其中的挂载点）时，或在典型路径无效的备用启动环境中，可以使用此属性。`altroot` 不是持久属性。它仅在系统运行期间有效。设置 `altroot` 即表示缺省情况下使用 `cachefile=none`，但此设置可以使用显式设置来覆盖。

以下属性可在导入时设置：

readonly=on|off

控制是否可以修改池。启用后，在以读写模式导入池之前，任何仅存在于意图日志中的同步数据都不可访问。

以只读模式导入池具有以下限制：

- 导入期间试图设置其他池属性将被忽略。
- 所有文件系统挂载都将转换为包括只读 (`ro`) 挂载选项。

通过导出和导入池可将已以只读模式导入的池恢复为读写模式。

以下属性可在创建时和导入时设置，并可在以后使用 `zpool set` 命令更改：

autoexpand=on|off

在底层 LUN 增长时控制自动池扩展。如果设置为 **on**，池将根据扩展设备的大小重新调整大小。如果设备是某个镜像或 **raidz** 的一部分，则在新空间对池可用之前，该镜像/**raidz** 组内的所有设备都必须扩展。缺省行为为 **off**（关闭）。该属性也可通过其简短列名 **expand**（扩展）来引用。

autoreplace=on|off

控制设备的自动替换。如果设置为 **off**（关闭），则必须由管理员使用 **zpool replace** 命令来启动设备替换。如果设置为 **on**（打开），则会自动对与先前属于池的设备在同一物理位置的任何新设备进行格式化和替换。缺省行为为 **off**（关闭）。此属性也可通过其简短列名 **replace**（替换）来引用。

bootfs=pool/dataset

标识根池的缺省可引导数据集。此属性应主要通过安装和升级程序设置。

cachefile=path|none

控制高速缓存池配置的位置。要在系统启动时发现所有池，需要存储在根文件系统上的配置数据的高速缓存副本。系统引导时会自动导入此高速缓存中的所有池。有些环境（如安装和群集）需要将此信息高速缓存在其他位置，从而使池不会被自动导入。设置此属性后，池配置将高速缓存在一个不同的位置，以后可以使用 **zpool import -c** 将其导入。如果将其设置为特殊值 **none**（无），将创建一个永远不会高速缓存的临时池；如果设置为特殊值 **''**（空字符串），则将使用缺省位置。

多个池可以共享同一个高速缓存文件。因为在添加或移除池时，内核会销毁并重新创建此文件，因此在试图访问此文件时应小心。在导出或销毁了最后一个使用 **cachefile** 的池后，将删除该文件。

dedupditto=number

设置副本数的阈值。如果某个消重后的块的引用数超过了此阈值，将自动存储该块的另一个相同副本。缺省值为 **0**。

delegation=on|off

控制是否根据在数据集上定义的数据集权限授予非特权用户访问权。有关 ZFS 委托管理的更多信息，请参见 **zfs(1M)**。

failmode=wait|continue|panic

控制在出现灾难性池故障时的系统行为。这种情况通常是由于失去与基础存储设备的连接或池中所有设备出现故障而导致的。此类事件的行为按以下方式确定：

wait

阻止对池的所有 I/O 访问，直到恢复设备连接并清除错误。在设备问题得到解决之前，池一直处于 **wait** 状态。这是缺省行为。

continue

使 EIO 返回到任何新的写入 I/O 请求，但允许对其余任何运行状况良好的设备执行读取操作。将会阻止必须提交到磁盘的任何写入请求。

panic

向控制台输出一则消息，并生成一个系统故障转储。

`listsnapshots=on|off`

控制在不使用 `-t` 选项运行 `zfs list` 时，是否输出与该池关联的快照的相关信息。缺省值为 "off"。

`version=version`

池的当前盘上版本。此版本可以提高，但绝不能降低。尽管在为了实现向后兼容性而需要一个特定版本时可以使用此属性，但首选的池更新方法是使用 `zpool upgrade` 命令。此属性可以是 1 与 `zpool upgrade -v` 所报告的当前版本之间的任何数值。

子命令

修改状态的所有子命令将以原始形式永久记录到池中。

`zpool` 命令提供了一些子命令，可以创建和销毁存储池，为存储池增加容量，以及提供有关存储池的信息。支持以下子命令：

`zpool -?`

显示帮助消息。

`zpool help command | help | property property-name`

显示 `zpool` 命令用法。可以显示特定命令或属性的帮助。如果显示特定命令或属性的帮助，将会显示命令语法或可用的属性值。使用不带任何参数的 `zpool help` 可显示 `zpool` 命令的完整列表。

