



Solaris Common Messages and Troubleshooting Guide

Sun Microsystems, Inc.
901 San Antonio Road
Palo Alto, CA 94303-4900
U.S.A.

Part Number 806-1075-10
February 2000

Copyright 2000 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, California 94303-4900 U.S.A. All rights reserved.

This product or document is protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or document may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any. Third-party software, including font technology, is copyrighted and licensed from Sun suppliers.

Parts of the product may be derived from Berkeley BSD systems, licensed from the University of California. UNIX is a registered trademark in the U.S. and other countries, exclusively licensed through X/Open Company, Ltd.

Sun, Sun Microsystems, the Sun logo, docs.sun.com, AnswerBook, AnswerBook2, SunOS, SunPC, Java, Java Workshop, NFS, and Solaris are trademarks, registered trademarks, or service marks of Sun Microsystems, Inc. in the U.S. and other countries. All SPARC trademarks are used under license and are trademarks or registered trademarks of SPARC International, Inc. in the U.S. and other countries. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc. OpenGL is a registered trademark of Silicon Graphics, Inc.

The OPEN LOOK and Sun™ Graphical User Interface was developed by Sun Microsystems, Inc. for its users and licensees. Sun acknowledges the pioneering efforts of Xerox in researching and developing the concept of visual or graphical user interfaces for the computer industry. Sun holds a non-exclusive license from Xerox to the Xerox Graphical User Interface, which license also covers Sun's licensees who implement OPEN LOOK GUIs and otherwise comply with Sun's written license agreements.

RESTRICTED RIGHTS: Use, duplication, or disclosure by the U.S. Government is subject to restrictions of FAR 52.227-14(g)(2)(6/87) and FAR 52.227-19(6/87), or DFAR 252.227-7015(b)(6/95) and DFAR 227.7202-3(a).

DOCUMENTATION IS PROVIDED "AS IS" AND ALL EXPRESS OR IMPLIED CONDITIONS, REPRESENTATIONS AND WARRANTIES, INCLUDING ANY IMPLIED WARRANTY OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE OR NON-INFRINGEMENT, ARE DISCLAIMED, EXCEPT TO THE EXTENT THAT SUCH DISCLAIMERS ARE HELD TO BE LEGALLY INVALID.

Copyright 2000 Sun Microsystems, Inc. 901 San Antonio Road, Palo Alto, Californie 94303-4900 Etats-Unis. Tous droits réservés.

Ce produit ou document est protégé par un copyright et distribué avec des licences qui en restreignent l'utilisation, la copie, la distribution, et la décompilation. Aucune partie de ce produit ou document ne peut être reproduite sous aucune forme, par quelque moyen que ce soit, sans l'autorisation préalable et écrite de Sun et de ses bailleurs de licence, s'il y en a. Le logiciel détenu par des tiers, et qui comprend la technologie relative aux polices de caractères, est protégé par un copyright et licencié par des fournisseurs de Sun.

Des parties de ce produit pourront être dérivées du système Berkeley BSD licenciés par l'Université de Californie. UNIX est une marque déposée aux Etats-Unis et dans d'autres pays et licenciée exclusivement par X/Open Company, Ltd.

Sun, Sun Microsystems, le logo Sun, docs.sun.com, AnswerBook, AnswerBook2, SunOS, SunPC, Java, Java Workshop, NFS, et Solaris sont des marques de fabrique ou des marques déposées, ou marques de service, de Sun Microsystems, Inc. aux Etats-Unis et dans d'autres pays. Toutes les marques SPARC sont utilisées sous licence et sont des marques de fabrique ou des marques déposées de SPARC International, Inc. aux Etats-Unis et dans d'autres pays. Les produits portant les marques SPARC sont basés sur une architecture développée par Sun Microsystems, Inc. OpenGL est une marque dpose de Silicon Graphics, Inc.

L'interface d'utilisation graphique OPEN LOOK et Sun™ a été développée par Sun Microsystems, Inc. pour ses utilisateurs et licenciés. Sun reconnaît les efforts de pionniers de Xerox pour la recherche et le développement du concept des interfaces d'utilisation visuelle ou graphique pour l'industrie de l'informatique. Sun détient une licence non exclusive de Xerox sur l'interface d'utilisation graphique Xerox, cette licence couvrant également les licenciés de Sun qui mettent en place l'interface d'utilisation graphique OPEN LOOK et qui en outre se conforment aux licences écrites de Sun.

CETTE PUBLICATION EST FOURNIE "EN L'ETAT" ET AUCUNE GARANTIE, EXPRESSE OU IMPLICITE, N'EST ACCORDEE, Y COMPRIS DES GARANTIES CONCERNANT LA VALEUR MARCHANDE, L'APTITUDE DE LA PUBLICATION A REpondre A UNE UTILISATION PARTICULIERE, OU LE FAIT QU'ELLE NE SOIT PAS CONTREFAISANTE DE PRODUIT DE TIERS. CE DENI DE GARANTIE NE S'APPLIQUERAIT PAS, DANS LA MESURE OU IL SERAIT TENU JURIDIQUEMENT NUL ET NON AVENU.



Contents

Preface

- 1. About Error Messages 27**
 - Searching for Messages 27
 - Choosing What to Look For 27
 - In the Printed Book 28
 - In the AnswerBook Navigator 28
 - Combining Search Techniques 30
 - Understanding the Message Explanations 30
- 2. Alphabetical Message Listing 33**
 - Numbers and Symbols 33
 - ***** FILE SYSTEM WAS MODIFIED ***** 33
 - ** Phase 1- Check Blocks and Sizes 34
 - ** Phase 1b- Rescan For More DUPS 34
 - ** Phase 2- Check Pathnames 35
 - ** Phase 3- Check Connectivity 35
 - ** Phase 4- Check Reference Counts 36
 - ** Phase 5- Check Cyl groups 36
 - @@ 37
 - 29a00 illegal instruction 38

451 timeout waiting for input during *source* 38

501 MAIL FROM: unrecognized address: @@*hostname* 39

550 *hostname*... Host unknown 40

550 Security server failed to perform requested command 41

550 *username*... User unknown 42

554 *hostname*... Local configuration error 42

"A" 43

A command window has exited because its child exited. 43

access violation unknown host IP address 43

Accessing a corrupted shared library 44

Address already in use 44

Address family not supported by protocol family 44

admintool: Received communication service error 4 45

Advertise error 45

answerbook: XView error: NULL pointer passed to xv_set 46

apdb: Resource temporarily unavailable 46

Arg list too long 47

Argument out of domain 48

Arguments too long 48

assertion failed: *string*, file *name*, line *int* 49

Attempting to link in more shared libraries than system limit 49

automount[*int*]: *name*: Not a directory 49

automountd[*int*]: server *hostname* responding 50

"B" 50

Bad address 50

BAD/DUP FILE I=i OWNER=o MODE=m SIZE=s MTIME=t CLEAR?
51

Bad file number 51

block no. BAD I=*inode no.* 52
BAD_MESSAGE (error code 100) from X.400 52
bad module/chip at: *position* 53
Bad request descriptor 53
BAD SUPER BLOCK: *string* 53
BAD TRAP 55
/bin/sh: *file*: too big 55
Block device required 56
Boot device: /iommu/sbus/*directory/directory/sd@3,0* 56
Broadcast Message from root (*pts/int*) on server [date] 57
Broken pipe 57
Bus Error 58
"C" 58
Cannot access a needed shared library 58
Cannot allocate colormap entry for "*string*" 59
Cannot assign requested address 59
Cannot bind to domain *domainname*: can't communicate with ypbind 59
Cannot boot after install, error that points to an .rc file 60
cannot change passwd, not correct passwd 61
cannot establish nfs service over /dev/tcp: transport setup problem 61
Cannot exec a shared library directly 62
Cannot find SERVER *hostname* in network database 62
cannot install bootblock 63
Cannot open FCC file 64
Cannot send after transport endpoint shutdown 64
can't communicate with ypbind 65
Can't create public message device (Device busy) 65
Can't invoke /etc/init, error *int* 66

can't open /dev/rdisk/string: (null): UNEXPECTED
INCONSISTENCY 67

can't synchronize with hayes 67

cd: Too many arguments 68

Channel number out of range 68

chmod: ERROR: invalid mode 68

Command not found 69

Communication error on send 70

config error: mail loops back to myself. 70

connect from *hostIP* to callit(ypserv): request from non-local host 71

connect from *hostIP* to callit(ypserv): request from unauthorized host 71

connect from *hostIP* to callit(ypserv): request from unprivileged port 72

connect from *hostIP* to callit(ypserv): request not forwarded 72

Connection closed. 72

Connection closed by foreign host. 73

[Connection closed. Exiting] 73

Connection refused 73

Connection reset by peer 74

Connection timed out 74

console login: ^J^M^Q^K^K^P 75

core dumped 76

corrupt label - wrong magic number or corrupt label or corrupt label -
label checksum failed 77

could not grant slave pty 78

Could not initialize tooltalk (tt_open): TT_ERR_NOMP 78

Could not open ToolTalk Channel 79

Could not start new viewer 80

Could not start NFS service for any protocol. Exiting 80

cpio: Bad magic number/header. 81

cpio : can't read input : end of file encountered prior to expected end of archive. 81

Cross-device link 82

"D" 82

data access exception 82

Data fault 83

Deadlock situation detected/avoided 83

Destination address required 84

destination component full 84

/dev/fd/int: /dev/fd/int: cannot open 84

/dev/rdisk/c0t6d0s2: No such file or directory 85

Device busy 86

device busy 86

/dev/rdisk/string: CAN'T CHECK FILE SYSTEM. 88

/dev/rdisk/string: UNEXPECTED INCONSISTENCY; RUN fsck MANUALLY. 88

Directory not empty 89

Disc quota exceeded 89

disk does not appear to be prepared for encapsulation 90

diskN not unique 90

dlopen (libxfs.so) failed 91

driver is already installed 91

dtmail: cannot open mailfile on 2.5.1 /var/mail server 92

DUMP: Cannot open dump device '/dev/rdisk/c2t0d0s1': Permission denied 93

dumpton: Cannot open '/dev/rmt/string': Device busy 93

DUP/BAD I=i OWNER=o MODE=m SIZE=s MTIME=t FILE=f REMOVE? 94

int DUP I=int 94

"E" 95

Encapsulation of root disk is not supported on systems with old PROM versions 95

ENOMEM The available data space is not large enough to accommodate the shared memory segment 96

error 15 initializing 97

Error 76 97

Error 88 97

error code 2: access violation 98

error: DPS has not initialized or server connection failed 98

Error: Error adding OS service Solaris 2.6 sparc sun4u: 99

Error Host Unknown: 100

ERROR: missing file arg (cm3) 101

ERROR [SCCS/*s.string*]: '*SCCS/p.string*' nonexistent (ut4) 101

ERROR [SCCS/*s.string*]: writable '*string*' exists (ge4) 101

Error: you don't have a license to run this program 102

esp0: data transfer overrun 102

ether_hostton errors from cb_reset 102

Event not found 103

EXCESSIVE BAD BLKS I=*int* CONTINUE? 104

EXCESSIVE DUP BLKS I=*int* CONTINUE? 104

Exec format error 105

"F" 106

failed to initialize adapter 106

Failed to Load Security Policy: Invalid argument 106

fast access mmu error 107

fbconsole: ioctl SRIOCSREDIR: Device Busy. 108

fd0: unformatted diskette or no diskette in the drive 109

File descriptor in bad state 109

File exists 109
File locking deadlock 110
File name too long 110
file system full 111
FILE SYSTEM STATE IN SUPERBLOCK IS WRONG; FIX? 111
File table overflow 112
File too large 112
filemgr: mknod: Permission denied 113
FREE BLK COUNT(S) WRONG IN SUPERBLK SALVAGE? 113
fsck & ufsdump - cannot read block/sector errors 114
fsck: Can't open /dev/dsk/*string* 115
fsck: Can't stat /dev/dsk/*string* 115
ftp: ftp/tcp: unknown service 116
fw_ipinput: q fc5fddc0:illegal interface 116
FW1: log message queue is full 117
fwm: no license 118
fwskip_parse_headers: invalid peer n 118
"G" 119
giving up 119
Graphics Adapter device /dev/fb is of unknown type 120
group.org_dir: NIS+ servers unreachable 120
"H" 121
hang console 121
/home/*string*: No such file or directory 121
Host is down 122
host name configuration error 123
hosts.org_dir: NIS+ servers unreachable 123
"I" 124

I can't read your attachments. What mailer are you using? 124

Identifier removed 124

ie0: Ethernet jammed 125

ie0: no carrier 125

If pipe/FIFO, don't sleep in stream head 126

ifconfig: bad address 126

ifconfig bad address le0 127

ifconfig: *host name* bad space address 127

ifconfig: SIOCGIFFLAGS: hme0: no such interface 128

Illegal Instruction 128

Illegal instruction "0xhex" was encountered at PC 0xhex 129

Illegal seek 129

Image Tool: Unable to open XIL Library. 130

Inappropriate ioctl for device 130

INCORRECT BLOCK COUNT I=*int* (should be *int*) CORRECT? 131

index failed:full:index preceded by saveset name 131

inetd[*int*]: execv /usr/sbin/in.uucpd: No such file or directory 132

inetd[*int*]: *string*/tcp: unknown service 132

inetd[*int*]: *string*/udp: unknown service 133

inetd: Too many open files 133

INIT: Cannot create /var/adm/utmpx 133

InitOutput: Error loading module for /dev/fb 134

insertion failed: a problem with the filesystem has been detected:
filesystem is probably full 135

Interrupted system call 135

Invalid argument 136

Invalid null command 136

Invalid_SS_JWS_HOME:no C:\\lib\\basicframe.properties 137

I/O error 137

IP: Hardware address '08:00:20:xx:xx:xx' trying to be our address
xxx.xxx.xxx.xxx! 138

Is a directory 140

"J" 140

java.lang.UnsatisfiedLinkError: 140

"K" 141

kernel read error 141

killed 141

Killed 142

kmem_free block already free 142

"L" 143

last message repeated *int* times 143

late initialization error 143

ld.so.1 fatal: can't set protection on segment 144

ld.so.1: *string*: fatal: *string*: can't open file: errno=2 144

ld.so.1: *string*: fatal: *string*: open failed: No such file or directory 145

ld.so.1: *string*: fatal: relocation error: *string*: *string*: referenced symbol not
found 146

ld.so.1: *string*: fatal: relocation error: symbol not found: *string* 146

le0: Memory error! 147

le0: No carrier- cable disconnected or hub link test disabled? 148

le0: No carrier- transceiver cable problem? 148

level 15 interrupt 149

.lib section in a.out corrupted 149

LINK COUNT FILE I=i OWNER=o MODE=m SIZE=s MTIME=t
COUNT... ADJUST? 149

Link has been severed 150

LL105W: Protocol error detected. 150

ln: cannot create /dev/fb: Read-only filesystem 151
 lockd[*int*]: create_client: no name for inet address 0x*hex* 151
 log_get: len is not a multiple of 4 from FW-1 152
 Login incorrect 152
 lp hang 153
 "M" 154
 Machine is not on the network 154
 Mail Tool is confused about the state of your Mail File. 154
 mail: Your mailfile was found to be corrupted (Content-length mismatch). 155
 mailtool: Can't create dead letter: Permission denied 155
 mailtool: Could not initialize the Classing Engine 156
 Management Server is VPN while client is NON-VPN 156
 file name may contain holes - can't swap on it. 157
 mbuf map full 157
 Memory address alignment 157
 memory leaks 158
 Message too long 158
 mount: /dev/dsk/*string* is already mounted, /*string* is busy, or... 159
 mount: giving up on: /*string* 159
 mount: mount-point /*string* does not exist. 160
 mount: the state of /dev/dsk/*string* is not okay 160
 Multihop attempted 160
 "N" 161
 Name not unique on network 161
 named [*pid*]: *hostname.domainname* has CNAME and other data (illegal) 161
 /net/*string*: No such file or directory 162
 Network dropped connection because of reset 163

Network is down 163

Network is unreachable 163

NFS getattr failed for server *string*: RPC: Timed out 164

nfs mount: Couldn't bind to reserved port 165

nfs mount: mount: *string*: Device busy 165

NFS mount: /*string* mounted OK 165

NFS mounted callog file Unsupported. 166

NFS read failed for server *string* 167

nfs_server: bad getargs for *int/int* 167

NFS server *string* not responding still trying 168

NFS server *string* ok 168

NFS *string* failed for server *string*: error *int (string)* 169

nfs umount: *string*: is busy 169

NFS write error on host *string*: No space left on device. 169

NFS write failed for server *string*: RPC: Timed out 170

NIS+ authentication failure 170

nis_cachmgr: Error in reading NIS cold start file : '/var/nis/
NIS_COLD_START' 171

No buffer space available 171

No child processes 172

No default media available 173

No directory! Logging in with home=/ 173

No message of desired type 174

No recipients specified 174

No record locks available 175

No route to host 175

No shell Connection closed 176

No space left on device 176

No such device 177
No such device or address 177
No such file or directory 178
no such map in server's domain 178
No such process 179
No such user as *string*- cron entries not created 179
No utmpx entry 179
no valid fm license 180
no VTOC 181
Not a data message 182
Not a directory 182
Not a stream device 182
Not enough space 183
not found 184
Not login shell 184
Not on system console 185
Not owner 185
Not supported 185
NOTICE: /*string*: out of inodes 186
NOTICE: vxvm: unexpected status on close 187
nsrck: SYSTEM error, more space needed to compress [*client*] index, 8.1 MB required 188
"O" 189
Object is remote 189
ok 189
open: no such device or address from FW-1 190
Operation already in progress 190
Operation canceled 191

operation failed [error 185], unknown group error 0, *string* 191

Operation not applicable 191

Operation not supported on transport endpoint 192

Operation now in progress 192

/opt/bin/jws: /solaris/bin/locate_dirs: not found 193

Option not supported by protocol 193

out of memory 193

Out of stream resources 194

overlapping swap volume 194

"P" 195

Package not installed 195

page_create: invalid flag 195

Panic 195

panic -boot: Could not mount filesystem 196

Panic on cpu 0: valloc'd past tmpptes 198

PARTIALLY ALLOCATED INODE I=*int* CLEAR? 199

passwd: Changing password for *string* 199

passwd (SYSTEM): System error: repository out of range 200

passwd.org_dir: NIS+ servers unreachable 201

Password does not decrypt secret key for unix.uid@*string* 201

password file busy - try again later. 201

pdbadmin start node fails cluster_establish join not allowed 202

Permission denied 203

Please specify a recipient. 203

Protocol error 203

protocol error, *string* closed connection 204

Protocol family not supported 204

Protocol not supported 205

Protocol wrong type for socket 205

"Q" 206

quotactl: open Is a directory 206

"R" 206

Read error from network: Connection reset by peer 206

Reading configuration data 207

Read-only file system 207

rebooting... 208

Recipient names must be specified 208

refused connect from *hostIP* to callit(ypserv) 209

Reset tty pgrp from *int* to *int* 209

Resource temporarily unavailable 209

Restartable system call 210

Result too large 210

rlogin: no directory! connection closed 211

rmdir: *string*: Directory not empty 211

ROOT LOGIN /dev/console 212

ROOT LOGIN /dev/pts/*int* FROM *string* 212

route: socket: Protocol not supported 212

RPC: Program not registered 213

rx framing error 213

"S" 214

save: SYSTEM error, Arg list too long 214

SCSI bus DATA IN phase parity error 214

SCSI transport failed: reason 'reset' 215

Security exception on host *string*: USER ACCESS DENIED. 216

Segmentation Fault 216

sendmail[]: can't lookup data via name server "dns" or sendmail[]: can't
lookup data via name server "nis" 217

sendmail[init]: NOQUEUE: SYSERR(root): Cannot bind to domain
<domain>: no such map in server's domain: Bad file number 218

sendmail[int]: NOQUEUE: SYSERR: net hang reading from *string* 218

Service wouldn't let us acquire selection 219

setmnt: Cannot open /etc/mnttab for writing 219

share_nfs: /home: Operation not applicable 219

Signal 8 error 220

SIMS license error: licenses invalid 220

Slice c0t1d0s0 is too small to contain 1 replicas 221

snmpdx: bind() failed on udp on 161 [errno: address already in use] 125

snmpdx dmide: unable to connect to snmpdx 222

Socket type not supported 222

Soft error rate (*int*%) during writing was too high 222

Software caused connection abort 223

Srmount error 223

Stale NFS file handle 224

start up failure no such file or directory 224

statd: cannot talk to statd at *string* 224

stty: TCGETS: Operation not supported on socket 225

su: No shell 225

su: 'su root' failed for *login* on /dev/pts/*int* 226

su: 'su root' succeeded for *login* on /dev/pts/*int* 227

SunPC may NOT run correctly as root 227

syncing file systems... 228

syslog service starting. 228

System booting after fatal error FATAL 229

SYSTEM error, Arg list too long 229

system hang 230

SYSTEM HANGS DURING BOOT 230

system will not connect to port 80 230

"T" 230

tar: /dev/rmt/0: No such file or directory 230

tar: directory checksum error 231

tar: tape write error 231

Text file busy 232

Text is lost because the maximum edit log size has been exceeded. 232

tftpd: nak: Transport endpoint is already connected 233

THE FOLLOWING FILE SYSTEM(S) HAD AN UNEXPECTED INCONSISTENCY: 235

The SCSI bus is hung. Perhaps an external device is turned off. 235

THE SYSTEM IS BEING SHUT DOWN NOW !!! 236

The system will be shut down in *int* minutes 237

This gateway does not support Unix Password. 237

This mail file has been changed by another mail reader. 238

Timeout waiting for ARP/RARP packet 238

Timer expired 239

token ring hangs 239

Too many links 239

Too many open files 240

Transport endpoint is already connected 240

Transport endpoint is not connected 241

TRAP 3E 241

"U" 242

ufsdump 4mm commands 242

umount: warning: /*string* not in mnttab 243

Unable to connect to license server. Inconsistent encryption code. 243

unable to get pty! 243

Unable to install/attach driver '*string*' 244

Unable to open nwrecover, Error: nwrecover: NSR: please start a server on *client_name* 244

uname: error writing name when booting 245

undefined control 245

unknown host exception: unknown host 245

Unmatched ' 246

UNREF FILE I=i OWNER=o MODE=m SIZE=s MTIME=t CLEAR? 246

UnsatisfiedLinkError 247

Use "logout" to logout. 247

user unknown 248

/usr/dt/bin/rpc.ttdbserverd:Child Status' changed 248

/usr/openwin/bin/xinit: connection to X server lost 249

/usr/ucb/cc: language optional software package not installed 249

UX: userdel: error: Cannot update system files login cannot be deleted 251

"V" 251

Value too large for defined data type 251

Volume Manager reports error: Configuration daemon can't speak protocol version 252

Volume too large for defined data type 253

vxconfigd error: segmentation fault 253

vxfs filesystems not mounting 254

vxvm:vx slicer:ERROR unsupported disk layout 254

"W" 255

WARNING: add_spec: No major number for sf 255

warning:cacheofs:invalid cache version 256

WARNING: Clock gained *int* days- CHECK AND RESET THE DATE! 256

Warning: Could not find matching rule in rules.ok 257

WARNING: FAN FAILURE check if fans are still spinning 257

WARNING: FAN FAILURE still sensed 258

WARNING: No network locking on *string*: contact admin to install server change 258

WARNING: processor level 4 interrupt not serviced 259

WARNING: /tmp: File system full, swap space limit exceeded 259

WARNING: TOD clock not initialized- CHECK AND RESET THE DATE! 260

WARNING: Unable to repair the / filesystem. Run fsck 260

WARNING: vxvm:vxio: Illegal vminor encountered 261

Watchdog Reset 262

Who are you? 262

Window Underflow 263

"X" 264

X connection to *string*:0.0 broken (explicit kill or server shutdown). 264

xinit: not found 264

XIO: fatal IO error 32 (Broken pipe) on X server "*string*:0.0" 265

Xlib: connection to "*string*:0.0" refused by server 265

Xlib: extension "GLX" missing on display "0.0" 266

xntpd: clnt_dg_create: out of memory 267

xterm: fatal IO error 32 (Broken Pipe) or KillClient on X server "*string*:0.0" 267

XView warning: Cannot load font set '*string*' (Font Package) 268

"Y" 268

yp_all RPC clnt_call (transport level) failure 268

ypbind[*int*]: NIS server for domain "*string*" OK 269

`ypbind[int]`: NIS server not responding for domain "*string*"; still trying 269

`ypserv[int]`: restarting resolv server. old one not responding 270

`ypwhich`: can't communicate with `ypbind` 271

"Z" 271

`zsint`: silo overflow 271

Preface

System administrators and advanced users can use the *Common Messages and Troubleshooting Guide* to find explanations of some of the more common error messages in the Solaris™ operating environment.

Look up the messages and explanations here when you see a system message that you do not understand. If the message you are searching for is fairly common, it might be documented in this book.

How This Book Is Organized

Chapter 1 explains how to find messages in both the AnswerBook™ Navigator and in the printed book.

Chapter 2 lists messages alphabetically, with troubleshooting information following each message listing.

Ordering Sun Documents

Fatbrain.com, an Internet professional bookstore, stocks select product documentation from Sun Microsystems, Inc.

For a list of documents and how to order them, visit the Sun Documentation Center on Fatbrain.com at <http://www1.fatbrain.com/documentation/sun>.

Accessing Sun Documentation Online

The docs.sun.comSM Web site enables you to access Sun technical documentation online. You can browse the docs.sun.com archive or search for a specific book title or subject. The URL is `http://docs.sun.com`.

What Typographic Conventions Mean

The following table describes the typographic changes used in this book.

TABLE P-1 Typographic Conventions

Typeface or Symbol	Meaning	Example
AaBbCc123	The names of commands, files, and directories; on-screen computer output	Edit your <code>.login</code> file. Use <code>ls -a</code> to list all files. <code>machine_name%</code> you have mail.
AaBbCc123	What you type, contrasted with on-screen computer output	<code>machine_name%</code> su Password:
<i>AaBbCc123</i>	Command-line placeholder: replace with a real name or value	To delete a file, type rm <i>filename</i> .
<i>AaBbCc123</i>	Book titles, new words, or terms, or words to be emphasized.	Read Chapter 6 in <i>User's Guide</i> . These are called <i>class</i> options. You must be <i>root</i> to do this.

Shell Prompts in Command Examples

The following table shows the default system prompt and superuser prompt for the C shell, Bourne shell, and Korn shell.

TABLE P-2 Shell Prompts

Shell	Prompt
C shell prompt	machine_name%
C shell superuser prompt	machine_name#
Bourne shell and Korn shell prompt	\$
Bourne shell and Korn shell superuser prompt	#

About Error Messages

This book explains some of the more common error messages in the Solaris operating environment. Most messages explained here come from the operating system and the window system, but some come from commands, networking, and system administration (the *man pages section 1: User Commands* and *man pages section 1M: System Administration Commands*).

Searching for Messages

Choosing What to Look For

How you choose to look up a particular message depends on:

- How the message is constructed
- Whether you are searching in a printed book or in the AnswerBook documentation

Variable Words and Numbers

Remember as you are searching that some words and numbers in messages vary when the messages are displayed. For example, the following message uses the name of the server affected, `b5server` in this case:

```
NFS read failed for server b5server
```

When message words or numbers vary, this book uses the words *variable* and *number* in the italics type face. So the previous message is listed in this book as:

```
NFS read failed for server variable
```

Variable words and numbers can appear anywhere in a message, even at the beginning. Because of this, messages are alphabetized by the first nonreplaced word or number in the message.

Frequently Duplicated Parts of Messages

Many messages you see are actually combined messages, often beginning with a program name. The five error messages in the following example are basically the same, even though the command names are different.

- find: out of memory
- grep: out of memory
- ls: out of memory
- mount: out of memory
- fsck: out of memory

Rather than document this message at least five times, it appears in this book as the message “out of memory.” Messages that contain colons (:) are often combined messages, and you might find that explanations of message sections are available separately.

If you don't find the beginning of a message in the book, and the message contains colons, search for other parts of the message.

In the Printed Book

Methods for finding a particular message vary depending on whether you are looking at a printed book or are searching online with the AnswerBook Navigator.

To find a message in the printed book, you can search the table of contents (which is an alphabetical listing of the messages) or the main body of the manual, as shown in Chapter 2.

In the AnswerBook Navigator

While print search methods work in AnswerBook, too, you can search for messages through the search utility in the AnswerBook Navigator much faster.

1. Bring up AnswerBook.

```
$ answerbook
```
2. Click Select on the Search button.
3. Enter the words or pattern to search for in the “Search Library For:” pane.

4. Double-click Select on an entry in the resulting list. Although any of the entries might contain the information you are looking for, those from this book are most likely to be what you want.

If your first search does not find the message, consider altering the search pattern. Remember that this book contains only a small percentage of possible messages.

In general, you are most likely to find a documented message in the AnswerBook search pane when you enclose the searched-for words in quotation marks (“”) or in parentheses ().

Using Pattern Matching

You can search in the AnswerBook Navigator for text containing specific single words, phrases that contain spaces, words near one another, and word variations.

See “Using the AnswerBook Software” in *OpenWindows User’s Guide* for more detailed information about the AnswerBook search.

TABLE 1-1 AnswerBook Search Pattern Matching

To search for	Such as	Use
Single words	Installing, le0, group	The words
Phrases with spaces	Installing Packages	Quotation marks (“”) around the phrase
Words near one another	Installing...Server	Parentheses (())
Word variations	Delete, deleting, deletion	Asterisks (*) and hyphens (-)

Table 1-2 shows some of the possible matches for specific AnswerBook Navigator searches.

TABLE 1-2 AnswerBook Search Results

Searching With	Finds These (for Example)
Installing	Installing XIL Device Handlers (XIL Device Porting and Extensibility Guide) Installing Packages on a Server for... (Application Packaging Developer's Guide)
"Installing Packages"	Installing Packages on a Server for... (Application Packaging Developer's Guide) Installing Packages for Clients on a Server (Software and AnswerBook Packages...)
(Installing Server)	Installing Packages on a Server for... (Application Packaging Developer's Guide) Creating an Install Server (SPARC Installing Solaris Software)
Delet*	Delete All Silence (Solaris Advanced User's Guide) Deleting a Line (Solaris Advanced User's Guide) Deletion of the New Selection (OLIT Reference Manual)

Combining Search Techniques

Combine the above search techniques to further refine your search. For example, "chang* mail-tool" finds documents containing phrases such as "change mailtool," "change mail tool," "change mail-tool," "changing mailtool," and so on.

Understanding the Message Explanations

Each message in this book contains at least one of the following areas:

- *Cause*: What might have happened to cause the message
- *Action*: What you can do to fix the problem or, to continue with your work
- *Technical Notes*: Background information that might be interesting or helpful to a technical audience. This often contains information specifically for programmers.

Whenever you see part of a message that says “`errno=`” and then a number, look up the number on the `Intro(2)` man page to see what it indicates. System error messages on the `Intro(2)` man page are organized numerically.

- *See Also:* Suggests further reading

Alphabetical Message Listing

Messages listed here are in the current Solaris 8 operating environment. Many were also present in earlier Solaris operating environments.

Messages are listed alphabetically.

Numbers and Symbols

******* FILE SYSTEM WAS MODIFIED *******

Cause

This comment from the `fsck(1M)` command tells you that it changed the file system it was checking.

Action

If `fsck(1M)` was checking the root file system, reboot the system immediately to avoid corrupting the `/` partition. If `fsck(1M)` was checking a mounted file system, unmount that file system and run `fsck(1M)` again, so that work done by `fsck(1M)` is not undone when in-memory file tables are written out to disk.

**** Phase 1– Check Blocks and Sizes**

Cause

The `fsck(1M)` command is checking the file system shown in the messages that are displayed before this one. The first phase checks the inode list, finds bad or duplicate blocks, and verifies the inode size and format.

Action

If more than a dozen errors occur during this important phase, you might want to restore the file system from backup tapes. Otherwise, it is fine to proceed with `fsck(1M)`.

See Also

For more information, see the chapter on checking file system integrity in the *System Administration Guide, Volume 1*.

**** Phase 1b– Rescan For More DUPS**

Cause

The `fsck(1M)` command detected duplicate blocks while checking a file system, so `fsck(1M)` is rescanning the file system to find the inode that originally claimed that block.

Action

If `fsck(1M)` executes this optional phase, you will see additional DUP/BAD messages in phases 2 and 4.

See Also

For more information, see the chapter on checking file system integrity in the *System Administration Guide, Volume 1*.

**** Phase 2– Check Pathnames**

Cause

The `fsck(1M)` command is checking a file system, and `fsck(1M)` is now removing directory entries pointing to bad inodes that were discovered in phases 1 and 1b. This phase might ask you to remove files, salvage directories, fix inodes, reallocate blocks, and so on.

Action

If more than a dozen errors occur during this important phase, you might want to restore the file system from backup tapes. Otherwise it is fine to proceed with `fsck(1M)`.

See Also

For more information, see the chapter on checking file system integrity in the *System Administration Guide, Volume 1*.

**** Phase 3– Check Connectivity**

Cause

The `fsck(1M)` command is checking a file system, and `fsck(1M)` is now verifying the integrity of directories. You might be asked to adjust, create, expand, reallocate, or reconnect directories.

Action

You can usually answer "yes" to all these questions without harming the file system.

See Also

For more information, see the chapter on checking file system integrity in the *System Administration Guide, Volume 1*.

** Phase 4– Check Reference Counts

Cause

The `fsck(1M)` command is checking a file system, and `fsck(1M)` is now checking link count information obtained in phases 2 and 3. You might be asked to clear or adjust link counts.

Action

You can usually answer "yes" to all these questions without harming the file system.

See Also

For more information, see the chapter on checking file system integrity in the *System Administration Guide, Volume 1*.

** Phase 5– Check Cyl groups

Cause

The `fsck(1M)` command is checking a file system, and `fsck(1M)` is now checking the free-block and used-inode maps. You might be asked to salvage free blocks or summary information.

Action

You can usually answer "yes" to all these questions without harming the file system.

See Also

For more information, see the chapter on checking file system integrity in the *System Administration Guide, Volume 1*.

@@

Cause

This message is about how to fix the common @@token sendmail errors. There are instances when you receive email bounce messages because of syntax errors complaining that it does not know how to send email to @@token. Probably a site is NOT running NIS and is generating these errors or is talking to another site that is generating the errors and then passing the email on to your site. This happens because a single token is changed into a null ("") token. As a result, ruleset 3 (S3) changes null tokens into @@token. There are two key issues here. First, you do not want to be the host responsible for generating these errors, and, second, you do not want to pass along any errors that were generated by other hosts.

Action

To fix this problem, modify rules S3 and S22. (You'll only have S22, if using main.cf.) First, so you do not cause these errors, comment out the `invert aliases` rule in S22:

```
S22
R$*<@LOCAL>$*      $:$1
#R$-<@$->          $:$>3${Z$1@$2$}  invert aliases
R$*<@$+.$*>$*     @$1<@$2.$3>$4    already ok
R$+<@$+>$*        @$1<@$2.$m>$3   tack on our domain
R$+                @$1<@$w.$m>   tack on our full name
```

Next, so you do not pass on errors caused by other hosts, modify ruleset S3 from:

```
S3
# handle "from:<>"  special case
R$*<>$*  $@@      turn into magic token
```

To:

```
S3
# handle "from:<>"  special case
R$*<>$*  $@$n     turn into magic token
```

29a00 illegal instruction

Cause

When trying to boot a client from a boot/jumpstart server to install or upgrade a workstation, it fails with the following message:

```
boot net - install
Rebooting with command: net - install
Boot device: /iommu/sbus/ledma@f, 400010/le@f, 8c0000 File and args: -
install
29a00 Illegal Instruction
(0) ok
```

Action

The problem lies in the `/tftpboot` directory of the boot server. Confirm that the `HOSTID` and `HOSTID.ARCH` files are linked to the correct `inetboot.*` file for your architecture. The following is an example of how a symbolic link should look:

```
# cd /tftpboot
# ls -l 81971904*
81971904 -> inetboot.sun4m.Solaris_2.4
81971904.SUN4M -> inetboot.sun4m.Solaris_2.4
```

If the entries are not correct, remove the entry for the particular client in this directory, using `rm_install_client` or `rm_client` commands, and re-add the client with the `add_install_client(1M)` or `add_client` command or through `Solstice` giving the correct architecture.

451 timeout waiting for input during *source*

Cause

When `sendmail(1M)` reads from anything that might time out, such as an SMTP connection, it sets a timer to the value of the `r` processing option before reading begins. If the read does not complete before the timer expires, this message appears and reading stops. (Usually this happens during RCPT.) The mail message is then queued for later delivery.

Action

If you see this message often, increase the value of the `x` processing option in the `/etc/mail/sendmail.cf` file. If the timer is already set to a large number, look for hardware problems, such as poor network cabling or connections.

See Also

For more information about setting the timer, see the section describing the `sendmail(1M)` configuration options in the *System Administration Guide, Volume 3*. If you are using AnswerBook online documentation, the term "timeouts" is a good search string.

501 MAIL FROM: unrecognized address: @@hostname

Cause

A Sun machine running Sendmail 8.6 is used as a mailhost to send mail to the Internet in an environment that has MS Mail exchanger or a cc:Mail gateway. Mail from the MS exchange/cc:Mail gateway for the Internet is relayed to the mailhost, which actually delivers the mail. The mail from the Internet is accepted on the mailhost and forwarded to the MS exchanger/cc:mail gateway. The postmaster on the mailhost sees bounced messages with error messages, such as the following:

```
The original message was received at Thu, 29 May 1997 12:30:41 -0700
from artemis [206.189.46.3]

----- The following addresses had delivery problems -----
<Joe_Smith@cc.test.com> (unrecoverable error)

----- Transcript of session follows -----
... while talking to cc:
>>> MAIL From:<hermes>
>>> 501 MAIL FROM: unrecognized address: <hermes>
554 <Joe_Smith@cc.test.com> Remote protocol error
```

When analyzed, this mail turns out to be mail that has bounced from the Internet (for any reason) and was on its way back to the MS Exchange/cc:Mail gateway by the mailhost. The MS Exchange/cc:Mail gateway does not want to accept the mail because the "MAIL FROM:" address does not stick to the standards. `@@hostname` is an illegal SMTP address. Sendmail does not have a restriction on sender's address; however, other SMTP gateways, which need to translate the address to their native

address formats, are rather strict in adhering to the SMTP address format and would not accept the address in the `@@hostname` format.