`zpool add [-f] [-n [-l]] pool vdev ...`

将指定的虚拟设备添加到给定的池。`vdev` 规范在“虚拟设备”部分中介绍。`-f` 选项的行为以及所执行的设备检查在“`zpool create`”子命令中介绍。

`-f`

强制使用 `vdev`，即使他们显示为正在使用；或者指定一个冲突的复制级别。并非所有设备都可以通过这种方式覆盖。

`-n`

显示在不实际添加 `vdev` 的情况下使用的配置。在实际创建池时，因为特权不足或设备共享的原因，操作仍然可能失败。

`-l`

如果可能，让 `-n` 以当前 `/dev/chassis` 位置格式显示配置。

不要添加当前配置为 `zpool` 的法定设备的磁盘。某个磁盘进入池中后，就可将该磁盘配置为法定设备。

`zpool attach [-f] pool device new_device`

将 `new_device` 连接到现有的 `zpool` 设备。现有设备不能是 `raidz` 配置的一部分。如果 `device` 当前不是某个镜像配置的一部分，`device` 将自动转换为 `device` 和 `new_device` 的双向镜像。如果 `device` 是某个双向镜像的一部分，则连接 `new_device` 将创建一个三向镜像，依此类推。无论哪种情况，`new_device` 都会立即开始重新同步。

`-f`

强制使用 `new_device`，即使它显示为正在使用。并非所有设备都可以通过这种方式覆盖。

zpool clear [-F [-n]] pool [device] ...

清除某个池中的设备错误。如果未指定任何参数，将清除池中的所有设备错误。如果指定了一个或多个设备，将仅清除与指定设备关联的错误。

-F

启动不可打开池的恢复模式。尝试放弃池中最后少数几个事务会将其恢复到可打开状态。并非所有损坏的池都能使用此选项恢复。如果成功，将不可挽回地丢失所放弃事务中的数据。

-n

与 **-F** 标志结合使用。检查放弃事务是否会使池可打开，但不会实际放弃任何事务。

zpool create [-f] [-n [-l]] [-o property=value] ... [-O file-system-property=value] ... [-m mountpoint] [-R root] pool vdev ...

创建一个新的存储池，其中包含在命令行上指定的虚拟设备。池名称必须以字母开头，并且只能包含字母数字字符以及下划线 (“_”)、连字符 (“-”) 和句点 (“.”)。池名称 **mirror**、**raidz**、**spare** 和 **log** 是保留名称，与以模式 **c[0-9]** 开头的名称相同。**vdev** 规范在“虚拟设备”部分中介绍。

该命令确认指定的每个设备可访问，并且其他子系统当前未使用这些设备。有一些使用情况（如当前正在挂载，或者被指定为专用转储设备）会阻止设备被 ZFS 使用。其他使用情况（例如，有预先存在的 UFS 文件系统）可以用 **-f** 选项覆盖。

该命令还会检查池的复制策略是否一致。除非指定 **-f**，否则试图将冗余和非冗余存储组合在一个池中或者混合磁盘和文件会导致错误。除非指定 **-f**，否则在单个 **raidz** 或镜像组中使用不同大小的设备也会加上错误标志。

除非指定 **-R** 选项，否则缺省挂载点为 **/pool**。不得有挂载点，或者即使存在，也必须为空，否则将无法挂载根数据集。此选项可用 **-m** 选项覆盖。

-f

强制使用 **vdev**，即使他们显示为正在使用；或者指定一个冲突的复制级别。并非所有设备都可以通过这种方式覆盖。

-n

显示将在不实际创建池的情况下使用的配置。实际的池创建操作仍可能因特权不足或设备当前正处于使用状态而失败。

-l

如果可能，让 **-n** 以当前 **/dev/chassis** 位置格式显示配置。

-o property=value [-o property=value] ...

设置指定的池属性。有关可设置的有效属性列表，请参见“属性”部分。

-O file-system-property=value

[-O file-system-property=value] ...