Another situation: The user with `cc:Mail` sends mail to the Internet, and, due to one of many possible errors (user not found, host not found, and so forth), the message is sent back to the sender (bounces back). When a message is sent back, its recipient's address is replaced by the sender's address and the sender's address is erased (contains only "<>"). When the bounced sender's address goes through ruleset 3 and then 11 on the user's mail gateway (as it has to return it to the `cc:Mail` gateway, which is in the local domain => `mailer=ether`), it is transformed to `@@mail-gateway-name`.

Action

Insert the following line in the S11 ruleset after the line starting with `R$=D&`:

```
R@          $@mailer_daemon<@$w>          for @@hostname problem
```

After the insertion, S11 looks like this:

```
S11
R$*<@$+>$*      $1<@$2>$3          already ok
R$=D            $@$1<@$w>          tack on my hostname
R@             $@mailer_daemon<@$w> for @@hostname problem
R$+           $@$1<@$m>          tack on my mbox hostname
```

550 *hostname...* Host unknown

Cause

This `sendmail(1M)` message indicates that the destination host machine, specified by the portion of the address after the at-sign (`@`), was not found during domain naming system (DNS) lookup.

Action

Use the `nslookup(1M)` command to verify that the destination host exists in that or other domains, perhaps with a slightly different spelling. Failing that, contact the intended recipient and ask for a proper address.

Sometimes this return message indicates that the intended host is inoperable, rather than unknown. If a DNS record contains an unknown alternate host, and the primary host is inoperable, `sendmail(1M)` returns a "Host unknown" message from the alternate host.¹

For `uucp(1C)` mail addresses, the "Host unknown" message probably means that the destination host name is not listed in the `/etc/uucp/Systems` file.

See Also

For information on how `sendmail(1M)` works, see the *System Administration Guide, Volume 3*

550 Security server failed to perform requested command

Cause

While using the 3.x FW-1 FTP Security Server, the user sees the following error message when trying to use FTP `get` or `put` commands:

```
550 Security server failed to perform requested command
```

Action

FW-1's FTP Security Server sends a `pwd` command prior to any data connection command (such as `get`, `put`, `ls`), since it needs to know the current directory for purposes such as logging, virus inspection, and resources. FW-1 assumes that these commands are blocked whenever the `pwd` command is blocked. Therefore, do not disable `pwd` on your FTP server.

¹ This is a known `sendmail(1M)` version 8.6.7 bug.

550 *username...* User unknown

Cause

This `sendmail(1M)` message indicates that the intended recipient, specified by the portion of the address before the at-sign (`@`), could not be located on the destination host machine.

Action

Check the email address and try again, perhaps with a slightly different spelling. If this does not work, contact the intended recipient and ask for a proper address.

See Also

For information on how `sendmail(1M)` works, see the *System Administration Guide, Volume 3*.

554 *hostname...* Local configuration error

Cause

This `sendmail(1M)` message usually indicates that the local host is trying to send mail to itself.

Action

Check the value of the `$j` macro in the `/etc/mail/sendmail.cf` file to ensure that this value is a fully qualified domain name.

Technical Notes

When the sending system provides its host name to the receiving system (in the SMTP `HELO` command), the receiving system compares its name to the sender's name. If these are the same, the receiving system issues this error message and closes the connection. The name provided in the `HELO` command is the value of the `$j` macro.

See Also

For information on how `sendmail(1M)` works, see the *System Administration Guide, Volume 3*.

"A"

A command window has exited because its child exited.

Cause

The argument to a `cmdtool(1)` or a `shelltool(1)` window looks like it is supposed to be a command, but the system cannot find the command.

Action

To run this command inside a `cmdtool(1)` or a `shelltool(1)`, make sure the command is spelled correctly and is in your search path. If necessary, use a full path name. If you intended this argument as an option setting, use a minus sign (-) at the beginning of the option.

Technical Notes

Both the `cmdtool(1)` and the `shelltool(1)` are OpenWindows terminal emulators.

access violation unknown host IP address

Cause

Solstice backup utility fails and displays the following error: `access violation unknown host IP address` on Networker 4.2.2. This error is usually caused by a corrupted host name in the host NIS/NIS+ map/table.

Action

Check the Networker client configuration for an incorrect host name. If all else fails, as a workaround, add the entry to `/etc/hosts`.

Accessing a corrupted shared library

Cause

The system is trying to `exec(2)` an `a.out` that requires that it be linked in a static shared library, and `exec(2)` could not load the static shared library. The static shared library is probably corrupted.

Technical Notes

The symbolic name for this error is `ELIBBAD`, `errno=84`.

Address already in use

Cause

The user attempted to use an address already in use, and the protocol does not allow this.

Technical Notes

The symbolic name for this error is `EADDRINUSE`, `errno=125`.

Address family not supported by protocol family

Cause

An address incompatible with the requested protocol was used.

Technical Notes

The symbolic name for this error is `EAFNOSUPPORT`, `errno=124`.

admintool: Received communication service error 4

Cause

AdminTool could not start a display method, because a remote procedure, which had been called, timed out; therefore, it could not send the request. You receive this error when `admintool(1M)` tries to access the NIS or NIS+ tables and networking is not enabled.

Action

Verify the system network status with `ifconfig -a` to make sure the system is connected to the network. Make sure the Ethernet cable is connected and the system is configured to run NIS or NIS+.

Advertise error

Cause

This error is RFS specific. It occurs when users try to advertise a resource already advertised, try to stop RFS while there are resources still advertised, or try to forcefully unmount a resource that is still advertised.

Technical Notes

The symbolic name for this error is `EADV`, `errno=68`.

answerbook: XView error: NULL pointer passed to xv_set

Cause

The AnswerBook navigator window comes up, but the document viewer window does not. This message appears on the console, and the message `Could not start new viewer` appears in the navigator window. This situation indicates that you have an unknown client or a problem with the network naming service.

Action

Run the `ypmatch(1)` or `nismatch(1)` command to determine if the client host name is in the host's map. If not, add it to the NIS hosts map on the NIS master server. Then, make sure the `/etc/hosts` file on the client contains an IP address and entry for that host name, which is followed by `loghost`.

Note - Reboot, if you changed the `/etc/hosts` file.

Check that the `ypmatch(1)` or `nismatch(1)` *client* hosts command returns the same IP host address as in the `/etc/hosts` file. Finally, quit all existing AnswerBooks and restart.

See Also

For more information on the NIS hosts map, see the section on the default search criteria in the *NIS+ and FNS Administration Guide*. If you are using AnswerBook online documentation, "NIS hosts map" is a good search string.

apdb: Resource temporarily unavailable

Cause

This error can occur when attempting to add or remove AP databases with the `apdb` command.

Action

From `/var/adm/messages` you find the reason for the `apdb` command failure, as shown below:

```
Jan 15 14:00:51 Starfire2 apd[683]: /etc/system: could not find:
* End AP database info (do not edit)
Jan 15 14:00:52 Starfire2 apd[683]: failed to patch the system file!
```

Unfortunately, this error from the `netcon` session does not get an `echo` to the console; therefore, it can easily be missed. To correct it, simply edit the `/etc/system` file so that it has the correct comments before and after setting `ap:apdb_dblist`. See below:

```
* Begin AP database info (do not edit)
set ap:apdb_dblist="sd:5 sd:8"
* End AP database info (do not edit)
```

Arg list too long

Cause

The system could not handle the number of arguments given to a command or program when it combined those arguments with the environment's exported shell variables. The argument list limit is the size of the argument list plus the size of the environment's exported shell variables.

Action

The easiest solution is to reduce the size of the parent process environment by unsetting extraneous environment variables. (See the man page for the shell you are using to find out how to list and change your environment variables.) Then run the program again.

Technical Notes

An argument list longer than `ARG_MAX` bytes was presented to a member of the `exec(2)` family of system calls.

The symbolic name for this error is `E2BIG`, `errno=7`.

Argument out of domain

Cause

This message is a programming error or a data input error.

Action

Ask the program's author to fix this condition or to supply data in a different format.

Technical Notes

This indicates an attempt to evaluate a mathematical programming function at a point where its value is not defined. The argument of a programming function in the math package is out of the domain of the function. This could happen when taking the square root, power, or log of a negative number, when computing a power to a non-integer, or when passing an out-of-range argument to a hyperbolic programming function.

To help pinpoint a program's math errors, use the `matherr(3M)` facility.

The symbolic name for this error is `EDOM`, `errno=33`.

Arguments too long

Cause

This C shell error message indicates that too many arguments follow a command. For example, this can happen by invoking `rm *` in a huge directory. The C shell cannot handle more than 1706 arguments.

Action

Temporarily start a Bourne shell with `sh(1)` and run the command again. The Bourne shell dynamically allocates command line arguments. Return to your original shell by typing `exit`.

assertion failed: *string*, file *name*, line *int*

Cause

An unexpected condition in the program has occurred.

Action

Contact the vendor or author of the program to ask why it failed. If you have the source code for the program, you can look at the file and line number where the assertion failed. This might give you an idea of how to run the program differently.

Technical Notes

This message is the result of a diagnostic macro called `assert(3C)` that a programmer inserted into the specified line of a source file. The untrue expression precedes the file name and line number.

Attempting to link in more shared libraries than system limit

Cause

The system is trying to `exec(2)` an `a.out` that requires more static shared libraries than is allowed on the current configuration of the system.

Technical Notes

The symbolic name for this error is `ELIBMAX`, `errno=86`.

automount[*int*]: *name*: Not a directory

Cause

The file specified after the first colon is not a valid mount point, because it is not a directory.

Action

Ensure that the mount point is a directory and not a regular file or a symbolic link.

automountd[*int*]: server *hostname* responding

Cause

This automounter message indicates that the system tried to mount a file system from an NFS™ server that is either down or extremely slow to respond. In some cases, this message indicates that the network link to the NFS server is broken, although that condition produces other error messages as well.

Action

If you are the system administrator responsible for the non-responding NFS server, check to see whether the machine needs repair or rebooting. Encourage your user community to report such problems quickly, but only once. When the NFS server is back in operation, the automounter can access the requested file system.

See Also

For more information on NFS failures, see the section on NFS troubleshooting in the *System Administration Guide, Volume 3*. If you are using AnswerBook online documentation, a good search string is "NFS Service."

"B"

Bad address

Cause

The system encountered a hardware fault in attempting to access a parameter of a programming function.

Action

Check the address to see if it resulted from supplying the wrong device or option to a command. If that is not the problem, contact the vendor or author of the program for an update.

Technical Notes

This error could occur any time a function that takes a pointer argument is passed an invalid address. Because processors differ in their ability to detect bad addresses, on some architectures, passing bad addresses can result in undefined behaviors.

The symbolic name for this error is `EFAULT`, `errno=14`.

BAD/DUP FILE I=i OWNER=o MODE=m SIZE=s MTIME=t CLEAR?

Cause

While checking inode link counts during phase 4, `fsck(1M)` found a file (or directory) that either does not exist or exists somewhere else.

Action

To clear the inode of its reference to this file or directory, answer "yes." With the `-p` (green) option, `fsck(1M)` automatically clears bad or duplicate file references. Answering "yes" to this question seldom causes a problem.

Bad file number

Cause

Generally this message is a program error, not a usage error.

Action

Contact the vendor or author of the program for an update.

Technical Notes

Either a file descriptor refers to no open file, or a `read(2)`—or a `write(2)`—request is made to a file that is open only for writing or reading.

The symbolic name for this error is `EBADF`, `errno=9`.

block no. BAD I=inode no.

Cause

Upon detecting an out-of-range block, `fsck(1M)` prints the bad block number and its containing inode (after `I=`).

Action

In `fsck(1M)` phases 2 and 4, you decide whether or not to clear these bad blocks. Before committing to repair with `fsck(1M)`, you could determine which file contains this inode by passing the inode number to the `ncheck(1M)` command:

```
# ncheck -i inum filesystem
```

See Also

For more information, see the chapter on checking file system integrity in the *System Administration Guide, Volume 1*.

BAD_MESSAGE (error code 100) from X.400

Cause

In this situation, X.400 software had been working without problems. Suddenly, the message exchanges failed in `ma_start_delivery()`. It was returning an error code of 100 (`BAD_MESSAGE`).

The `ma_start_delivery()` call fails when trying to exchange a file of more than 900 bytes.

Action

X.400 was restarted with the wrong umask. To fix, set the umask to 0022 and restart the software.

bad module/chip at: *position*

Cause

This message from the memory management system often appears with parity errors and indicates a bad memory module or chip at the position listed. Data loss is possible, if the problem occurs other than at boot time.

Action

Replace the memory module or chip at the indicated position. Refer to the vendor's hardware manual for help finding this location.

Bad request descriptor

Cause

This message is apparently only used in NIS+ to indicate corrupted or missing tables.

Technical Notes

The symbolic name for this error is EBADR, `errno=51`.

BAD SUPER BLOCK: *string*

Cause

This message from `fsck(1M)` indicates that a file system's super block is damaged beyond repair and must be replaced. At boot time (with the `-p` option) this message is prefaced by the file system's device name. After this message comes the actual damage recognized (see Action). Unfortunately, `fsck(1M)` does not print the number of the damaged super block.

Action

The most common cause of this error is overlapping disk partitions. Do not immediately rerun `fsck(1M)` as suggested by the lines that display after the error message. First, make sure that you have a recent backup of the file system involved; if not, try to back up the file system now using `ufsdump(1M)`. Then, run the `format(1M)` command, select the disk involved, and print out the partition information.

```
# format
: N
> partition
> print
```

Note whether the overlap occurs at the beginning or end of the file system involved. Then, run `newfs(1M)` with the `-N` option to print out the file system parameters, including the location of backup super blocks.

```
# newfs -N /dev/dsk/device
```

Select a super block from a non-overlapping area of the disk, but note that in most cases you have only one chance to select the proper replacement super block, which `fsck(1M)` soon propagates to all the cylinders. If you select the wrong replacement super block, data corruption will probably occur, and you will have to restore from backup tapes. After you select a new super block, provide `fsck(1M)` with the new master super block number:

```
# fsck -o b=NNNN /dev/dsk/device
```

Technical Notes

Specific reasons for a damaged super block include: a wrong magic number, an out-of-range number of cylinder groups (NCG) or cylinders per group (CPG), the wrong number of cylinders, a preposterously large super block size, and trashed values in super block. These reasons are generally not meaningful, because a corrupt super block is usually extremely corrupt.

See Also

For more information on bad super blocks, see the sections on restoring bad super blocks in the *System Administration Guide, Volume 1*. If you are using AnswerBook online documentation, "super block" is a good search string.

BAD TRAP

Cause

A bad trap can indicate faulty hardware or a mismatch between hardware and its configuration information. Data loss is possible if the problem occurs other than at boot time.

Action

If you recently installed new hardware, verify that the software was correctly configured. Check the kernel traceback displayed on the console to see which device generated the trap. If the configuration files are correct, you probably have to replace the device.

In some cases, the bad trap message indicates a bad or down-rev CPU.

Technical Notes

A hardware processor trap occurred, and the kernel trap handler was unable to restore the system state. This message is a fatal error that usually precedes a panic, after which the system performs a sync, dump, and reboot. The following conditions can cause a bad trap: a system text or data access fault, a system data alignment error, or certain kinds of user software traps.

/bin/sh: *file*: too big

Cause

This Bourne shell message indicates a classic "no memory" error. While trying to load the program specified after the first colon, the shell noticed that the system ran out of virtual memory (swap space).

Action

For information on reconfiguring your system to add more swap space, refer to "Not enough space" on page 183.

Block device required

Cause

A raw (character special) device was specified where a block device was required, such as during a call to the `mount(1M)` command.

Action

To see which block devices are available, use `ls -l` to look in `/devices`. Then specify a block device instead of a character device. Block device modes start with a `b`, whereas raw character device modes start with a `c`.

Technical Notes

The symbolic name of this error is `ENOTBLK, errno=15`.

Boot device: /iommu/sbus/*directory*/*directory*/ sd@3,0

Cause

This message always appears at the beginning of rebooting. If there is a problem, the system hangs, and no other messages appear. This condition is caused by conflicting SCSI targets for the boot device, which is almost always target 3.

Action

The boot device is usually the machine's internal disk drive, target 3. Make sure that external and secondary disk drives are targeted to 1, 2, or 0, and do not conflict with each other. Also make sure that the tape drives are targeted to 4 or 5, and CD drives to 6, avoiding any conflict with each other or with the disk drives. You can set a device's target number using push-button switches or a dial on the back near the SCSI cables. If the targeting of the internal disk drive is in question, check it by powering off the machine, removing all external drives, turning the power on, and running the `probe-scsi-all` or `probe-scsi` command from the PROM monitor.

Broadcast Message from root (pts/*int*) on server [date]

Cause

This message from the `wall(1M)` command is transmitted to all users logged into a system. You could see it during a `rlogin(1)` or `telnet(1)` session, or on terminals connected to a timesharing system.

Action

Carefully read the broadcast message. Often this broadcast is followed by a shutdown warning.

For details about system shutdown, refer to “The system will be shut down in *int* minutes” on page 237.

See Also

For more information on bringing down the system, see the section on halting the system in the *System Administration Guide, Volume 1*. If you are using AnswerBook online documentation, “halting the system” is a good search string.

Broken pipe

Cause

This condition is often normal, and the message is merely informational (as when piping many lines to the `head(1)` program). The condition occurs when a write on a pipe does not find a reading process. This usually generates a signal to the executing program, but this message displays when the program ignores the signal.

Action

Check the process at the end of the pipe to see why it exited.

Technical Notes

The symbolic name of this error is `EPIPE`, `errno=32`.

Bus Error

Cause

A process has received a signal indicating that it attempted to perform I/O to a device that is restricted or that does not exist. This message is usually accompanied by a core dump, except on read-only file systems.

Action

Use a debugger to examine the core file and determine what program fault or system problem led to the bus error. If possible, check the program's output files for data corruption that might have occurred before the bus error.

Technical Notes

Bus errors can result from either a programming error or device corruption on your system. Some common causes of bus errors are: invalid file descriptors, unreasonable I/O requests, bad memory allocation, misaligned data structures, compiler bugs, and corrupt boot blocks.

"C"

Cannot access a needed shared library

Cause

The system is trying to `exec(2)` an `a.out` that requires a static shared library, and the static shared library does not exist or the user does not have permission to use it.

Technical Notes

The symbolic name for this error is `ELIBACC`, `errno=83`.

Cannot allocate colormap entry for "*string*"

Cause

This message from `libXt` (X Intrinsic library) indicates that the system color map was full, even before the color name specified in quotes was requested. Some applications can continue after this message. Other applications, such as workspace properties color, fail to come up when the color map is full.

Action

Exit the programs that make heavy use of the color map, then restart the failed application and try again.

Cannot assign requested address

Cause

An attempt was made to create a transport endpoint with an address not on the current machine.

Technical Notes

The symbolic name for this error is `EADDRNOTAVAIL`, `errno=126`.

Cannot bind to domain *domainname*: can't communicate with ypbind

Cause

While running the `ypinit -m` script for the setup of an NIS Master Server, you get this error message.

Action

You could be using the wrong `nsswitch` template for `/etc/nsswitch.conf`. During setup, you should be using `/etc/nsswitch.files` as the name services switch

template. After setup is complete, you would then want to use `/etc/nsswitch.nis`. Do the following to verify that you are using `nsswitch.files`:

```
# head /etc/nsswitch.conf
# -->
# /etc/nsswitch.files:
```

If you are not using the `nsswitch.files`, copy it over as shown below:

```
# cp /etc/nsswitch.files /etc/nsswitch.conf
```

Run the `ypinit -m` script, again.