设置池的顶层文件系统的指定属性。有关可设置的有效属性列表，请参见 [zfs\(1M\)](#) 的“属性”部分。

-R root

等效于 `-o cachefile=none,altroot=root`

-m mountpoint

设置池的顶层文件系统的挂载点。如果指定了 `altroot`，则缺省挂载点为 `/pool` 或 `altroot/pool`。挂载点必须是一个绝对路径 `legacy`（传统）或 `none`（无）。有关数据挂载点的更多信息，请参见 [zfs\(1M\)](#)。

zpool destroy [-f] pool

销毁指定的池，释放所有设备，使之可用于其他用途。此命令会尝试在销毁池之前卸载所有活动的数据集。

-f

强制卸载池中包含的所有活动数据集。

zpool detach pool device

从镜像存储池分离 *device* 或备件。如果以物理方式替换了现有设备，则也可以从 RAID-Z 存储池分离备件。或者，可以分离 RAID-Z 存储池中的现有设备（如果已替换为备件）。如果没有数据的其他有效副本，将拒绝该操作。

zpool export [-f] pool ...

从系统中导出指定的池。所有设备将标记为已导出，但其他子系统仍认为它们在使用。只要存在足够的设备数量，就可以在系统之间移动设备（即使这些设备采用不同的字节存储顺序 (*endianness*)）以及导入这些设备。

在导出池之前，将卸载池中的所有数据集。

要使池可移植，必须在 `zpool` 命令中指定整个磁盘（而不仅仅是分片），从而使 ZFS 可以使用可移植的 EFI 标签来设置磁盘标签。否则，采用不同字节存储顺序的平台上的磁盘驱动程序将无法识别磁盘。

-f

使用 `umount -f` 命令强制卸载所有数据集。

此命令将强制导出池。

zpool get "all" | property[...] pool ...

检索指定存储池的指定属性列表或所有属性（如果使用了 `all`）。这些属性以下列字段显示：

<code>name</code>	Name of storage pool
<code>property</code>	Property name
<code>value</code>	Property value
<code>source</code>	Property source, either 'default' or 'local'.

有关可用池属性的更多信息，请参见“属性”部分。

zpool history [-il] [pool] ...

显示指定池或所有池（如果没有指定任何池）的命令历史记录。

-i
显示内部记录的 ZFS 事件以及用户启动的事件。

-l
以长格式显示日志记录；长格式除了标准格式以外，还包括用户名、主机名以及执行操作所在的区域。

zpool import [-d *path* ... | -c *cachefile*] [-D]

列出可导入的池。如果未指定 **-d** 选项，此命令将在 `/dev/dsk` 中搜索设备。可以多次指定 **-d** 选项，将搜索所有目录和设备路径。如果设备显示为导出池的一部分，则此命令将显示池的摘要，其中包括池的名称、数字标识符以及 *vdev* 布局 and 每个设备或文件的当前运行状况。除非指定 **-D** 选项，否则不会列出销毁的池（即以前使用 **zpool destroy** 命令销毁的池）。

数字标识符是唯一的，当有多个名称相同的导出池时，可用于代替池名称。

-c *cachefile*
从使用 *cachefile* 池属性创建的指定 *cachefile* 中读取配置。将使用此 *cachefile* 代替搜索设备。

-d *path*
在 *path* 中搜索设备或文件，其中 *path* 可以是目录或设备路径。可以多次指定 **-d** 选项。

-D
仅列出销毁的池。

zpool import [-o *mntopts*] [-o *property=value*] ... [-d *path* ... | -c *cachefile*] [-D] [-f] [-m] [-N] [-R *root*] [-F [-n [-l]]] -a

导入在搜索目录或设备路径中找到的所有池。与前一个命令相同，但会导入具有足够的可用设备数量的所有池。除非指定 **-D** 选项，否则不会导入销毁的池（即以前使用 **zpool destroy** 命令销毁的池）。

-o *mntopts*
挂载池中的数据集时要使用的挂载选项列表（用逗号分隔）。有关数据集属性和挂载选项的说明，请参见 [zfs\(1M\)](#)。

-o *property=value*
对导入的池设置指定属性。有关可用池属性的更多信息，请参见“属性”部分。

-c *cachefile*
从使用 *cachefile* 池属性创建的指定 *cachefile* 中读取配置。将使用此 *cachefile* 代替搜索设备。

-d *path*
在 *path* 中搜索设备或文件。可以多次指定 **-d** 选项。此选项与 **-c** 选项不兼容。

-D
仅导入销毁的池。还需要 **-f** 选项。

- f
强制导入，即使池看上去可能处于活动状态。
- F
不可导入池的恢复模式。尝试通过放弃最后几个事务将池恢复到可导入状态。并非所有损坏的池都能使用此选项恢复。如果成功，将不可挽回地丢失所放弃事务中的数据。如果池是可导入的，或者已经导入，将忽略此选项。
- a
搜索并导入所有找到的池。
- m
允许在缺少日志设备时导入池。
- R *root*
将 *cachefile* 属性设置为 none（无），将 *altroot* 属性设置为 *root*。
- N
在不挂载任何文件系统的情况下导入池。
- n
与 -F 恢复选项一起使用。确定是否可将不可导入的池重新变得可导入，但不实际执行池恢复。有关池恢复模式的更多详细信息，请参见前面的 -F 选项。
- l
如果可能，让 -n 以当前 */dev/chassis* 位置格式显示信息。

```
zpool import [-o mntopts] [-o property=value] ... [-d path ...] [-c cachefile] [-D] [-f] [-m] [-R root] [-F [-n [-l]]] pool | id [newpool]
```

导入特定的池。池可以用名称或数字标识符来标识。如果指定了 *newpool*，将使用持久名称 *newpool* 导入池。否则，使用与其导出名称相同的名称导入。请不要使用新名称导入根池。否则，系统可能无法引导。

如果没有先运行 `zpool export` 就从系统中移除了设备，则设备看上去可能处于活动状态。这将无法确定这是失败的导出，还是真的有其他主机在使用该设备。要导入处于此状态的池，必须使用 -f 选项。

-o *mntopts*
挂载池中的数据集时要使用的挂载选项列表（用逗号分隔）。有关数据集属性和挂载选项的说明，请参见 [zfs\(1M\)](#)。

-o *property=value*
对导入的池设置指定属性。有关可用池属性的更多信息，请参见“属性”部分。

-c *cachefile*
从使用 *cachefile* 池属性创建的指定 *cachefile* 中读取配置。将使用此 *cachefile* 代替搜索设备。

-d *path*
在 *path* 中搜索设备或文件。可以多次指定 -d 选项。此选项与 -c 选项不兼容。

- D
导入销毁的池。还需要 **-f** 选项。
- f
强制导入，即使池看上去可能处于活动状态。
- F
不可导入池的恢复模式。尝试通过放弃最后几个事务将池恢复到可导入状态。并非所有损坏的池都能使用此选项恢复。如果成功，将不可挽回地丢失所放弃事务中的数据。如果池是可导入的，或者已经导入，将忽略此选项。
- R *root*
将 *cachefile* 属性设置为 *none*（无），将 *altroot* 属性设置为 *root*。
- n
与 **-F** 恢复选项一起使用。确定是否可将不可导入的池重新变得可导入，但不实际执行池恢复。有关池恢复模式的更多详细信息，请参见前面的 **-F** 选项。
- l
如果可能，让 **-n** 以当前 */dev/chassis* 位置格式显示信息。
- m
允许在缺少日志设备时导入池。

zpool iostat [-T *d|u*] [-v [-l]] [*pool*] ... [*interval* [*count*]]

显示指定池的 I/O 统计信息。如果指定了一个时间间隔，则每隔 *interval* 秒就会输出统计信息，直到按 **Ctrl-C** 为止。如果未指定 *pools*，将显示系统中每个池的统计信息。如果指定了 *count*，该命令将在输出 *count* 报告后退出。

-T *d|u*
显示时间戳。

指定 *d* 表示标准日期格式。请参见 [date\(1\)](#)。指定 *u* 表示时间的内部表示形式的印刷表示形式。请参见 [time\(2\)](#)。

- v
详细统计信息。除了报告池范围的统计信息外，还报告池中单个 *vdev* 的使用统计信息。
- l
如果可能，让 **-v** 以当前 */dev/chassis* 位置格式显示 *vdev* 统计信息。

zpool list [-H] [-o *props*[,...]] [-T *d|u*] [*pool*] ...

列出指定池以及运行状态和空间使用情况。如果未指代任何参数，将列出系统中的所有池。

如果指定了时间间隔，将每隔 *interval* 秒就会显示状态和空间使用情况，直到按 **Ctrl-C** 为止。如果指定了 *count*，该命令将在显示 *count* 报告后退出。

- H
脚本模式。不显示标题，以单个制表符（而不是任意空格）来分隔字段。

-o props

要显示的属性的逗号分隔列表。有关有效属性列表，请参见“属性”部分。缺省列表为 `name, size, allocated, free, capacity, health, altroot`。

-T d|u

显示时间戳。

指定 `d` 表示标准日期格式。请参见 [date\(1\)](#)。指定 `u` 表示时间的内部表示形式的印刷表示形式。请参见 [time\(2\)](#)。

zpool offline [-t] pool device ...