Cannot boot after install, error that points to an `.rc` file

Cause

The user completes the installation of the Solaris 2.6 IA software. Upon reboot, the user gets an error referencing an `.rc` file (example: `11045.rc`). This file has probably been deleted or placed in a different directory. As the Solaris software looks for this file during the bootup sequence and cannot find it, the system hangs, because it cannot complete the boot process.

Action

During the installation process, there is an option to save the configuration assistant choices to a file. The error is pointing to the saved configuration file. The user was never supposed to have the option to save these choices to a file. Users should exit the setup after making their choices. If the users do save these choices to a file and if this file gets deleted or moved, the system hangs during the boot process. To solve this problem, the user boots in single user mode. From the `#` prompt, the user should do the following:

1. `cd /platform/i86pc/boot/solaris/machines`
2. Delete all files in this directory.
3. Reboot the system.

This corrects the problem and allows the Solaris software to complete loading.

cannot change passwd, not correct passwd

Cause

While running `yppasswd(1)` and trying to change a user's password, the system responded with this message: `cannot change passwd, not correct passwd`.

Also, the user was getting `yppasswd user string` does not exist on the server console, but by running `ypcat passwd | grep user` it returns the user name. It was verified that `yppasswdd(1M)` was running.

Action

Check the `passwd(4)` file with `pwck(1M)` and verify that `yppasswdd(1M)` is running on the right server. Then verify where the `passwd(4)` file is located and, if changed, check that `yppasswdd(1M)` has the location in the process line. The password located in `/etc/yp` should read

```
/usr/lib/yp/rpc.yppasswdd -D /etc/yp. The -D option with the passwd files directory location tells yppasswdd(1M) where to update and verify password changes.
```

cannot establish nfs service over /dev/tcp: transport setup problem

Cause

During boot strap of a SunOS 2.5.2 system, `nfsd(1M)` displays the following:

```
netdir_getbyname (transport tcp, host/serv \1/  
nfs), No such file or directory  
Cannot establish NFS service over /dev/tcp: transport setup problem.
```

The problem: The NIS maps have been populated from older systems, and the `nfs/tcp` entry of the services map is missing. (The user is running NIS+, but this problem can also occur with NIS.)

Action

Either put a `files` entry before the `nis` or `nisplus` in the services line of the `/etc/nsswitch.conf` file, or, better, merge the changes to the services file into the services map.

It is a good idea to always merge in the new entries to `/etc/services`, `/etc/inet/protocols`, and `/etc/rpc` into their respective maps whenever a new OS is installed.

Cannot exec a shared library directly

Cause

The system is attempting to `exec(2)` a shared library, directly.

Technical Notes

The symbolic name for this error is `ELIBEXEC`, `errno=87`.

Cannot find SERVER *hostname* in network database

Cause

A brief description: the user is on a different subnet and is running permanent licenses:

```
ultral(50)% cc -o hello hello.c
License Error : Cannot find the license server (fry)
in the network database for product(Sun WorkShop Compiler C)
Cannot find SERVER hostname in network database (-14,7)
cc: acomp failed for hello.c
ultral(51)%
```

Action

Check the following:

1. Make sure that the server is up and running.
2. Make sure that the server is in the `/etc/hosts` file of the client system by typing: `ping servername`.
3. Make sure the license daemon on the server is running.
4. Make sure there is an elementary license file on the client:

```
cd /etc/opt/licenses
more sunpro.loc
```

5. Make sure there are only text license files, such as `sunpro.lic.1` in the `sunpro,loc` directory.
6. For the client check, see below:

```
% cd /etc
% more nsswitch.conf | grep hosts
hosts:      nis [NOTFOUND=return] files
```

This means that it is using the NIS server to look up the IP address. If it is set first for `nis` and the `/etc/hosts` file has the server listed by name, change the line to

```
hosts:      files nis
```

Then, see if it can be found. If not, try `truss` and `snoop` to see what is happening.

cannot install bootblock

Cause

In this case, the user installs the Solaris IA software on the Intel platform and the install seems fine. When the system is rebooted after the installation, the user receives the above error message at startup. At this point, the user cannot gain access to the system.

Action

This error occurs when you use the `fdisk` utility in the Solaris operating environment, do a `newfs`, and then do a `restore`, but forget to do the `install` for the boot block. When you do a `newfs` and then a `restore` operation, you need to perform an `installboot` before installing the OS. Otherwise, you get the above error. There is no guarantee, but the `installboot` procedure might or might not work after booting into single user mode from the CD-ROM.

To install the UFS boot block and partition the boot program on slice 2 of target 0 on controller 1 of the platform, where the command is being run, use the following:

```
# installboot /usr/platform/uname -i/lib/fs/ufs/pboot \  
/usr/platform/uname -i/lib/fs/ufs/bootblk /dev/rdisk/c1t0d0s2
```

Cannot open FCC file

Cause

When trying to send mail by Netscape, this message is displayed. Netscape is trying to save the outbound message to a file that has been specified by the user, but does not exist.

Action

To correct this problem do the following: go to options Mail and News Preferences, then go to Compose. A template pops up. There is a section that specifies where to save outgoing mail and news files. Make sure that these files exist or remove them from the template, if you do not care about logging which messages are sent through Netscape.

Cannot send after transport endpoint shutdown

Cause

A request to send data was disallowed, because the transport endpoint has already been shut down.

Technical Notes

The symbolic name for this error is ESHUTDOWN, errno=143.

can't communicate with ypbind

Cause

ypcat passwd returns with the error message, can't communicate with ypbind, but ypbind is running.

```
ls -l /var/yp/binding/ypbind.pid
-r----- 1 root    root          3 Dec  1 07:40 ypbind.pid
```

umask for root is set to 077.

Action

Set umask for root back to 022. /var/yp/binding/ypbind.pid must be readable by all groups.

Refer to the following example:

```
ls -l /var/yp/binding/ypbind.pid
-r--r--r-- 1 root    root          3 Dec  1 07:40 ypbind.pid
```

Can't create public message device (Device busy)

Cause

This message comes from the lp(1) print scheduler, indicating that it is either extremely busy or hanging.

Action

If print jobs are coming out of the printer in question, wait until they are finished and then resubmit this print job. If you see this message again, the lp(1) system is probably hanging.

For a procedure to clear the queue, refer to "lp hang" on page 153.

Technical Notes

If `lp(1)` is unable to create a device for printer messages, the message FIFO could already be in use or could be locked by another print job.

See Also

For more information on the print scheduler, see the section on administrating printers in the *System Administration Guide, Volume 2*.

Can't invoke `/etc/init`, error `int`

Cause

This message can appear while a system is booting, indicating that the `init(1M)` program is missing or corrupted. Note that `/etc/init` is a symbolic link to `/sbin/init`.

Action

Do the following:

1. Boot the mini-root so you can replace `init(1M)`.
2. Halt the machine by typing `Stop-A` or by pressing the reset button.
3. Reboot as a single user from the CD-ROM, the net, or a diskette. For example, type `boot cdrom -s` at the `ok` prompt to boot from a CD-ROM.
4. After the system comes up and gives you a `#` prompt, mount the device corresponding to the original `root (/)` partition somewhere, with a command similar to the `mount(1M)` command, as shown below:

```
# mount /dev/dsk/c0t3d0s0 /mnt
# cp /sbin/init /mnt/sbin/init
# reboot
```

5. Then copy the `init(1M)` program from the mini-root to the original `root (/)` partition.
6. Reboot the system.

If this does not work, other files might be corrupted, and you might need to reinstall the entire system.

Technical Notes

The error number is 2 if `/sbin/init` is missing, or 8 if `/sbin/init` has an incorrect executable format. This message is usually followed by a `panic: icode` message. The system tries to reboot itself, but goes into a loop, because rebooting is impossible without `init(1M)`.

See Also

For more information on booting the system, see the section on halting and booting the system in the *System Administration Guide, Volume 1*.

can't open `/dev/rdisk/string`: (null): UNEXPECTED INCONSISTENCY

Cause

In the SunOS™ 4.1.x release, this message indicated that the device containing the `/dev` file system has become disconnected.

A particular response from the Solaris operating environment has not been defined.

can't synchronize with hayes

Cause

This message sometimes appears when using a modem that the system regards as a "Hayes" type modem, which includes most modems manufactured today. The message can be caused by incorrect switch settings, by poor cable connections, or by not turning the modem on.

Action

Check that the modem is on and that the cables between the modem and your system are securely connected. Check the internal and external modem switch settings. If necessary, turn the modem off and then on again.

cd: Too many arguments

Cause

The C shell's `cd(1)` command takes only one argument. Either more than one directory was specified, or a directory name containing a space was specified. Directory names with spaces are easy to create with File Manager.

Action

Use only one directory name. To change to a directory whose name contains spaces, enclose the directory name in double (") or single (') quotes, or use File Manager.

Channel number out of range

Cause

The system has run out of stream devices. This error results when a stream head attempts to open a minor device that does not exist or is currently in use.

Action

Check that the stream device in question exists and was created with an appropriate number of minor devices. Make sure that the hardware corresponds to this configuration. If the stream device configuration is correct, try again later when more system resources might be available.

Technical Notes

The symbolic name for this error is `ECHRNG`, `errno=37`.

chmod: ERROR: invalid mode

Cause

This message from the `chmod(1)` command indicates a problem in the first non-option argument.

Action

If you are specifying a numeric file mode, you can provide any number of digits (although only the final one-to-four are considered), but all digits must be between 0 and 7. If you are specifying a symbolic file mode, use the syntax provided in the `chmod(1)` usage message to avoid the "invalid mode" error message: Usage: `chmod [ugoa][+ -=][rwxlstugo] file ...`

Some combinations of symbolic key letters produce no error message, but fail to have any effect. The first group, `[ugoa]`, is truly optional. The second group, `[+ -=]`, is mandatory for `chmod(1)` to have an effect. The third group, `[rwxlstugo]`, is also mandatory for effect and can be used in combination when that combination does not conflict.

Command not found

Cause

The C shell could not find the program you gave as a command.

Action

Check the form and spelling of the command line. If that looks correct, use `echo $path` to see if the user's search path is correct. When communications are garbled, it is possible to unset a search path to such an extent that only built-in shell commands are available. Below is a command to reset a basic search path:

```
% set path = (/usr/bin /usr/ccs/bin /usr/openwin/bin .)
```

If the search path looks correct, check the directory contents along the search path to see if programs are missing or if directories are not mounted.

See Also

For more information about the C shell, see `csh(1)`.

Communication error on send

Cause

This error occurs when the current process is waiting for a message from a remote machine, but the link connecting the machines breaks.

Technical Notes

The symbolic name for this error is `ECOMM`, `errno=70`.

config error: mail loops back to myself.

Cause

User sees this message when sending mail:

```
# dle@g3... Connecting to g3.xyz.edu. (ether)...
220 xyz.edu Sendmail SMI-8.6/
SMI-SVR4 ready at Wed, 7 Jan 1998 14:28:20 -0600
>>> HELO xyz.edu
250 xyz.edu Hello g1.xyz.edu [129.106.16.1], pleased to meet you
xyz.edu config error: mail loops back to myself
>>> QUIT
221 g1.xyz.edu closing connection
dle@g3... Local configuration error
Saving message in /dead.letter
/dead.letter... Sent
```

The sending system (see line 220) and the receiving system (see the HELO line) both think they are known as "xyz.edu."

Action

Edit the `sendmail.cf` file as follows:

1. Type the official host name.
2. For the domain, you have choices: If you want the gateway machine to identify itself as the domain, use `Dj$m`; if you want the gateway machine to appear to be inside the domain, use `Dj$w.$m`; and if you are using `sendmail.mx` (or have a fully-qualified host name), use `Dj$w`.

3. Uncomment `Dj$w.$m` and comment `Dj$m`. This gives each system a unique name. `$w` is the system host name, and `$m` is the domain.

connect from *hostIP* to `callit(ypserv)`: request from non-local host

Refer to “connect from *hostIP* to `callit(ypserv)`: request from unauthorized host” on page 71.

connect from *hostIP* to `callit(ypserv)`: request from unauthorized host

Cause

An example of a message from SunOS:

```
Jan  5 14:45:37 host1 portmap[86]: connect from 158.175.36.135 to
callit(ypserv): request from unauthorized host
```

Other possibilities for the end portion of the error message include:

- request from unprivileged port
- request from non-local host
- request not forwarded

In the Solaris operating environment, the error might look similar to the following:

```
Jan  5 14:45:37 host1 rpcbind[86]: refused connect from 158.175.36.135
to callit(ypserv)
```

In all cases, the `ypserv` part of the message might actually be any RPC service, such as `mount` or `nfs` or `status`.

Action

The user has a replacement `portmap` or `rpcbind`. The version is enhanced to add access controls, and the error in question is reporting an access violation. The

replacements are third-party and are not supported by Sun. The user must locate the access control configuration files and change them to the desired access controls.

connect from *hostIP* to callit(ypserv): request from unprivileged port

Refer to “connect from *hostIP* to callit(ypserv): request from unauthorized host” on page 71.

connect from *hostIP* to callit(ypserv): request not forwarded

Refer to “connect from *hostIP* to callit(ypserv): request from unauthorized host” on page 71.

Connection closed.

Cause

When using `rlogin(1)`, this message can appear under the following circumstances:

- If the remote host cannot create a process for this user
- If the user takes too long to type the correct password
- If the user interrupts the network connection
- If the remote host goes down

Data loss is possible if files were modified and not saved before the connection closed.

Action

Try again. If the other system has gone down, wait for it to reboot first.

Connection closed by foreign host.

Cause

When a user applies `telnet(1)` to another system, this message can appear under the following circumstances:

- If the user takes too long to type the correct password
 - If the remote host cannot create a login for this user
 - If the remote host goes down or terminates the connection
- Data loss is possible if files were modified and not saved before the connection closed.

Action

Try again. If the other system has gone down, wait for it to reboot first.

[Connection closed. Exiting]

Cause

After using the `talk(1)` command to communicate with another user, the other person enters an interrupt (usually Control-C), and this message appears on your screen.

Action

Sending an interrupt is the usual way of exiting the `talk` program. The `talk(1)` session is over, and you can return to your work.

Connection refused

Cause

No connection could be made because the target machine actively refused it. This happens either when trying to connect to an inactive service or when a service process is not present at the requested address.

Action

Activate the service on the target machine, or start it up again if it has disappeared. If, for security reasons, you do not intend to provide this service, inform the user community, possibly suggesting an alternative.

Technical Notes

The symbolic name for this error is `ECONNREFUSED`, `errno=146`.

Connection reset by peer

Cause

A connection was forcibly closed by a peer. This is normally due to a remote host connection loss from a timeout or a reboot.

Technical Notes

The symbolic name for this error is `ECONNRESET`, `errno=131`.

Connection timed out

Cause

This error occurs either when the destination host is down or when problems in the network cause a loss in transmission.

Action

Do the following:

1. Check the operation of the host system, for example by using `ping(1M)` and `ftp(1)`.
2. Repair or reboot as necessary.
3. If the above does not solve the problem, check the network cabling and connections.

Technical Notes

No connection was established in a specified time. A connect or send request failed because the destination host did not properly respond after a reasonable interval. (The time-out period is dependent on the communication protocol.)

The symbolic name for this error is ETIMEDOUT, errno=145.

console login: ^J^M^Q^K^K^P

Cause

This error usually occurs because OpenWindows exited abnormally, leaving the system's keyboard in the wrong mode. The characters that appear when someone attempts to login are garbage transliterations of what someone typed.

Action

If you are on a SPARC™ system, do the following:

1. Find another machine and remote log in to this system
2. Run the following command:

```
$ /usr/openwin/bin/kbd_mode -a
```

This puts the console back into ASCII mode.

Note - `kbd_mode` is not a windows program; it fixes the console mode.

If you are on an IA system, do the following:

1. Log in remotely and `start`
2. `kill` the X server or reboot the system

Technical Notes

The usual reason for this problem occurring is an automated script run from `cron(1M)` that clears the `/tmp` directory periodically. Ensure that any such scripts do not remove the `/tmp/.X11-pipe` or `/tmp/.X11-unix` directories, or any files in them.

core dumped

Cause

A `core(4)` file contains an image of memory at the time of software failure and is used by programmers to find the reason for the failure.

Action

To see which program produced a `core(4)` file, run either the `file(1)` command or the `adb(1)` command. The following examples show the output of the `file(1)` and `adb(1)` commands on a core file from the `dtmail` program.

```
$ file core
core: ELF 32-bit MSB core file SPARC Version 1, from 'dtmail'
```

```
$ adb core
core file = core -- program 'dtmail'
SIGSEGV 11: segmentation violation
^D      (use Control-d to quit the program)
```

Ask the vendor or author of this program for a debugged version.

Technical Notes

Some signals, such as `SIGQUIT`, `SIGBUS`, and `SIGSEGV`, produce a core dump. See the `signal(5)` man page for a complete list.

If you have the source code for the program, you can try compiling it with `cc -g`, and debugging it yourself using `dbx` or a similar debugger. The `where` directive of `dbx` provides a stack trace.

On mixed networks, it can be difficult to discern which machine architecture produced a particular core dump, since `adb(1)` on one type of system generally cannot read a `core(4)` file from another type of system and can produce an unrecognized file message. Run `adb(1)` on various machine architectures until you find the right one.

See Also

For information on saving and viewing crash information, see the *System Administration Guide, Volume 2*. If you are using AnswerBook online documentation, "system crash" is a good search string.

corrupt label - wrong magic number or corrupt label or corrupt label - label checksum failed

Cause

After a power cycle, the machine displays either of the following error messages:

- corrupt label - label checksum failed
- corrupt label - wrong magic number

`format(1M)` displayed the following:

0	unassigned	wm	0	0	(0/0/0)	0
1	unassigned	wm	0	0	(0/0/0)	0
2	backup	wm	0 - 5460	4.2G	(5460/0/0)	4154160
3	unassigned	wm	0	0	(0/0/0)	0
4	unassigned	wm	0	0	(0/0/0)	0
5	unassigned	wm	0	0	(0/0/0)	0
6	unassigned	wm	0 - 2730	2.1G	(0/0/0)	0
7	unassigned	wm	2730-5460	2.1G	(0/0/0)	0

The disks were using raw partitions beginning at block 0 (cylinder 0). The disk label (VTOC) is kept on the block 0 of cylinder 0. The label eventually gets overwritten by database programs using raw partitions, if the raw partition begins at cylinder 0. (UNIX® file systems avoid this area of the partition.)

Action

As a workaround, do the following:

1. Go into `format(1M)` and get the backup label using the `backup` command.
2. Relabel the disk using this backup label. You should then be able to access the disk.
3. Backup the data on this disk.
4. Go back to the disk and relabel it, starting the raw partition at cylinder 1. (This loses one cylinder, but prevents corrupting the VTOC.)
5. Label again.

6. Restore the data from your backup.

could not grant slave pty

Cause

User gets the error message `could not grant slave pty` when attempting a `telnet(1)`, `rlogin(1)`, or `rsh(1)` session (anything that requires a shell) or when trying to bring up an x-term.

Action

The user's file permissions were set wrong on `/usr/lib/pt_chmod`. The user had:

```
# ls -la /usr/lib/pt_chmod
---s--x--x  1 bin      bin           3120 May  3  1996
```

The permissions should be:

```
# ls -la /usr/lib/pt_chmod
---s--x--x  1 root      bin           3120 May  3  1996
```

Note - The owner should be `root`; the user had `bin` as the owner. Also, the `setuid` bit must be set.

By using `chown root pt_chmod`, the problem was corrected.

Could not initialize tooltalk (tt_open): TT_ERR_NOMP

Cause

Various desktop tools display or print this message when the `ttsession(1)` process is not available. The ToolTalk service generally tries to restart `ttsession(1)`, if it is not

running. Thus, this error indicates that the ToolTalk service is either not installed or is not installed correctly.

Action

Verify that the *ttsession(1)* command exists in `/usr/openwin/bin` or `/usr/dt/bin`. If this command is not present, ToolTalk is not installed correctly. The packages constituting ToolTalk are the runtime `SUNWt1tk`, developer support `SUNWt1tkd`, and the manual pages `SUNWt1tkm`. CDE ToolTalk packages have the same names with ".2" appended.

Technical Notes

The full `TT_ERR_NOMP` message string reads as follows: "No *ttsession(1)* is running, probably because *tt_open(3)* has not been called yet. If this is returned from *tt_open(3)*, it means *ttsession(1)* could not be started, which generally means ToolTalk is not installed on the system."

Could not open ToolTalk Channel

Cause

This error message is displayed while attempting to remotely run `workshop`.

Action

Do the following:

1. Make sure `workshop` is no longer running.
2. In the `telnet/rlogin` session window, type:
`/bin/ps -ef | grep tsession`. If one is running in the system that belongs to the `telnet` user, type `kill pid_of_tsession`.
3. In the `telnet rlogin` session, type
`/usr/dt/bin/ttsession -s -d machine_telnetting_from:0.0.`
4. Start `workshop`.

Could not start new viewer

Cause

This message appears in the AnswerBook navigator window, along with an XView error message on the console.

Action

For details, refer to “answerbook: XView error: NULL pointer passed to xv_set” on page 46.

Could not start NFS service for any protocol. Exiting

Cause

The following errors occur at boot time:

```
/usr/lib/nfs/nfsd[478]: t_bind to wrong address
/usr/lib/nfs/nfsd[478]: t_bind to wrong address
/usr/lib/nfs/nfsd[478]: Cannot establish NFS service over /dev/
udp: transport setup problem.
/usr/lib/nfs/nfsd[478]: Cannot establish NFS service over /dev/
udp: transport setup problem.
/usr/lib/nfs/nfsd[478]: t_bind to wrong address
/usr/lib/nfs/nfsd[478]: t_bind to wrong address
/usr/lib/nfs/nfsd[478]: Cannot establish NFS service over /dev/
tcp: transport setup problem.
/usr/lib/nfs/nfsd[478]: Cannot establish NFS service over /dev/
tcp: transport setup problem.
/usr/lib/nfs/
nfsd[478]: Could not start NFS service for any protocol. Exiting.
/usr/lib/nfs/
nfsd[478]: Could not start NFS service for any protocol. Exiting.
```

In this situation, a backup copy of the `S15nfs.server` script in `/etc/rc3.d` was made. However, the backup copy was renamed to `S15nfs.server.BAK`. Since the backup copy starts with a upper case “S,” it was also executed at boot time. The errors occurred when a second NFSD was attempted.

Action

If a backup copy of any startup script is made, it should be renamed with a lower case "s," so as not to be executed at boot.

cpio: Bad magic number/header.

Cause

A `cpio(1)` archive has either become corrupted or was written out with an incompatible version of `cpio(1)`.

Action

Use the `-k` option to `cpio(1)` to skip I/O errors and corrupted file headers. This might permit you to extract other files from the `cpio(1)` archive. To extract files with corrupted headers, try editing the archive with a binary editor such as *emacs(1)*. Each `cpio(1)` file header contains a filename as a string.

See Also

For more information on magic numbers, see `magic(4)`.

cpio : can't read input : end of file encountered prior to expected end of archive.

Cause

This message appears when trying to read a multi-volume floppy in bar format using the following command:

```
# cpio -id -H bar -I /dev/diskette0
```

Action

kill `/usr/sbin/vold` by running `/etc/init.d/volmgt stop` and use the device name `/dev/rfd0`.

Cross-device link

Cause

An attempt was made to make a hard link to a file on another device, such as on another file system.

Action

Establish a symbolic link using `ln -s` instead. Symbolic links are permitted across file system boundaries.

Technical Notes

The symbolic name for this error is `EXDEV`, `errno=18`.

"D"

data access exception

Cause

This message appears when running an old version of the operating system that does not support new hardware or when running an operating system that is not configured for new hardware. It can also be the result of an incorrectly installed DSIMMs or a disk problem.

Action

Upgrade your operating system to a version that supports the new hardware or machine architecture.

See Also

For more information on upgrades, see the section describing system and device configuration in the *Solaris Transition Guide*.

Data fault

Cause

This error is a kind of `BAD TRAP` that usually causes a system panic. When this message appears after a `BAD TRAP` message, a system text or data access fault probably occurred.² In the absence of a `BAD TRAP` message, this message might indicate a user text or data access fault. Data loss is possible, if the problem occurs other than at boot time.

Action

Make sure the machine can reboot, then check the log file `/var/adm/messages` for hints about what went wrong.

Deadlock situation detected/avoided

Cause

A programming deadlock situation was detected and avoided.

Action

If the system had not detected and avoided a deadlock, a piece of software would have hung. Run the program again. The deadlock might not reoccur.

Technical Notes

This error usually relates to file and record locking, but can also apply to mutexes, semaphores, condition variables, and read/write locks.

The symbolic name for this error is `EDEADLK`, `errno=45`.

See Also

See the section on deadlock handling in the *System Interface Guide*. See also the section on avoiding deadlock in the *Multithreaded Programming Guide*.

². See the message `BAD TRAP` for more information.

Destination address required

Cause

A required address was omitted from an operation on a transport endpoint.
Destination address required.

Technical Notes

The symbolic name for this error is `EDESTADDRREQ`, `errno=96`.

destination component full

Cause

Solstice backup is reporting `destination component full`.

This message appears when a manual operation is performed on the jukebox/autochanger (for example, physically unloading the tape drive by means of the buttons on the autochanger, rather than using SBU to unmount the volume). This operation causes SBU to lose track of the status of the media in the autochanger.

Action

The following command should resolve the problem: `/usr/sbin/nsr/nsrjb -H`.

`/dev/fd/int: /dev/fd/int: cannot open`

Cause

`setuid` and `setgid` shell scripts refuse to run. They return an error message similar to `/dev/fd/3: /dev/fd/3: cannot open`. (The number following `/dev/fd/` is not necessarily 3.) The first line of the script properly starts a shell, and the file system containing the script is not mounted with the `nosuid` option.

Running `truss` on the shell script reveals that a call to `open(2)` is failing with error number 6 (`ENXIO`):

<code>open("/dev/fd/3", O_RDONLY)</code>	<code>Err#6 ENXIO</code>
--	--------------------------

Action

`setuid` and `setgid` shell scripts use the file descriptors in `/dev/fd`. The contents of `/dev/fd` are a file descriptor file system (FDFS) and have no connection with floppy disks!

Ensure that the `fdfs` is mounted as `/dev/fd`. Before the machine is next rebooted, the following line should appear in `/etc/vfstab`, exactly like this (with no initial comment symbol):

```
fd - /dev/fd fd - no -
```

It might be possible to remount `/dev/fd` without rebooting by running the following as root:

```
# mount fd /dev/fd
```

Otherwise, to make `setuid`/`setgid` shell scripts available, the machine must be rebooted after editing `/etc/vfstab` as detailed above.

Some administrators, unaware of what `/dev/fd` is for, comment out the entry in `/etc/vfstab` that mounts the FDFS (file descriptor file system). This can go unnoticed until an attempt is made to run a `setuid` or `setgid` shell script.

`/dev/rdisk/c0t6d0s2`: No such file or directory

Cause

When attempting to eject a CD-ROM on a Ultra 450 system, the `eject cdrom` command fails, displaying the error message.

This happens when the CD-ROM is on controller 1, not 0. When using the `eject(1)` command, the CD-ROM "nickname" equates to `/dev/rdisk/c0t6d0s2`. On an Ultra 450, the CD-ROM equates to `/dev/rdisk/c1t6d0s2`. Therefore, using `cdrom` does not work.

Action

Use the following command instead:

```
# eject cdrom0
```

If volume manager `/usr/sbin/vold` is not running, you can use the following:

```
# eject /dev/rdisk/clt6d0s2
```

Note - Make sure that the front panel of the system is unobstructed so the CD-ROM tray is not blocked. Otherwise, the `eject(1)` command appears to hang since the tray is trying to open, but is physically blocked.

Device busy

Cause

An attempt was made to mount a device that was already mounted or to unmount a device containing an active file (such as an open file, a current directory, a mount point, or a running program). This message also occurs when trying to enable accounting that is already enabled.

Action

To unmount a device containing active processes, close all the files under that mount point, quit any programs started from there, and change directories out of that hierarchy. Then try to unmount again.

Technical Notes

Mutexes, semaphores, condition variables, and read/write locks set this error condition to indicate that a lock is held.

The symbolic name for this error is `EBUSY`, `errno=16`.

device busy

Cause

If you perform an `eject cdrom` and then receive the above message, it could be due to a number of problems. Below is a list of things that you can check and do to permit ejection of the CD from the device.

Action

Step A: Ensure that the current directory is not somewhere in the CD:

```
% cd
% eject cdrom
```

Step B: As root:

```
# cd /etc/init.d
# ./volmgt stop
# eject cdrom
```

If this works, then try:

```
# ./volmgt start
```

If this does not work, go to step C.

Step C: As root:

```
# fuser /cdrom
```

Kill any processes you feel you have already terminated. A note of caution: If this is an NFS-mounted CD-ROM and there are other users who access this drive, make sure you know what process you are killing and why.

```
# ./volmgt stop
# ps -ef | grep vold
```

If vold still is running, kill the process.

```
# eject cdrom
```

If this does not work, then:

```
# cd /vol
```

Make sure that dev, dsk, rdsd, rmt are in the directory. If not, probably your /vol directory is corrupt and a reboot might be needed for proper rebuild.

Step D: The last three options are:

1. Reboot.
2. If the CD drive is external to the system, try power cycling the drive and pressing the eject button.

3. If all else fails and the CD-ROM is external, on the right hand side of the eject button is a small hole into which you can insert a small straight device which forces manual ejection of the caddy.

`/dev/rdisk/string`: CAN'T CHECK FILE SYSTEM.

Cause

The system cannot automatically clean (preen) this file system because it appears to be set up incorrectly or is having hard-disk problems. This message asks that you run `fsck(1M)` manually, since data corruption might already have occurred.

Action

Run `fsck` to clean the file system in question. For proper procedures, refer to “`/dev/rdisk/string`: UNEXPECTED INCONSISTENCY; RUN `fsck` MANUALLY.” on page 88.

`/dev/rdisk/string`: UNEXPECTED INCONSISTENCY; RUN `fsck` MANUALLY.

Cause

During a boot, the `/etc/rcS` script runs the `fsck(1M)` command to check the integrity of file systems marked “fsck” in `/etc/vfstab`. If `fsck(1M)` cannot repair a file system automatically, it interrupts the boot procedure and produces this message. When `fsck(1M)` gets into this state, it cannot repair a file system without losing one or more files, so it wants to defer this responsibility to you, the administrator. Data corruption has probably already occurred.

Action

First run `fsck -n` on the file system, to see how many and what type of problems exist. Then run `fsck(1M)` again to repair the file system. If you have a recent backup of the file system, you can generally answer “y” to all the `fsck(1M)` questions. It is a good idea to keep a record of all problematic files and inode numbers for later reference. To run `fsck(1M)` yourself, specify options as recommended by the boot script. For example:

```
# fsck /dev/rdisk/c0t4d0s0
```

Usually the files lost during `fsck(1M)` repair are those that were created just before a crash or power outage, and they cannot be recovered. If you lose important files, you can recover them from backup tapes.

If you do not have a backup, ask an expert to run `fsck(1M)` for you.

See Also

For more information on file checking, see the section on checking file system integrity in the *System Administration Guide, Volume 1*.

Directory not empty

Cause

The directory operation that was attempted, such as directory removal with `rmdir(1)`, can be performed only on an empty directory.

Action

To remove the directory, first remove all the files that it contains. A quick way to remove a non-empty directory hierarchy is with the `rm -r` command.

Technical Notes

The symbolic name for this error is `ENOTEMPTY`, `errno=93`.

Disc quota exceeded

Cause

The user's disk limit has been exceeded on a user file system, usually because a file was just created or enlarged beyond the limit. This almost always refers to a magnetic disk, and not to an optical disc. Any data created after this condition occurs can be lost.

Action

The user can delete files to bring disk usage under the limit, or the server administrator can use the `edquota(1M)` command to increase the user's disk limit.

Technical Notes

The symbolic name for this error is `EDQUOT, errno=49`.

disk does not appear to be prepared for encapsulation

Cause

When attempting to encapsulate the root disk during `vxinstall`, the user gets this error message.

The disk was sliced properly for encapsulation; however, the `prtvtoc` command was non-executable, because the permissions had been changed.

diskN not unique

Cause

During boot, the system displays `disk0 not unique`. The error happens before the kernel loads.

Action

There are more than one `devalias` entries for `disk0`. Use `devalias` at the `OK` prompt to see the entries.

To remove the duplicate, run the following command at the `OK` prompt:

```
nvunalias disk0
```

and reset the system.

dlopen (libxfn.so) failed

Cause

The `SUNWfns` package was left out of the End User Cluster. If only this cluster is installed and automounter is used, it fails with the above message. `libxfn.so` is the shared library for the Federated Naming System.

Action

Install package `SUNWfns` from of the distribution CD.

driver is already installed

Cause

The SunPC™ 4.1 package and then necessary patches (102924) were added. When trying to run `sunpc_install`, the user got the above error message. `prtconf(1M)` shows that the driver is not attached, and `modinfo(1M)` displays 4 modules.

After removing the package, backing out the patch, and reinstalling, the user still received the same error message.

Action

SunPC had previously been installed on the system. When removing the package with the `pkgrm(1M)` command, not all components were removed, because `pkgrm(1M)` is not aware of changes made by the `sunpc_install` script.

To resolve this problem it is necessary to remove sections in the files pertaining to SunPC: `/etc/devlink.tab`, `/etc/driver_aliases`, and `/etc/rc2.d/S10storekernname`, and then reinstall the package.

dtmail: cannot open mailfile on 2.5.1 /var/mail server

Cause

/var/mail is mounted onto client machine A, which is running CDE 1.2 (the Solaris 2.6 release), from machine B, a server running the Solaris 2.5.1 release.

OpenWindow's mailtool can read/write mailfiles on the server without any problems. However, CDE's dtmail does not open the mailbox.

Action

The bug's permissions and ownership have to be checked. The mail directory should have the following permissions:

```
skywalker$ ls -lad /var/mail
drwxrwsrwt  3 root      mail          512 Feb 10 14:40 /var/mail/
```

while the mailbox itself should look something like this:

```
-rw-----  1 zvinakis mail      3206838 Feb 19 11:51 /var/mail/zvinakis
```

If the directory's permissions are not set properly, issue these commands on the mail server:

```
chmod a+t /var/mail
chmod g+s /var/mail
```

If the permissions (or group) are not correct on the mailbox itself, using "joe" as an example mailbox, type:

```
chgrp mail /var/mail/joe
```

To change the permissions, type:

```
chmod 600 /var/mail/joe
```

DUMP: Cannot open dump device ‘/dev/rdisk/c2t0d0s1’: Permission denied

Cause

When using `ufsdump(1M)` as user `sys` (UID 3) on a disk drive in an SSA, the `ufsdump(1M)` command fails with this message.

Action

Six-hundred (600) permissions were created on the SSD “instance path” for a disk drive in an SSA. For a non-root user to read them, there should have been 0640. For example, if you see this:

```
# ls -lL /dev/rdisk/c2t0d0s1
crw----- 1 root sys 192,241 Jul 10 1996 /dev/rdisk/c2t0d0s1
```

Change it to this:

```
crw-r----- 1 root sys 192,241 Jul 10 1996 /dev/rdisk/c2t0d0s1
```

You might also want to add the following line:

```
ssd:* 0640 root sys
```

to the `/etc/minor_perm` file, so subsequently added arrays do not have the same problem.

dumptm: Cannot open ‘/dev/rmt/string’: Device busy

Cause

During file system backup, the `dump` program cannot open the tape drive, because some other process is holding it open.

Action

Find the process that has the tape drive open, and either `kill(1)` the process or wait for it to finish.

```
# ps -ef | grep /dev/rmt
# kill -9 processID
```

DUP/BAD I=i OWNER=o MODE=m SIZE=s MTIME=t FILE=f REMOVE?

Cause

During phase 1, `fsck(1M)` found duplicate blocks or bad blocks associated with the file or directory specified after `FILE=` whose inode number appears after `I=` (with other information).

Action

To remove this file or directory, answer “yes.” If you have to remove more than a few files in this manner, data can be lost. Therefore, it might be preferable to restore the file system from backup tapes.

See Also

For more information on checking file systems, see the section on checking file system integrity in the *System Administration Guide, Volume 1*.

int DUP I=*int*

Cause

Upon detecting a block that is already claimed by another inode, `fsck(1M)` prints the duplicate block number and its containing inode (after `I=`).

Action

In `fsck(1M)` phases 2 and 4, you decide whether or not to clear these bad blocks. Before committing to repair with `fsck(1M)`, you could determine which file contains this inode by passing the inode number to the `ncheck(1M)` command:

```
# ncheck -inum filesystem
```

See Also

For more information, see the chapter on checking file system integrity in the *System Administration Guide, Volume 1*.

"E"

Encapsulation of root disk is not supported on systems with old PROM versions

Cause

When encapsulating the root disk with Volume Manager, this error message is printed.

Action

This error message probably has nothing at all to do with the user's system PROM version. It most likely is related to the fact that the file `/dev/vx/config` (and the pseudo device that it is linked to) does not exist on the system. A few things could contribute to this file not being created:

1. Make sure these lines are in the `/etc/system` file:

```
forceload: drv/vxdmp      (only needed for SEVM 2.5 and above)
forceload: drv/vxio
forceload: drv/vxspec
```

2. Make sure the `vx` entries are in the `/etc/name_to_major` file.

```
grep vx /etc/name_to_major
```

This should come back with 2 or 3 lines (`vxio` and `vxspec`, also `vxdump`, if running SEVM 2.5 and above). The major number values might be different from machine to machine; however, if the entries are in there, that should be sufficient.

3. If you have not performed a `boot -r` since doing a `pkgadd` to the Volume Manager software, this might have contributed to the problem.

ENOMEM The available data space is not large enough to accommodate the shared memory segment

Cause

ENOMEM errors occur after 80 segments have been allocated by Lotus Notes.

Action

The design and implementation of the Solaris ISM (intimate shared memory)—which limits the number of shared memory segments that can be attached to a particular process—caused the ENOMEM failures to the Lotus Notes application.

There is a limit because all shared memory segments are attached in the intimate shared memory (ISM) mode by a system variable that is set in the `shmsys:share_page_table` system file.