将指定的物理设备脱机。当 *device* 处于脱机状态时，不会尝试对设备进行任何读或写操作。

此命令不适用于备件或高速缓存设备。

-t

临时。重新引导时，指定的物理设备将恢复到其以前的状态。

zpool online [-e] pool device...

将指定的物理设备联机。

此命令不适用于备件或高速缓存设备。

-e

扩展设备以使用所有可用空间。如果设备是某个镜像或 `raidz` 的一部分，则在新空间对池可用之前，所有设备都必须扩展。

zpool remove pool device ...

从池中移除指定的设备。此命令当前仅支持移除热备件设备、高速缓存设备和日志设备。通过为日志指定顶层镜像，可以移除一个镜像日志设备。可以使用 `zpool detach` 命令移除属于某个镜像配置一部分的非日志设备。不能从池中移除非冗余设备和 `raidz` 设备。

zpool replace [-f] pool old_device [new_device]

将 *old_device* 替换为 *new_device*。该命令等效于连接 *new_device*，等待它重新同步，然后分离 *old_device*。

new_device 的大小必须大于或等于镜像配置或 `raidz` 配置中所有设备的最小大小。

如果池不是冗余池，则必须要有 *new_device*。如果不指定 *new_device*，将缺省为 *old_device*。在现有磁盘出现故障并已进行物理更换后，这种形式的替换很有用。在这种情况下，新磁盘可能有着与旧设备相同的 `/dev/dsk` 路径，尽管它实际上是一个不同的磁盘。ZFS 会识别这一情况。

在 `zpool status` 输出中，*old_device* 显示在 `replacing` 一词下，其后附有字符串 `/old`。重新同步完成后，将自动移除 `replacing` 和 *old_device*。如果新设备在重新同步之前发生故障，并且有第三个设备安装在它的位置，则两个故障设备都将显示为后面附有 `/old`，并且重新同步过程将重新开始。重新同步完成后，两个 `/old` 设备将与 `replacing` 一词一起被移除。

-f

强制使用 *new_device*，即使它显示为正在使用。并非所有设备都可以通过这种方式覆盖。

zpool scrub [-s] pool ...

开始一次清理。scrub 将检查指定池中的所有数据以验证其校验和是否正确。对于复制的（镜像或 raidz）设备，ZFS 会自动修复在清理过程中发现的所有损坏。zpool status 命令报告清理的进度，并在完成时汇总清理的结果。

清理和重新同步是非常相似的操作。不同之处在于，重新同步仅检查 ZFS 已知将要过期的数据（例如，在将新设备连接到某个镜像或者替换现有设备时），而清理将检查所有数据以发现由于硬件故障或磁盘故障而产生的无提示错误。

由于清理和重新同步都是 I/O 密集型操作，因此 ZFS 一次仅允许使用其中的一个。如果已经在进行清理了，则后续的 **zpool scrub** 将返回错误，设备将使用 **zpool scrub -s** 取消当前的清理。如果正在进行重新同步，ZFS 将不允许启动清理，直到重新同步完成为止。

-s

停止清理。

zpool set property=value pool

对指定池设置指定属性。有关可设置哪些属性和可接受的值的更多信息，请参见“属性”部分。

zpool split [-R altroot] [-n [-l]] [-o mntopts] [-o property=value] pool newpool [device ...]

从某个池中每个镜像的顶层 vdev 中分割出一个磁盘，然后使用分割出的磁盘新建一个池。原始池必须包含一个或多个镜像，并且不能处在重新同步过程中。split 子命令会选择每个镜像 vdev 中的最后一个设备，除非被命令行上的设备指定覆盖。

使用 *device* 参数时，split 会将指定的设备包括在新池中；如果有任何设备仍处于未指定状态，则会将每个镜像 vdev 中的最后一个设备分配到该池，就像其正常操作时一样。如果不确定 split 命令的结果，请使用 -n（“预运行”）选项确保命令会产生您想要的结果。

-R altroot

在分割后自动导入新创建的池，将指定的 altroot 参数用于新池的备用根。请参见前面“属性”部分中的 altroot 说明。

-n

显示在不实际分割池的情况下会创建的配置。实际的池分割操作仍可能因特权不足或设备状态的原因而失败。

-l

如果可能，让 -n 以当前 /dev/chassis 位置格式显示配置。

-o *mntopts*

挂载池中的数据集中要使用的挂载选项列表（用逗号分隔）。有关数据集属性和挂载选项的说明，请参见 [zfs\(1M\)](#)。仅在与 **-R** 选项结合时有效。

-o *property=value*

对新池设置指定的属性。有关可用池属性的更多信息，请参见前面的“属性”部分。

zpool status [-l] [-v] [-x] [-T d|u] [*pool*] ... [*interval* [*count*]]

显示指定池的详细运行状态。如果未指定任何 *pool*，将显示系统中每个池的状态。有关池和设备运行状况的更多信息，请参见“设备故障和恢复”部分。

如果指定了时间间隔，将每隔 *interval* 秒就会显示状态和空间使用情况，直到按 Ctrl-C 为止。如果指定了 *count*，该命令将在显示 *count* 报告后退出。

如果正在清理或重新同步，此命令将报告完成百分比以及估计的完成时间。这两项数据都只是近似值，因为池中的数据量以及系统上的其他工作负荷可能会发生变化。

-l

如果可能，以当前 `/dev/chassis` 位置格式显示 *vdev* 状态。

-x

仅显示呈现错误或者因其他原因不可用的池的状态。

-v

显示详细数据错误信息，输出自上次完成池清理以来所有数据错误的完整列表。

-T d|u

显示时间戳。

指定 **d** 表示标准日期格式。请参见 [date\(1\)](#)。指定 **u** 表示时间的内部表示形式的印刷表示形式。请参见 [time\(2\)](#)。

zpool upgrade

标识池的盘上版本，该版本确定当前正在运行的软件版本中的可用池功能。可以继续使用较旧的池版本，但有些功能可能不可用。可以通过使用 `zpool upgrade -a` 命令升级池。无法在运行较早软件版本的系统上访问较新版本的池。

zpool upgrade -v

显示当前软件支持的 ZFS 池版本。当前 ZFS 池版本以及所有以前支持的版本都会显示，还会显示随每个版本提供的功能的说明。

zpool upgrade [-V *version*] -a | *pool* ...

将指定的池升级到最新的盘上版本。如果此命令显示某个池已过期，随后可使用 `zpool upgrade -a` 命令来升级该池。升级后的池在运行较早软件版本的系统上将不可访问。

-a

升级所有池。

-v *version*

升级到指定版本，该版本必须高于当前版本。如果不指定 -v 标志，则会将池升级到最新的版本。

示例

示例1 创建 RAID-Z 存储池

以下命令创建一个池，其中包含由六个磁盘组成的单个 raidz 根 *vdev*。

```
# zpool create tank raidz c0t0d0 c0t1d0 c0t2d0 c0t3d0 c0t4d0 c0t5d0
```

示例2 创建镜像存储池

以下命令创建包含两个镜像的池，其中每个镜像各包含两个磁盘。

```
# zpool create tank mirror c0t0d0 c0t1d0 mirror c0t2d0 c0t3d0
```

或者，可以使用描述磁盘当前位置的 `/dev/chassis` 路径指定整个磁盘。

```
# zpool create tank \
    mirror \
        /dev/chassis/RACK29.U01-04/DISK_00/disk \
        /dev/chassis/RACK29.U05-08/DISK_00/disk \
    mirror \
        /dev/chassis/RACK29.U01-04/DISK_01/disk \
    /dev/chassis/RACK29.U05-08/DISK_01/disk
```

示例3 使用分片创建 ZFS 存储池

以下命令使用两个磁盘分片创建一个取消镜像的池。

```
# zpool create tank /dev/dsk/c0t0d0s1 c0t1d0s4
```

或者，可以使用描述磁盘当前位置的 `/dev/chassis` 路径指定整个磁盘。

```
# zpool add tank mirror \
    /dev/chassis/RACK29.U01-04/DISK_02/disk \
    /dev/chassis/RACK29.U05-08/DISK_02/disk
```

示例4 将镜像添加到 ZFS 存储池

以下命令将两个镜像磁盘添加到池 *tank* 中（假定该池已经由双向镜像组成）。增加的空间可立即供池中的任何数据集使用。

```
# zpool add tank mirror c1t0d0 c1t1d0
```