When a shared memory segment is attached in ISM mode, the OS locks that segment into physical memory and arranges the virtual/physical address mappings such that only one copy of the mapping information is shared among all attaching processes. To accomplish this, the OS requires that the virtual starting address of the segment be aligned on a 16 Mbytes (hex `0x1000000`) = 16777216-bytes address boundary.

The NULL address lets the system decide what virtual address the segment should be attached to. The system also assigns addresses at `0x3000000` apart, unless forced to attach addresses at `0x1000000` apart.

A `sun4d` could create and attach up to 220 1-Mbyte ISM segments, and a `sun4m` could create and attach up to 235 1-Mbyte ISM segments, providing the segments were `0x1000000` apart.

Having established that ISM is the cause of the limit, below are some options:

First, the limit only gives Lotus Notes the ability to attach a total of 80 Mbytes of shared memory. By increasing the segment size to 10 Mbyte, as Lotus has already

recommended, 8 ISM segments can handle the load previously needing 80 1-Mbyte segments. The load could conceivably grow to 800 Mbytes now without running into the ISM addressing limit.

Second, the `share_page_table` (ISM) flag could be turned off. This would give a sun4m the ability to create in excess of 3000 1-Mbyte segments. The problem here is that ISM does improve the performance of shared memory accesses, and, if the user intends to move up to 2.5.1, ISM is required to get around another set of problems related to shared memory loads of this kind.

Third, Lotus could change the Notes server so that it kept track of the attach addresses and always attached at 0x1000000 boundary addresses, instead of having the system default to the 0x3000000 address boundary. This would allow a Notes server to grow to 235 segments on a sun4m.

error 15 initializing

Cause

It is caused by a bad `/boot` or 4.1 on ss2 - level 15 interrupt.

Error 76

Cause

This error is RFS-specific. The server is telling the client that a process has transferred back from mount point.

Technical Notes

The symbolic name for this error is `EDOTDOT`, `errno=76`.

Error 88

Cause

This error is caused by an illegal byte sequence.

Action

You need to handle multiple characters as a single character.

Technical Notes

The symbolic name for this error is `EILSEQ`, `errno=88`.

error code 2: access violation

Cause

The user receives this message when trying to do a `tftp get`.

Action

Do not use a relative path when using `tftp`. For example:

```
tftp> get /tftpboot/testfile
```

fails, and

```
tftp> get testfile
```

succeeds.

error: DPS has not initialized or server connection failed

Cause

This message appears when trying to run AnswerBook on a generic X11 window server or on a generic X terminal.

Action

Running AnswerBook requires Display PostScript (DPS), or a NeWS server, or the Adobe DPS NS remote display software. Additionally, a complete LaserWriter II Type-1 font set (including Palatino) should be installed on the X server. To find out if

the X server has DPS, run *xdpyinfo(1)* to verify the presence of an "Adobe-DPS-Extension" line. X servers without this line do not know about DPS.

Error: Error adding OS service Solaris 2.6 sparc sun4u:

Cause

While trying to add OS services to a newly installed Solaris 2.6 environment and using Solstice Adminsuite 2.3, the process fails with the following error message:

```
Error: Error adding OS service Solaris 2.6 sparc sun4u:
inconsistent revision, installed package SUNWpppk revision 3.0.1
does not match revision 11.6.0,REV=1997.07.15.21.46 for sparc
architecture.
```

This error is caused by the optional Solstice PPP 3.0.1 packages from the "Solaris Server Intranet Extension" CD-ROM installed on the system.

Action

As a workaround, remove the PPP 3.0.1 packages and replace them with the PPP packages from the Solaris 2.6 release CD-ROM. For example:

```
# pkgrm SUNWlicsw SUNWlit SUNWpppk SUNWpppm SUNWpppr SUNWppps SUNWpppu
:
: {package remove info}
:
# cd /cdrom/cdrom0/s0/Solaris_2.6/Product
# pkgadd -d . SUNWapppr SUNWapppu SUNWpppk
:
: {package add info}
:
```

Then, use Adminsuite to add the OS services, which should then work without error.

Note - If the Solstice PPP 3.0.1 package is configured and currently in use on the system, the user should save any of the previously entered PPP configuration information for restoration after the OS services have been installed. (`pkgrm(1M)` the 3 PPP packages installed from the 2.6 CD release, and again `pkgadd(1M)` all of the PPP packages from the Intranet Extension CD-ROM, then redo the configuration.) If the Solstice PPP 3.0.1 package was not used on the system, there is no reason to reinstall it. Use `/usr/bin/pkginfo` to check the installed packages.

This is documented in Chapter 9 of the *Solaris Server Intranet Extension Installation and Release Notes Solaris 2.6* manual.

Error Host Unknown:

Cause

In this case, the user is on Windows 95, running PC-NFS pro2.0. The user uses `ping(1M)` to reach another computer on the network. `ping(1M)` returns Host Unknown.

This happens when name services are not set up correctly.

Action

1. Click the Windows 95 Start button, click Programs, click PC-NFSpro, then click Configuration.
2. Click TCP/IP and make sure all settings are entered correctly.
3. If NIS is enabled, click Configure NIS and make sure the NIS domain and server names are correct.
4. If DNS is enabled, click Configure DNS and make sure the DNS domain and server names are correct.
5. Click `edit hosts` and add the name and IP address of the machine you are trying to `ping(1M)`, along with the authentication server.

If you make any changes, click OK, then click Save and Exit on the Configuration dialog box. Shut down and restart Windows 95.

ERROR: missing file arg (cm3)

Cause

An attempt was made to run some `sccs(1)` operation that requires a file name, such as `create`, `edit`, `delget`, or `pvt`.

Action

Supply the appropriate file name after the SCCS operation.

ERROR [SCCS/*s.string*]: 'SCCS/*p.string*' nonexistent (ut4)

Cause

An attempt was made to `sccs(1) edit` or `sccs get` a file that was not yet under SCCS control.

Action

Run `sccs(1) create` on that file to place it under SCCS control.

ERROR [SCCS/*s.string*]: writable '*string*' exists (ge4)

Cause

An attempt was made to `sccs(1) edit` a file that is writable, probably because it was already checked out.

Action

Run `sccs(1) info` to see who has the file checked out. If it is you, go ahead and edit it. If it is somebody else, ask that person to check-in the file.

Error: you don't have a license to run this program

Cause

The user tries to mount the `/export` file system with Volume Manager 2.1.1 and gets this message.

Action

Run `vxserial -p` to print the available Volume Manager licenses in the system.

Also, check the `/etc/vfstab` file to make sure that the file system is not a `vxfs` file system.

esp0: data transfer overrun

Cause

When a user tries to mount a CD-ROM on a third-party CD drive, `mount(1M)` fails with the above error, followed by the `sr0: SCSI transport failed message`. The CD drive probably comes from a vendor unknown to the system.

Action

Third-party CD drives generally have an 8192 block size, as opposed to the 512 block size on supported Sun drives. Check with the vendor to see if any special configuration is possible to allow the drive to operate on a Sun workstation.

ether_hostton errors from cb_reset

Cause

You issue `cb_reset` on an SSP and get the following:

```
cb_reset
Resetting host snax-cb0...
warning: ether_hostton(SrcHost:beer): Bad file number
```

(continued)

(Continuation)

```
warning: ether_hosttton(SrcHost:beer): Bad file number
warning: ether_hosttton(SrcHost:beer): Bad file number
Resetting host snax-cbl...
warning: ether_hosttton(SrcHost:beer): Resource temporarily unavailable
warning: ether_hosttton(SrcHost:beer): Resource temporarily unavailable
warning: ether_hosttton(SrcHost:beer): Resource temporarily unavailable
snax-cb0 is ready...
snax-cbl is ready...
```

The `cb_reset` actually completes, but the error messages are annoying.

Action

`/etc/nodename` is probably incorrect. The following details are from a machine getting this error message. Note that `/etc/nodename` contains an alias to the real name of the SSP. To correct the problem, edit `/etc/nodename` to match the true name and reboot.

```
# cat /etc/nodename
beer

# cat /etc/hostname.qfe0
snax-ssp

# cat /etc/hosts 127.0.0.1 localhost
129.153.49.179 snax-ssp beer loghost

# cat /etc/ethers
8:0:20:87:58:a5 snax-ssp beer
```

Event not found

Cause

This C shell message indicates that a user tried to repeat a command from the history list, but that command or number does not exist in the list.

Action

Run the C shell `history(1)` command to display recent events in the history list. If a user often tries to run commands that have disappeared from the history list, make the list longer by setting `history(1)` to a higher value.

See Also

For more information about the C shell, see `csh(1)`.

EXCESSIVE BAD BLKS I=*int* CONTINUE?

Cause

During phase 1, `fsck(1M)` found more than 10 bad (out-of-range) blocks associated with the specified inode number.

Action

With this many bad blocks, it might be preferable to restore the file system from backup tapes.

See Also

For more information on bad blocks, see the section on checking file system integrity in the *System Administration Guide, Volume 1*. If you are using AnswerBook online documentation, "bad blocks" is a good search string.

EXCESSIVE DUP BLKS I=*int* CONTINUE?

Cause

During phase 1, `fsck(1M)` found more than 10 duplicate (previously claimed) blocks associated with the specified inode number.

Action

With this many duplicate blocks, it might be preferable to restore the file system from backup tapes.

See Also

For more information on blocks, see the section on checking file system integrity in the *System Administration Guide, Volume 1*. If you are using AnswerBook online documentation, "bad blocks" is a good search string.

Exec format error

Cause

This often happens when trying to run software compiled for different systems or architectures, such as when executing the programs on a SunOS 4.1 system, or when trying to execute SPARC-specific programs on an IA machine. This error can also occur if the Binary Compatibility Package was not installed.

Action

Make sure that the software matches the architecture and system you are using. The `file(1)` command can help you determine the target architecture. If you are using SunOS 4.1 software on a later release, make sure that the Binary Compatibility Package is installed. You can check for it using this command:

```
$ pkginfo | grep SUNWbcp
```

Technical Notes

A request was made to execute a file that, although it has the appropriate permissions, does not start with a valid format.

The symbolic name for this error is `ENOEXEC`, `errno=8`.

See Also

See the `a.out(4)` man page for a description of executable files.

"F"

failed to initialize adapter

Cause

When using an Adaptec AHA-154x Cx SCSI HBA during installation of the IA release, you might see a message during the MDB device probe that says `failed to initialize adapter` after the probe has correctly identified the card. There are a variety of reasons for this error, but in all cases the error is because of misconfiguring the card.

Action

To correct the problem, press `Ctrl-A` during boot to enter the 154x BIOS configuration utility. Choose the `Configure/View Host Adapter Settings` option, then press the `F6` key to return the adapter to its factory default settings.

After doing this, reconfigure the adapter using the instructions contained in the IA Device Configuration Guide or Driver Update Guide, if applicable. It is especially important that the adapter be configured to use DMA 6. Note that it must be changed from the default of DMA 5.

Failed to Load Security Policy: Invalid argument

Cause

While installing a policy from the GUI (or the command line) the following error message is displayed:

```
default.W: Security Policy Script generated into default.pf
default:
Compiled OK.

Installing Security Policy default on all.all@lab-netra
Failed to Load Security Policy: Invalid argument <----- !!
Installing Security Policy on localhost(localhost) failed
```

If you truss the policy load, you receive the following:

```
truss -o /tmp/truss -f -vall -rall -wall /etc/fw/bin/fw
/etc/fw/conf/default.W
```

The following is near the end of the truss:

```
1226: open("/dev/fw0", O_RDWR|O_NONBLOCK) = 7
1226: ioctl(7, 0xC0C07A18, 0xEFFBCA0) Err#22 EINVAL
```

This problem is caused by someone "plumbing" or configuring a new Ethernet interface after Firewall-1 has already started (that is, plumbing an interface by hand after the system has been booted).

Action

This error can be resolved by configuring the interface to configure automatically at boot time (for example, by creating a `/etc/hostname.qe0` file) and rebooting the system.

The following is another solution:

```
/etc/fw/bin/fwstop # Stop firewall
modinfo | grep fw # Get kernel module ID

85 f5e19000 3cc0c 51 1 fw (fw)

modunload -i 85 # Unload kernel module

/etc/fw/bin/fwstart # Restart firewall
```

The policy installs correctly now with the following:

```
# ./fw load ../conf/default.W
default.W: Security Policy Script generated into default.pf
default:
Compiled OK.
```

fast access mmu error

Cause

The user receives this message while trying to boot the Ultra over the network by using the FDDI 5.0 card.

Action

Do the following:

1. Setenv auto-boot? to false.
2. Reset the system.
3. Boot the FDDI card.

fbconsole: ioctl SRIOCSREDIR: Device Busy.

Cause

When starting OpenWindows from the command line, the following error message is echoed on the Solaris "Welcome" screen: fbconsole: ioctl SRIOCSREDIR: Device Busy

Once inside OpenWindows, the following message is displayed in the background windows and when starting `cmdtool -C`:

```
SYSTEM WARNING: Object 0x340f8, Device busy, ioctl SRIOCSREDIR
returned -1, attempt to make tty the console failed (Tty
package)
```

Action

OpenWindows was probably started in the background (using the "&"). Exit OpenWindows, and run the command in foreground:

```
/usr/openwin/bin/openwin
```

If this does not help, then perhaps some daemon or process is "holding" the console. Type the command: `fuser /dev/console`.

A list of process IDs is returned. Examine these processes to determine if an application has hold of the console (using the `ps(1)` command helps).

fd0: unformatted diskette or no diskette in the drive

Cause

This message appears on the system console to indicate that the floppy driver `fd(4)` could not read the label on a diskette. Usually this is either because a new diskette has not yet been formatted, or a formatted diskette has become corrupted. This message often appears along with `read failed` and `bad format` messages after `volcheck(1)` has been run.

Action

If you are certain that the diskette contains no data, run `fdformat -d` to format the diskette in DOS format. (You can also format a diskette in UFS format if you like, although then it cannot be transported to most other systems.) When the diskette is formatted, you can write on it, if it has not been corrupted beyond repair.

File descriptor in bad state

Cause

Either a file descriptor refers to no open file or a read request was made to a file that is open only for writing.

Technical Notes

The symbolic name for this error is `EBADF`, `errno=81`.

File exists

Cause

The name of an existing file was mentioned in an inappropriate context. For example, establishing a link to an existing file, or overwriting an existing file are not allowed when the `csh(1)` `noclobber` option is set.

Action

Look at the names of files in the directory, then try again with a different name or after renaming or removing the existing file.

Technical Notes

The symbolic name for this error is `EEXIST`, `errno=17`.

File locking deadlock

Cause

This is a programming problem and, in some cases, is unavoidable.

Action

All a user can do is restart the program and hope deadlock does not reoccur.

Technical Notes

In the file locking subsystem, two processes tried to modify some lock at the same time. In the multi-threading subsystem, two threads became deadlocked and could not continue. When a program using the threads library encounters this error, it should restart the deadlocked threads.

The symbolic name for this error is `EDEADLOCK`, `errno=56`.

File name too long

Cause

The specified file name has too many characters.

Action

If a file name or path name component is too long, devise a shorter name. If the total path name is longer than `PATH_MAX` characters, first change to an intermediate

directory, then specify a shorter path name. Newly created data will be lost unless written to another file with a shorter name.

Technical Notes

In a UFS or NFS-mounted UFS file system, the length of a path name component exceeds `MAXNAMLEN` (255) characters, or the total length of the path name exceeds `PATH_MAX` (1024) characters. In a System V file system, the length of a path name component exceeds `NAME_MAX` (14) characters while no-truncation mode is in effect. These values are defined in the `/usr/include/limits.h` file.

The symbolic name for this error is `ENAMETOOLONG`, `errno=78`.

file system full

Cause

This error message is seen during a login. The login fails with the message `No utmpx entry`.

See Also

Refer to “No utmpx entry” on page 179.

FILE SYSTEM STATE IN SUPERBLOCK IS WRONG; FIX?

Cause

The `fsck(1M)` command has just checked a file system, and has determined that the file system is clean. The file system’s super block, however, still thinks the file system is “dirty” in some way.

Action

If you believe that the file system is adequately repaired, answer “yes” to mark the file system as clean.

Technical Notes

Different "dirty" file system types are listed in `/usr/include/sys/fs/ufs_fs.h`, and include `FSACTIVE`, `FSBAD`, `FSFIX`, `FSLOG`, and `FSSUSPEND`.

See Also

For more information on super blocks, see the section on checking file system integrity in the *System Administration Guide, Volume 1*. If you are using AnswerBook online documentation, "bad super block" is a good search string.

File table overflow

Cause

The kernel file table is full, because too many files are open on the system. Temporarily, no more files can be opened. New data created under this condition will probably be lost.

Action

Simply waiting often gives the system time to close files. However, if this message occurs often, reconfigure the kernel to allow more open files. To increase the size of the file table, increase the value of `MAXUSERS` in the `/etc/system` file. The default `MAXUSERS` value is the amount of main memory in Mbytes, minus 2.

Technical Notes

The symbolic name for this error is `ENFILE`, `errno=23`.

File too large

Cause

The file size exceeded the limit specified by `ulimit(1)`, or the file size exceeds the maximum supported by the file system. New data created under this condition can probably be lost.

Action

In the C shell, use the `limit(1)` command to see or set the default file size. In the Bourne or Korn shells, use the `ulimit -a` command. Even when the shells claim that the file size is unlimited, in fact the system limit is `FCHR_MAX` (usually 1 Gbyte).

Technical Notes

The symbolic name for this error is `EFBIG`, `errno=27`.

filemgr: mknod: Permission denied

Cause

File Manager issues this message and fails to come up whenever the `/tmp/.removable` directory is owned by another user and is not in `1777` mode. This can happen, for example, when multiple users share a workstation.

Action

Have the original owner use `chmod(1)` to change the mode of this file back to `1777`, its default creation mode. Rebooting the workstation also resolves this problem.

Technical Notes

This is a known problem that was fixed in the Solaris 2.4 release.

FREE BLK COUNT(S) WRONG IN SUPERBLK SALVAGE?

Cause

During phase 5, `fsck(1M)` detected that the actual number of free blocks in the file system did not match the super block's free block count. The `df(1M)` command accesses this free block count when measuring file system capacity.

Action

Generally you can answer "yes" to this question without harming the file system.

See Also

For more information on super blocks, see the section on checking file system integrity in the *System Administration Guide, Volume 1*. If you are using AnswerBook online documentation, "bad super block" is a good search string.

fsck & ufsdump - cannot read block/sector errors

Cause

If you have received the following messages from `fsck(1M)`:

```
CANNOT READ: BLK 196896
CONTINUE? y
THE FOLLOWING SECTORS COULD NOT BE READ: 196896 196897 196898 196899
```

Or the following warnings from `ufsdump(1M)`:

```
DUMP: Warning - cannot read sector 164016 of /dev/vx/rdisk/newdg/vol02
DUMP: Warning - cannot read sector 164017 of /dev/vx/rdisk/newdg/vol02
DUMP: Warning - cannot read sector 164018 of /dev/vx/rdisk/newdg/vol02
```

It could be that the size of this file system in this volume does not match the size of the regular file system.

Action

To check this, follow the example below:

Run the command:

```
# fstyp -v /dev/vx/rdisk/newdg/vol02 | head -30 | grep ncg
```

to print the following line (disregard any error or warning messages you might get):

```
ncg      25      size    102400  blocks  95983
```

Disregard everything but the number after the word `size`. This number tells you the file system is 102,400 Kbytes in length.

Next, find out the size of the volume. Run the command:

```
# vxprint -g newdg -vt vol02
```

which prints:

```
v NAME   USETYPE  KSTATE  STATE  LENGTH  READPOL  PREFPLEX
v vol02  fsgen    ENABLED ACTIVE  163840  SELECT
```

From this, you can see the volume is 163,840 sectors (divide this number by 2 to get it into Kbytes) or 81,920 Kbytes. As you can see from this example, the volume (80 Mbytes) is much smaller than the file system (100 Mbytes). This should be rectified immediately to avoid or minimize data loss.

To resolve this problem, back up the data as best you can, then either create a new volume or `newfs` this one and restore the data.

This problem can also occur on a DiskSuite metadvice. The difference is that you need to check the size of the metadvice using the `metastat` command. The `metastat` command shows the size of the metadvice in sectors, just like the `vxprint` does.

fsck: Can't open /dev/dsk/string

Cause

The `fsck(1M)` command cannot open the disk device, because although a similar file system exists, the partition specified does not.

Action

Run the `mount(1M)` or the `format(1M)` command to see what file systems are configured on the machine. Then run `fsck(1M)` again on an existing partition.

fsck: Can't stat /dev/dsk/string

Cause

The `fsck(1M)` command cannot open the disk device, because the specified file system does not exist.

Action

Run the `mount(1M)` or the `format(1M)` command to see what file systems are configured on the machine. Then run `fsck(1M)` again on an existing file system.

ftp: ftp/tcp: unknown service

Cause

The user received this error while using no naming service. The services file looked fine. The user could FTP as `root`, but not as a normal user.

Action

The permissions on the `/etc/services` file were wrong. To correct the problem, the user changed them to read access for everyone (644).

fw_ipinput: q fc5fddc0:illegal interface

Cause

The FW-1 kernel module displays this error message when a new network interface has been added to the FW-1 system while `fwd` is running.

Action

To resolve this problem, run the following to reinstall the FW kernel and the security policy:

```
# fw ctl uninstall
# fw ctl install
# fw fetch localhost
```

FW1: log message queue is full

Cause

The console reports FW1: Log message queue is full.

The message log is a queue that keeps all the firewall's event logs until FW-1 finishes processing them. If too many logs arrive, the buffer is full and the message FW-1: log message queue is full appears. It usually happens on loaded systems or firewalls that handle many network connections.

Action

Below are some suggestions to stop this warning message:

- Reduce the amount of logging in the security policy.

Note - ACCOUNTING logging is very heavy. Reducing logging from LONG to SHORT also helps.

- Increase the internal memory allocated to the FW kernel module. The default amount of memory is 524K. To increase to 1Mbyte, add the statement below to `/etc/system` and reboot:

```
set fw:fwmem=0x100000
```

- Set the Excessive Log Grace Period to 0. This is set through Properties -> Logging and Alerting. You must then reinstall the security policy for the change to take affect. The drawback for setting the Excessive Log Grace Period to 0: Your log now includes similar packets received at approximately the same time. When it was not zero, they were hidden (see *Managing FireWall-1 Using the OpenLook GUI*, p. 104). Thus, no packet disappears from the log, so your log might be a little bit bigger, but apart from that, no problem.
- Use `Renice fwd` for a higher priority. The default priority of the FW daemon is 0 (like most processes). To raise the priority you must give a negative priority, depending on the load on your system. See the man page on `nice(1)` for more information.

fwm: no license

Cause

Firewall-1 version 2.1 produces this message when the `fwstart` command is issued or when `fwm` is started from the command line.

There are two possible reasons for this:

1. When a firewall module is installed without a control station on the same machine, the messages are displayed on the console (under UNIX) or in the event log (under WinNT).
2. The messages might be legitimate. You might find that `fwm` has not started and you cannot do some crucial tasks. One possible problem: The license might be issued for the wrong host ID.

Action

Make sure the license daemon is running on the server. Then, consider the following cases:

Case one: As a workaround, ignore the present messages and get an upgrade to 2.1c or above.

Case two: To check for a misassigned license, run the command `hostid(1)`. Your `hostid` is displayed.

Next, run the command `fw printlic` to see output similar to the following:

```
This is FireWall-1 Version 2.1
Type           Expiration Features
id-649f152b    never          stdlight
```

The first field should list the correct `hostid`. Also check the expiration date and the features. A list of what is included with the features is provided in INFODOC 13215. If you find any inconsistencies, call the Sun License and Password Center and get a license reissued. Have your host ID and serial number ready.

fwskip_parse_headers: invalid peer n

Cause

In Firewall-1, the connections encrypted with SKIP are dropped at certain times, specifically near the top of the hour. For example, connections will be dropped from

10:55 to 11:15, then continue working normally until 11:55. These error messages appear on the console in pairs:

```
fwskip_parse_headers: invalid peer n
fw_skip_decrypt: cannot parse headers
```

These error messages are referring to the n counter. The n counter is the absolute number of hours in GMT time. It is included in the SKIP calculations as a safeguard against a playback attack. If the 2 hosts or firewalls exchanging encrypted packets are not in sync with respect to GMT time, they have different n counters and these errors appear.

Action

Keep the clocks on the encrypting hosts within one hour of each other, GMT time.

"G"

giving up

Cause

This message appears in the SCSI log to indicate that a read or write operation has been retried until it timed out. With SCSI disk the time-out period is usually 30 seconds; with tape, the period is usually 20 attempts. Time-out periods are generally coded into the drivers.

Action

Check that all SCSI devices are connected and powered on. Make sure that SCSI target numbers are correct and not in conflict. Verify that all cables are no longer than a total of six meters, and that all SCSI connections are properly terminated.

Technical Notes

The `scsi_log(9F)` routine usually displays messages on the system console and in the `/var/adm/messages` file. Run the `dmesg(1M)` command to see the most recent message buffer.

Graphics Adapter device `/dev/fb` is of unknown type

Cause

The `/dev/fb` driver is either missing or corrupted.

Action

For details, refer to “InitOutput: Error loading module for `/dev/fb`” on page 134.

group.org_dir: NIS+ servers unreachable

Cause

This is the second of three messages that an NIS+ client prints when it cannot locate an NIS+ server on the network.

Action

For details, refer to “hosts.org_dir: NIS+ servers unreachable” on page 123.

"H"

hang console

Cause

Console hangs, but all other operations are working, including `rlogin(1)` and `telnet(1)`. Rebooting the system (by way of a remote shell) clears the problem.

This problem occurs if another window is opened with the `-C` option, causing the console to hang. The other window could be another `cmdtool` window, `shelltool` window, or even an `xterm` window. Only one console window can be active at a time.

Action

The window/process that is causing the problem can be located by using the `ps(1)` command (`auxw` options might be necessary). The process can then be killed. Eliminate the console window running with `-C`, and control returns to the real console.

Machine hung in reboot process: when the user is booting the machine, it hangs at checking file systems.

As a possible workaround, do the following:

1. Boot `miniroot` from tape or CD-ROM.
2. Type: `mkdir mnt`.
3. Mount the root partition to some mount point (`/mnt`).
4. Change the directory to `/mnt/dev`.
5. Make sure the console is located in the `mnt/dev` directory.
6. If not, make the device `std` (`MAKEDEV std`).
7. Halt the system and reboot.

/home/string: No such file or directory

Cause

An attempt was made to change to a user's home directory, but either that user does not exist or the user's file server has not shared (exported) that file system.

Action

To check on the existence of a particular user, run the `ypmatch(1)` or `nismatch(1)` command, specifying the user name and then the `passwd(4)` map.

To export file systems from the remote file server, become superuser on that system and run the `share(1M)` command with the appropriate options. If that system is sharing (exporting) file systems for the first time, also invoke `/etc/init.d/nfs.server start` to begin NFS service.

See Also

For more information on sharing file systems, see the `share_nfs(1M)` man page.

Host is down

Cause

A transport connection failed because the destination host was down. For example, mail delivery was attempted over several days, but the destination machine was not available during any of these attempts.

Action

Report this error to the system administrator for the host. If you are the person responsible for this system, check to see if the machine needs repair or rebooting.

Technical Notes

This error results from status information delivered by the underlying communication interface. If there is no known connection to the host, a different message usually results. For details, refer to “No route to host” on page 175.

The symbolic name for this error is `EHOSTDOWN`, `errno=147`.

host name configuration error

Cause

This is an old `sendmail(1M)` message, which replaced `I refuse to talk to myself` and is now replaced by the `Local configuration error` message.

Action

For details, refer to “554 *hostname...* Local configuration error” on page 42.

hosts.org_dir: NIS+ servers unreachable

Cause

This is the third of three messages that an NIS+ client prints when it cannot locate an NIS+ server on the network.

Action

If other NIS+ clients are behaving normally, check the Ethernet cabling on the workstation showing this message. Note the following differences between architectures:

- On SPARC machines, disconnected network cabling also produces a series of `no carrier` messages.
- On IA machines, the NIS+ messages might be the only indication that network cabling is disconnected.

If many NIS+ clients on the network are giving this message, go to the NIS+ server in question and reboot or repair it, as necessary. When the server machine is back in operation, NIS+ clients give an `NIS server for domain OK` message.

"I"

I can't read your attachments. What mailer are you using?

Cause

The SunView *mailtool(1)* and prior 3.3 OpenWindows *mailtool(1)* produce this message when they cannot cope with an attachment. The attachment is probably in MIME (multipurpose internet mail extensions) format, using `base64` encoding.

Action

To read a mail message containing MIME attachments, use *mailtool(1)* from a system running at least the Solaris 2.3 release. If you are running an earlier version of the Solaris environment, `rlogin(1)` to a system running a later version, set the `DISPLAY` environment variable back to the first system, and run *mailtool* remotely. If those options prove impossible, ask the originator to send the message again using *mailtool(1)*, or using the CDE `dtmail compose File->SendAs->SunMailTool` option.

Technical Notes

Standard MIME attachments with `base64` encoding, for example, produce this message and fail to display in older *mailtool(1)*s.

See Also

Look into using `metamail`, available on the Internet, which allows you to send and receive MIME attachments.

Identifier removed

Cause

This message indicates an error in a System V IPC facility. Most likely a file associated with messaging, semaphores, or shared memory was deleted from the file system where it had been created.

Technical Notes

This error is returned to processes that resume execution after the removal of an identifier from the file system's name space. See `msgctl(2)`, `semctl(2)`, and `shmctl(2)` for details.

The symbolic name for this error is `EIDRM`, `errno=36`.

ie0: Ethernet jammed

Cause

This message can appear on SPARC servers or IA machines with an Intel 82586 Ethernet chip. It indicates that 16 successive transmission attempts failed, causing the driver to give up on the current packet.

Action

If this error occurs sporadically or at busy times, it probably means that the network is saturated. Wait for network traffic to clear. If bottlenecks arise frequently, think about reconfiguring the network or adding subnets.

Another possible cause of this message is a noise source somewhere in the network, such as a loose transceiver connection. Use `snoop(1M)` or a similar program to isolate the problem area, then check and tighten network connectors as necessary.

ie0: no carrier

Cause

This message can appear on SPARC servers or IA machines with an Intel 82586 Ethernet chip. It indicates that the chip has lost input to its carrier detect pin while trying to transmit a packet, causing the packet to be dropped.

Action

Check that the Ethernet connector is not loose or disconnected. Other possible causes include an open circuit somewhere in the network and noise on the carrier detect line from the transceiver. Use `snoop(1M)` or a similar program to isolate the problem area, then check the network connectors and transceivers, as needed.

If pipe/FIFO, don't sleep in stream head

Cause

This is a streams pipe error (not externally visible).

Technical Notes

The symbolic name for this error is `ESTRPIPE`, `errno=92`.

ifconfig: bad address

Cause

System fails to boot with this error message: `ifconfig: bad address`. When coming up to multi-user `ifconfig -a`, it indicates the following:

```
le0: flags=863<UP,BROADCAST,NOTRAILERS,RUNNING,MULTICAST> mtu 1500
inet 0.0.0.0 netmask 0
```

Once up, if this command succeeds, then all is well:

```
# ifconfig le0 inet hostname
```

Action

Check `/etc/hostname.*` for a possible bad entry.

`/etc/hosts` was linked to `/var/named/hosts` and `/var` was a separate file system. Until system comes up in multi-user to mount `/var`, host name could not be resolved to proper IP address.

ifconfig bad address le0

Cause

The user installed the recommended 2.5.1 patches. When booting, `rootuser.sh` presented the following errors:

```
ifconfig bad address le0
le0 arp - revarp failed no rarp replies
bad address hme0
hme0 auto-revarp failed: no rarp replies received.
```

The IP address of interface is set to `0.0.0.0`.

System fails to resolve host IP address from `/etc/host` and no other RARP servers responded to the system's request for its IP address.

Action

If `dns [NOTFOUND=return]` appears before files in `/etc/nsswitch.conf`, `ifconfig` complains at boot-time about bad address. In some cases this can cause the boot to fail.

ifconfig: *host name* bad space address

Cause

When the system is booted, this error message is displayed. The `/etc/nsswitch.conf` file had the following entries for the hosts line:

```
hosts: dns nis [NOTFOUND=return] files
```

Action

Move `files` to the first entry in the list. Now, when the system boots, it resolves the interface names from the `/etc/hosts` file.

ifconfig: SIOCGIFFLAGS: hme0: no such interface

Cause

If you just installed `hme` interface and are now manually configuring it, you could receive this error message when running the following:

```
ifconfig hme0 inet ipaddr netmask + broadcast + -trailers up
```

Action

If there is no `hostname.hme0` file, then the startup scripts do not execute the `ifconfig hme0 plumb` command. The user can either create the `hostname.hme0` file or issue the `ifconfig hme0 plumb` command manually before attempting to configure the interface.

Illegal Instruction

Cause

A process has received a signal indicating that it attempted to execute an instruction that is not allowed by the kernel. This usually results from running programs compiled for a slightly different machine architecture. This message is usually accompanied by a core dump, except on read-only file systems.

Action

If you are booting from a CD-ROM or from the net, check Readme files to make sure you are using an image appropriate for your machine architecture. Run `df(1M)` to make sure there is enough swap space on the system; too little swap space can cause this error. If you recently upgraded your CPU to a new architecture, replace your operating system with one that supports the new architecture (an operating system upgrade might be required).

Technical Notes

Sometimes this condition results from a programming error, such as when a program attempts to execute data as instructions. This condition can also indicate device file corruption on your system.

Illegal instruction "0xhex" was encountered at PC 0xhex

Cause

The machine is trying to boot from a non-boot device, or from a boot device for a different hardware architecture.

Action

If you are booting from the net, check Readme files to make sure you are using a boot image for that architecture. If you are booting from disk, make sure the system is looking at the right disk, which is usually SCSI target 3. If these solutions fail, connect a CD drive to the system and boot from CD-ROM.

Illegal seek

Cause

In this instance, using a pipe (|) on the command line does not work.

Action

Rather than using a pipe on the command line, redirect the output of the first program into a file and run the second program on that file.

Technical Notes

A call to `lseek(2)` was issued to a pipe. This error condition can also be fixed by altering the program to avoid using `lseek(2)`.

The symbolic name for this error is `ESPIPE`, `errno=29`.

Image Tool: Unable to open XIL Library.

Cause

This message follows multiple multi-line `XilDefaultErrorFunc` errors, indicating that ImageTool could not locate the X Imaging Library. Many OpenWindows and CDE deskset programs require XIL.

Action

Run `pkginfo(1)` to determine what packages are installed on the system. If the following packages are not present, install them from the CD-ROM or over the net: `SUNWxildg`, `SUNWxiler`, `SUNWxilow`, and `SUNWxilrt`.

Inappropriate ioctl for device

Cause

This is a programming error.

Action

Ask the program's author to fix this condition. The program needs to be changed so it employs a device driver that can accept special character device controls.

Technical Notes

The `ioctl(2)` system call was given as an argument for a file that is not a special character device. This message replaces the traditional, but puzzling `Not a typewriter` message.

The symbolic name for this error is `ENOTTY`, `errno=25`.

INCORRECT BLOCK COUNT I=*int* (should be *int*) CORRECT?

Cause

During phase 1, `fsck(1M)` determined that the specified inode pointed to a number of bad or duplicate blocks. The block count should be corrected to the actual number shown.

Action

Generally you can answer "yes" to this question without harming the file system.

See Also

For more information on bad blocks, see the section on checking file system integrity in the *System Administration Guide, Volume 1*.

index failed:full:index preceded by saveset name

Cause

This is a server that has several clients. It seems that when the backup kicks off, many of the savesets fail with the message listed below:

```
godzilla                               index failed:full:index
* godzilla:index 2 retries attempted
* godzilla:index sh: save: not found
```

Action

Edit the `/etc/init.d/networker` file and change the `nsrexecd` startup line to include a `-p` option to specify this command search path:

```
(/usr/sbin/nsr/nsrexecd -s masters -p /usr/sbin/nsr ) > /dev/console
```

inetd[*int*]: execv /usr/sbin/in.uucpd: No such file or directory

Cause

This message indicates that the Internet services daemon, `inetd(1M)`, tried to start up the UUCP service without the UUCP daemon existing on the system.

Action

The `SUNWbnuu` package must be installed before the machine can run UUCP. Run `pkgadd(1M)` to install this package from the distribution CD-ROM or over the network.

inetd[*int*]: *string*/tcp: unknown service

Cause

This message indicates that the Internet services daemon, `inetd(1M)`, could not locate the TCP service specified after the first colon.

Action

Check the current machine's `/etc/services` file, and the NIS `services` map, to see if the service is described. To start this service, add an appropriate entry into the `/etc/services` file and possibly the `services` map as well. Note that NIS+ does not consult the local `/etc/services` file unless you put files right after `nisplus` on the `services` line of the system's `/etc/nsswitch.conf` file.

If you do not want to start this service, edit the system's `/etc/inetd.conf` file and delete the entry that tries to start it up.

See Also

For more information about NIS+, see the *NIS+ and FNS Administration Guide*.

inetd[*int*]: *string*/udp: unknown service

Cause

This message indicates that the Internet services daemon, `inetd(1M)`, could not locate the UDP service specified after the first colon.

Action

For a solution, refer to “inetd[*int*]: *string*/tcp: unknown service” on page 132.

inetd: Too many open files

Cause

This message can appear when someone runs a command from the shell or uses a third-party application. The `sar(1)` command does not indicate that the system-wide open file limit has been exceeded.

The probable cause of this message is that the shell limit has been exceeded. The default open file limit is 64, but it can be raised to 256.

Action

For a solution, refer to “Too many open files” on page 240.

INIT: Cannot create /var/adm/utmpx

Cause

This console message indicates that `init(1M)` cannot write in the `/var` directory, which is usually part of the `/` (root) file system. Some other messages follow, and the system usually comes up single-user. The problem is often that `/` or `/var` is mounted read-only. Sometimes a brief power outage leaves the system believing that many file systems are still mounted.

Action

If `/var` is a separate file system on the machine and is not yet mounted, mount it now. If the file system containing `/var` is mounted read-only, remount it read-write with a command similar to this:

```
# mount -o rw,remount /
```

Then type Control-D and try to bring up the system multi-user. If that fails, the root file system is probably corrupted. Run `fsck(1M)` on the root file system, halt the machine, power cycle the CPU, and wait for the system to reboot. Should this problem still occur, restore the root file system from backup tapes, or re-install the system from net or CD-ROM to replace the root file system.