示例5 列出可用的 ZFS 存储池

以下命令列出系统上的所有可用池。

```
# zpool list
NAME      SIZE  ALLOC   FREE    CAP  DEDUP  HEALTH  ALTROOT
```

示例5 列出可用的 ZFS 存储池 (续)

```
pool      136G   109M   136G      0%   3.00x  ONLINE  -
rpool    67.5G  12.6G  54.9G     18%   1.01x  ONLINE  -
```

示例6 列出池的所有属性。
以下命令列出某个池的所有属性。

```
% zpool get all pool
NAME  PROPERTY      VALUE      SOURCE
pool  size         464G      -
pool  capacity     0%        -
pool  altroot      -         default
pool  health       ONLINE    -
pool  guid         9893442355858802973  default
pool  version      23        default
pool  bootfs       -         default
pool  delegation   on        default
pool  autoreplace  off       default
pool  cachefile    -         default
pool  failmode     wait      default
pool  listsnapshots off       default
pool  autoexpand   off       default
pool  dedupditto   0         default
pool  dedupratio   1.00x     -
pool  free         464G      -
pool  allocated    130K      -
```

示例7 销毁 ZFS 存储池
以下命令销毁池 *tank* 以及其中包含的所有数据集。

```
# zpool destroy -f tank
```

示例8 导出 ZFS 存储池
以下命令导出池 *tank* 中的设备，使这些设备可以重新定位，或在以后导入。

```
# zpool export tank
```

示例9 导入 ZFS 存储池
以下命令显示可用的池，然后导入池 *tank* 供系统使用。
此命令的结果与以下内容相似：

```
# zpool import
pool: tank
id: 7678868315469843843
state: ONLINE
action: The pool can be imported using its name or numeric identifier.
```

示例9 导入 ZFS 存储池 (续)

config:

```
tank    ONLINE
mirror-0 ONLINE
c1t2d0  ONLINE
c1t3d0  ONLINE
```

```
# zpool import tank
```

示例10 将所有 ZFS 存储池升级到当前版本

以下命令将所有 ZFS 存储池升级到软件的当前版本。

```
# zpool upgrade -a
```

This system is currently running ZFS pool version 22.

All pools are formatted using this version.

示例11 管理热备件

以下命令创建一个包含可用热备件的新池：

```
# zpool create tank mirror c0t0d0 c0t1d0 spare c0t2d0
```

如果其中的一个磁盘发生故障，池会缩减至降级状态。可以使用以下命令替换故障设备：

```
# zpool replace tank c0t0d0 c0t3d0
```

设备重新同步后，将自动分离备件，并使其在其他设备发生故障时可用。使用以下命令可将热备件从池中永久移除：

```
# zpool remove tank c0t2d0
```

示例12 创建具有单独的镜像日志设备的 ZFS 池

以下命令创建一个 ZFS 存储池，该池包含两个双向镜像和镜像日志设备。

```
# zpool create pool mirror c0d0 c1d0 mirror c2d0 c3d0 log mirror \
    c4d0 c5d0
```

示例13 将高速缓存设备添加到 ZFS 池中

以下命令添加两个磁盘，用作 ZFS 存储池的高速缓存设备：

```
# zpool add pool cache c2d0 c3d0
```

添加以后，高速缓存设备会逐渐被主内存中的内容填满。填充它们可能需要一个多小时的时间，具体取决于高速缓存设备的大小。使用 `iostat` 选项可以监视容量和读取操作，如下所示：

示例 13 将高速缓存设备添加到 ZFS 池中 (续)

```
# zpool iostat -v pool 5
```

示例 14 移除镜像日志设备
指定下面所示的配置后，以下命令即删除池 tank 中的镜像日志设备 mirror-2。

```
pool: tank
state: ONLINE
scrub: none requested
config:

    NAME      STATE    READ WRITE CKSUM
    tank      ONLINE    0     0     0
        mirror-0  ONLINE    0     0     0
            c6t0d0  ONLINE    0     0     0
            c6t1d0  ONLINE    0     0     0
        mirror-1  ONLINE    0     0     0
            c6t2d0  ONLINE    0     0     0
            c6t3d0  ONLINE    0     0     0
    logs
        mirror-2  ONLINE    0     0     0
            c4t0d0  ONLINE    0     0     0
            c4t1d0  ONLINE    0     0     0

# zpool remove tank mirror-2
```

示例 15 恢复发生故障的 ZFS 池
如果某个池发生了可恢复的故障，在该池已高速缓存（请参见前面的 cache file）时，zpool status 将提供指示此状态的消息，或者在因为池的 zpool import 失败而显示的错误输出中提供。

使用 zpool clear 命令恢复高速缓存的池：

```
# zpool clear -F data
Pool data returned to its state as of Tue Sep 08 13:23:35 2009.
Discarded approximately 29 seconds of transactions.
```

如果池配置未高速缓存，请使用带恢复模式标志的 zpool import：

```
# zpool import -F data
Pool data returned to its state as of Tue Sep 08 13:23:35 2009.
Discarded approximately 29 seconds of transactions.
```

示例 16 导入缺少日志设备的 ZFS 池
以下示例说明尝试导入缺少日志设备的池。-m 选项用于完成导入操作。
尽管无法确定其他设备的准确配置，但知道它们是该池的一部分。

示例 16 导入缺少日志设备的 ZFS 池 (续)

```
# zpool import tank
The devices below are missing, use '-m' to import the pool anyway:
    c5t0d0 [log]

cannot import 'tank': one or more devices is currently unavailable

# zpool import -m tank
# zpool status tank
    pool: tank
    state: DEGRADED
status: One or more devices could not be opened. Sufficient replicas
exist for
    the pool to continue functioning in a degraded state.
action: Attach the missing device and online it using 'zpool online'.
    see: http://www.sun.com/msg/ZFS-8000-2Q
    scan: none requested
config:

        NAME                                STATE      READ WRITE CKSUM
        tank                                DEGRADED    0     0     0
            c7t0d0                        ONLINE      0     0     0
        logs
            1693927398582730352 UNAVAIL      0     0     0 was
/dev/dsk/c5t0d0

errors: No known data errors
```

以下示例说明如何导入缺少镜像日志设备的池。

```
# zpool import tank
The devices below are missing, use '-m' to import the pool anyway:
mirror-1 [log]
c5t0d0
c5t1d0

# zpool import -m tank

# zpool status tank
    pool: tank
    state: DEGRADED
status: One or more devices could not be opened. Sufficient replicas
exist for the pool to continue functioning in a degraded state.
action: Attach the missing device and online it using 'zpool online'.
    see: http://www.sun.com/msg/ZFS-8000-2Q
    scan: none requested
config:
```

示例 16 导入缺少日志设备的 ZFS 池 (续)

```
NAME                                STATE    READ WRITE CKSUM
tank                                DEGRADED    0    0    0
  c7t0d0                            ONLINE      0    0    0
  logs
    mirror-1                        UNAVAIL      0    0    0
insufficient replicas
  46385995713041169                UNAVAIL      0    0    0 was
/dev/dsk/c5t0d0
  13821442324672734438            UNAVAIL      0    0    0 was
/dev/dsk/c5t1d0

errors: No known data errors
```

示例 17 通过特定路径导入池
以下命令通过标识 tank 池的特定设备路径（在本例中为 /dev/dsk/c9t9d9 和 /dev/dsk/c9t9d8）来导入该池。

```
# zpool import -d /dev/dsk/c9t9d9s0 /dev/dsk/c9t9d8s0 tank
```

现有的限制在于，即使该池由完整的磁盘组成，该命令也必须包括特定设备的分片标识符。

退出状态

将返回下列退出值：

- 0 成功完成。
- 1 出现错误。
- 2 指定的命令行选项无效。

属性

有关以下属性的说明，请参见 [attributes\(5\)](#)：

属性类型	属性值
可用性	system/file-system/zfs
接口稳定性	Committed（已确定）

另请参见

[zfs\(1M\)](#)、[attributes\(5\)](#)

引用名	zstreamdump – 过滤 zfs send 流中的数据
用法概要	zstreamdump [-C] [-v]
描述	zstreamdump 实用程序从 zfs send 命令的输出中读取数据，然后显示该输出中的头和某些统计信息。请参见 zfs(1M) 。
选项	所支持的选项如下： -C 禁用校验和的验证。 -v 详细模式。转储所有标头，而不仅仅是开始头和结束头。
属性	有关以下属性的说明，请参见 attributes(5) ：

属性类型	属性值
可用性	system/file-system/zfs
接口稳定性	Uncommitted（未确定）

另请参见 [zfs\(1M\)](#)、[attributes\(5\)](#)