InitOutput: Error loading module for `/dev/fb`

Cause

This fatal X server error message indicates that `/dev/fb`, the "dumb frame buffer," is either missing or corrupted. It is usually followed by a `giving up` message and a few `xinit(1)` errors.

Action

If other devices on the system are working correctly, the most likely reason for this error is that the `SUNWdfb` package was removed or never installed. Insert the installation CD-ROM, change to the `Solaris_2.*` directory, and run the following command to install the packages `SUNWdfbh` and `SUNWdfb` (for your machine architecture):

```
pkgadd -d .
```

If other devices on the system are not working correctly, the system might have a corrupt `/devices` directory. Halt the system and boot using the `-r` (reconfigure) option. The system will run `fsck(1M)` if the `/devices` file system is corrupted, most likely fixing the problem.

insertion failed: a problem with the filesystem has been detected: filesystem is probably full

Cause

With the use of automounter, `ls -l` of an automounted directory is giving the above error. This is a pop-up error message that forces you to press `continue`. However, the `ls -l` does not work properly.

Action

Do a `df -k` to see if the `/var` directory is completely full. Since the `/var/statmon` directory contains the locks for NFS, the automount fails if the `/var` is completely full. After the `/var` directory is reduced to less than 100% of the automount point, `ls -l` should work properly.

Interrupted system call

Cause

The user issued an interrupt signal (usually Control-C) while the system was in the middle of executing a system call. When network service is slow, interrupting `cd(1)` to a remote-mounted directory can produce this message.

Action

Proceed with your work; this message is strictly informational.

Technical Notes

An asynchronous signal (such as `interrupt` or `quit`), which a program was set up to catch, occurred during an internal system call. If execution is resumed after processing the signal, it will appear as if the interrupted programming function returned this error condition, so the program might exit with an incorrect error message.

The symbolic name for this error is `EINTR`, `errno=4`.

Invalid argument

Cause

An invalid parameter was specified that the system cannot interpret. For example, trying to mount an uncreated file system, printing without sufficient system support, or providing an undefined signal to a `signal(3C)` library function can all produce this message.

Action

If you see this message when you are trying to mount a file system, make sure that you have run `newfs(1M)` to create the file system.

If you see this message when you are trying to read a diskette, make sure that the diskette was properly formatted with `fdformat(1)`, either in DOS format, `pcfs(7FS)`, or as a UFS file system.

If you see this message while you are trying to print, make sure that the print service is configured correctly.

Technical Notes

The symbolic name for this error is `EINVAL`, `errno=22`.

Invalid null command

Cause

This C shell message results from a command line with two pipes (`|`) in a row or from a pipe without a command afterwards.

Action

Change the command line so that each pipe is followed by a command.

Invalid_SS_JWS_HOME:no C:\\lib\\basicframe.properties

Cause

The user was running WinNT 4.0 and received this error message when trying to launch Java Workshop™.

Action

Loaded software from marimba company was removed from the user's system. The product was castanet. Afterwards, the JWS worked without problems. Apparently, the product Sun™ Tuner came loaded with JDK™, and this conflicted with JWS.

See www.marimba.com for more details on marimba products.

Another possible solution:

Double-click `jws.exe` within the `C:\Java-WorkShop\jws\intel-win32\bin\` folder.

I/O error

Cause

Some physical Input/Output error has occurred. If the process was writing a file at the time, data corruption is possible.

Action

First, find out which device is experiencing the I/O error. If the device is a tape drive, make sure a tape is inserted into the drive. When this error occurs with a tape in the drive, it is likely that the tape contains an unrecoverable bad spot.

If the device is a floppy drive, an unformatted or defective diskette could be at fault. Format the diskette, or obtain a replacement.

If the device is a hard disk drive, you might need to run `fsck(1M)` and possibly even reformat the disk.

Technical Notes

In some cases this error might occur on a call following the one to which it actually applies.

The symbolic name for this error is `EIO`, `errno=5`.

IP: Hardware address '08:00:20:xx:xx:xx' trying to be our address xxx.xxx.xxx.xxx!

Cause

The above message appears in `/var/adm/messages`.

This can happen, for example, when the ATM lane device is set to promiscuous mode by running `snoop -d lane0`.

Action

Do not let the ATM lane device run in promiscuous mode and do not ignore the warning about it.

Technical Notes

A broadcast over ATM LAN Emulation is emulated by the broadcast and the unknown server (BUS) for the emulated LAN. If the Sun command transmits its ARP request, some switch implementations for LANE repeat the ARP request over the `bus_forward` channel, so that it can be seen on the local interface, again:

```
----- ATM AAL5 Header -----
Packet 1 arrived at 12:12:30.42
Packet size=66 bytes
TRANSMIT : VC=75
LANE Data Frame Type=0x0806 (ARP)
ARP: ----- ARP/RARP Frame -----
ARP:
ARP: Hardware type = 1
ARP: Protocol type = 0800 (IP)
ARP: Length of hardware address = 6 bytes
ARP: Length of protocol address = 4 bytes
ARP: Opcode 1 (ARP Request)
ARP: Sender's hardware address = 8:0:20:82:8f:91
ARP: Sender's protocol address = 192.168.31.54, lab054-lane0
ARP: Target hardware address = ?
```

(continued)

(Continuation)

```
ARP: Target protocol address = 192.168.31.50, lab050-lane0
ARP:

----- ATM AAL5 Header -----
Packet 2 arrived at 12:12:30.42
Packet size=66 bytes
RECEIVE : VC=76
LANE Data Frame Type=0x0806 (ARP)
ARP: ----- ARP/RARP Frame -----
ARP:
ARP: Hardware type = 1
ARP: Protocol type = 0800 (IP)
ARP: Length of hardware address = 6 bytes
ARP: Length of protocol address = 4 bytes
ARP: Opcode 1 (ARP Request)
ARP: Sender's hardware address = 8:0:20:82:8f:91
ARP: Sender's protocol address = 192.168.31.54, lab054-lane0
ARP: Target hardware address = ?
ARP: Target protocol address = 192.168.31.50, lab050-lane0
ARP:
```

Now the request is answered:

```
----- ATM AAL5 Header -----
Packet 3 arrived at 12:12:30.42
Packet size=66 bytes
RECEIVE : VC=84
LANE Data Frame Type=0x0806 (ARP)
ARP: ----- ARP/RARP Frame -----
ARP:
ARP: Hardware type = 1
ARP: Protocol type = 0800 (IP)
ARP: Length of hardware address = 6 bytes
ARP: Length of protocol address = 4 bytes
ARP: Opcode 2 (ARP Reply)
ARP: Sender's hardware address = 8:0:20:8c:4e:f0
ARP: Sender's protocol address = 192.168.31.50, lab050-lane0
ARP: Target hardware address = 8:0:20:82:8f:91
ARP: Target protocol address = 192.168.31.54, lab054-lane0
ARP:
```

Normally, the reflected ARP Request is suppressed. If the lane device is set to promiscuous mode, all packets are passed to upper layers, and so the upper instances receive Sun's own packet and raise this message:

```
Feb 10 12:12:30 sissi unix: IP: Hardware address '08:00:20:82:8f:91'  
trying to be our address 192.168.031.054!
```

Is a directory

Cause

An attempt was made to read or write a directory as if it were a file.

Action

Look at a listing of all the files in the current directory and try again, specifying a file instead of a directory.

Technical Notes

The symbolic name for this error is EISDIR, errno=21.

"J"

java.lang.UnsatisfiedLinkError:

Cause

When trying to start Java Workshop 2.0 (or some other Java applications), the following error is displayed:

```
java.lang.UnsatisfiedLinkError: setCursor  
    at sun.awt.motif.MComponentPeer.initialize(Compiled Code)  
    at sun.awt.motif.MTextAreaPeer.initialize(Compiled Code)  
    at sun.awt.motif.MComponentPeer.<init>(Compiled Code)  
    at sun.awt.motif.MTextAreaPeer.<init>(Compiled Code)  
    at sun.awt.motif.MToolkit.createTextArea(Compiled Code)
```

Action

The `LD_LIBRARY_PATH` is probably set up to include a Java lib directory that does not quite match the java bin command used. For example, in the Solaris 2.6 release `LD_LIBRARY_PATH = /usr/openwin/lib` results in Java Workshop running properly. But setting `LD_LIBRARY_PATH = /usr/java/lib:/usr/openwin/lib` results in the error being displayed, since Java Workshop uses its own version of JDK and the startup process picks up a mixture of versions.

To resolve, include `/usr/java/lib` in your `LD_LIBRARY_PATH`, since it is needed only in rare circumstances (like when you are using the Java Invocation API).

"K"

kernel read error

Cause

This message appears when `savecore(1M)`, if activated, tries to copy a debugging image of kernel memory to disk, but cannot read various kernel data structures correctly. Generally, this occurs after a system panic has corrupted the main memory. Data corruption on the system is possible.

Action

Look at the kernel error messages that preceded this one to try to determine the cause of the problem. Error messages such as `BAD TRAP` usually indicate faulty hardware. Until the problem that caused the kernel panic is resolved, a kernel core image cannot be saved for debugging.

killed

Cause

A process, which attempts to allocate large amounts of memory either as an array or by using `malloc`, fails when launched by the shell. This problem has been seen while allocating 240,000,000 elements as either an array of doubles or using `malloc` to allocate the 1,920,000,000 bytes of space.

Action

This can have one of two causes. Resolve it accordingly.

1. Lack of swap space

Try running the program as root on the console; if it runs, this is not the problem.

2. Stack size and data segment size are in conflict

If the stack size is set too large, this can conflict with the data segment, and the process cannot be started. Setting the stack size to the default value of 8192 resolves this problem and allows the programs to start.

Killed

Cause

This message is strictly informational. If the killed process was writing a file, some data might be lost.

Action

Continue with your work.

Technical Notes

This message from the signal handler or various shells indicates that a process has been terminated with a `SIGKILL`. However, if you do not see this message and cannot terminate a process with a `SIGKILL`, you might have to reboot the machine to remove that process.

kmem_free block already free

Cause

This is a programming error, probably from a device driver.

Action

Determine which driver is giving this message and contact the vendor for a software update, as this message indicates a bug in the driver.

Technical Notes

This message is from the DDI programming function `kmem_free(9F)`, which releases a block of memory at address `addr` of size `siz` that was previously allocated by the DDI function `kmem_alloc(9F)`. Both `addr` and `siz` must correspond to the original allocation. If you have source code for the driver, follow `kmem_alloc(9F)` and `kmem_free(9F)` in the code to make sure they allocate and free the same chunk of memory.

"L"

last message repeated *int* times

Cause

This message comes from `syslogd(1M)`, the facility that prints messages on the console and records them in `/var/adm/messages`. To reduce the log size and minimize buffer usage, `syslog` collapses any identical messages it sees during a 20 second period, then prints this message with the number of repetitions.

Action

Look above this message to see which message was repeated so often. Then consider the repeated message and take action accordingly. If repeated log entries such as `su ... failed` appear, consider the possibility of a security breach.

late initialization error

Cause

Netscape enterprise server 2.0 receives these error messages from the daemon:

- `late initialization error`

- start up failure no such file or directory
- system will not connect to port 80

Action

This is a file permission problem caused by someone changing the UID for the `httpd` user in `/etc/passwd`.

Change UID in `/etc/passwd` to the correct UID.

ld.so.1 fatal: can't set protection on segment

Cause

Applications have recently begun to fail with this error, `ld.so.1 fatal: can't set protection on segment`. The failures are random.

Action

This was happening because of the recent introduction of a rogue application that consumed most of the swap space on the system. The other applications, which failed randomly, were doing so because of having insufficient swap space to run. The error from `ld.so.1` occurred because there was no segment on which to set the protections.

ld.so.1: *string*: fatal: *string*: can't open file: errno=2

Cause

This message is produced in releases earlier than Solaris 2.5.1. It is not produced in releases after Solaris 2.5.1.

For more information about the cause, refer to “`ld.so.1: string: fatal: string: open failed: No such file or directory`” on page 145. It has the same cause.

Action

For the resolution, refer to “ld.so.1: *string*: fatal: *string*: open failed: No such file or directory” on page 145. Their resolutions are the same.

See Also

For more information about the Linker, see the *Linker and Libraries Guide*.

ld.so.1: *string*: fatal: *string*: open failed: No such file or directory

Cause

This message is produced in releases after Solaris 2.5.1. It is not produced in releases before Solaris 2.5.1.

This message indicates that the runtime linker, `ld.so.1(1)`, while running the program specified after the first colon, could not find the shared object specified after the third colon. (A shared object is sometimes called a dynamically linked library.)

Action

As a workaround, set the environment variable `LD_LIBRARY_PATH` to include the location of the shared object in question. For example:

```
/usr/dt/lib:/usr/openwin/lib
```

Better yet, if you have access to source code, recompile the program using the `-Rpath` loader option. Using `LD_LIBRARY_PATH` slows down performance.

See Also

For more information about the Linker, see the *Linker and Libraries Guide*.

ld.so.1: *string*: fatal: relocation error: *string*: *string*: referenced symbol not found

Cause

This message is produced in releases after the Solaris 2.5.1. It is not produced in the Solaris 2.5.1 or earlier releases.

The message from the runtime linker `ld.so.1(1)` indicates that in trying to execute the application given after the first colon, the specified symbol could not be found for relocation. The message goes on to say in what file the symbol was referenced. Because this is a fatal error, the application terminates with this message.

Action

Run the `ldd -d` command on the application to show its shared object dependencies and symbols that are not found. Probably your system contains an old version of the shared object that should contain this symbol. Contact the library vendor or author for an update.

Technical Notes

This error does not necessarily occur when you first bring up an application. It could take months to develop, if ordinary use of the application seldom references the undefined symbol.

See Also

For more information about the Linker, see the *Linker and Libraries Guide*.

ld.so.1: *string*: fatal: relocation error: symbol not found: *string*

Cause

This message is produced in the Solaris 2.5.1 release and earlier. It is not produced in releases after the Solaris 2.5.1.

Refer to “ld.so.1: *string*: fatal: relocation error: *string*: *string*: referenced symbol not found” on page 146. It has the same cause.

Action

For a resolution, refer to “ld.so.1: *string*: fatal: relocation error: *string*: *string*: referenced symbol not found” on page 146. Their resolutions are the same.

Technical Notes

This error does not necessarily occur when you first bring up an application. It could take months to develop, if ordinary use of the application seldom references the undefined symbol.

See Also

For more information about the Linker, see the *Linker and Libraries Guide*.

le0: Memory error!

Cause

This message indicates that the network interface encountered an access time-out from the CPU’s main memory. There is probably nothing wrong except system overload.

Action

If the system is busy with other processes, this error can occur frequently. If possible, try to reduce the system load by quitting applications or killing some processes.

Technical Notes

The Lance Ethernet chip timed out while trying to acquire the bus for a DVMA transfer. Most network applications wait for a transfer to occur, so generally no data gets lost. However, data transfer might fail after too many time-outs.

See Also

For more information about the Lance Ethernet chip, see the `le(7D)` man page.

le0: No carrier– cable disconnected or hub link test disabled?

Cause

Stand-alone machines with no Ethernet port connection get this error when the system tries to access the network. If the Ethernet cable is connected, this message could result from a mismatch between the machine's NVRAM settings and the Ethernet hub settings.

Action

If this message is continuous, try to save any work to a local disk.

When a machine is configured as a networked system, it must be plugged into the Ethernet with a twisted pair J45 connector.

If the Ethernet cable is plugged in, find out whether or not the Ethernet hub does a link integrity test. Then become superuser to check and possibly set the machine's NVRAM. If the hub's link integrity test is disabled, set this variable to `false`.

```
# eeprom | grep tpe
tpe-link-test?=true
# eeprom 'tpe-link-test?=false'
```

The default setting is `true`. If for some reason `tpe-link-test?` was set to `false`, and the hub's link integrity test is enabled, reset this variable to `true`.

le0: No carrier– transceiver cable problem?

Cause

Stand-alone machines with no Ethernet port connection get this error when the system tries to access the network.

Action

If this message is continuous, try to save any work to a local disk.

When a machine is configured as a networked system, it must be plugged into the Ethernet with either a twisted pair J45 connector or thicknet 10Base-T connector (depending on the building's Ethernet cable type).

Technical Notes

Older workstations have a thicknet connection on the back, instead of a twisted pair Ethernet connection; therefore, they require a thicknet to the twisted pair transceiver to translate between cabling types.

level 15 interrupt

Cause

This error occurred on an SS20.

.lib section in a.out corrupted

Cause

This occurred while trying to `exec(2)` an `a.out(4)`, which requires that a static shared library be linked in. Also, there was erroneous data in the `.lib` section of the `a.out(4)`. The `.lib` section tells `exec(2)` which static shared libraries are needed. The `a.out(4)` is probably corrupted.

Technical Notes

The symbolic name for this error is `ELIBSCN, errno=85`.

LINK COUNT FILE I=i OWNER=o MODE=m SIZE=s MTIME=t COUNT... ADJUST?

Cause

During phase 4, `fsck(1M)` determined that the inode's link count for the specified file is wrong and asks if you want to adjust it to the value given.

Action

Generally you can answer "yes" to this question without harming the file system.

See Also

For more information on `fsck(1M)`, see the section on checking file system integrity in the *System Administration Guide, Volume 1*.

Link has been severed

Cause

This error occurs when the connection to a remote machine is gone, for example after a remote procedure call is interrupted.

Technical Notes

The symbolic name for this error is `ENOLINK`, `errno=67`.

LL105W: Protocol error detected.

Cause

This error message comes from Lifeline™ Mail, an unbundled PC compatibility application.

Most likely, someone set up a user account without a password.

Action

To solve this problem, assign the user a password.

In: cannot create /dev/fb: Read-only filesystem

Cause

During device reconfiguration at boot time, the system cannot link to the frame buffer because `/dev` is on a read-only file system.

Action

Check that `/dev/fb` is a symbolic link to the hardware frame buffer, such as `cgsix(7D)` or `tcx(7D)`. Ensure that the file system containing `/dev` is mounted read-write.

lockd[*int*]: create_client: no name for inet address *0xhex*

Cause

This lock daemon message usually indicates that the NIS `hosts.byname` and `hosts.byaddr` maps are not coordinated.

Action

Wait a short time for the maps to synchronize. If they do not, take steps to coordinate them.

See Also

For information on updating NIS data, see the section on NIS maps in the *NIS+ and FNS Administration Guide*. If you are using AnswerBook online documentation, "hosts.byaddr" is a good search string.

log_get: len is not a multiple of 4 from FW-1

Cause

The Firewall-1™ log contains this message. It is logged when one of the log files is somehow damaged, usually after a power outage or violent reboot of the system.

Action

Try the following workaround:

```
# fwstop
# rename fw.log, fw.alog, fw.vlog
# fwstart
```

Login incorrect

Cause

This message from the `login(1)` program indicates an incorrect combination of login name and password. There is no way to tell whether the problem comes from the login name, the password, or both. Other programs such as `ftp(1)`, `rexecd(1M)`, `sulogin(1M)`, and `uucp(1C)` also give this error under similar conditions.

Action

Check the `/etc/passwd` file and the NIS or NIS+ `passwd` map on the local system to see if an entry exists for this user. If a user has simply forgotten the password, `su(1M)` and set a new one with the `passwd(1)` *username* command. This command automatically updates the NIS+ `passwd` map, but with NIS you will need to coordinate the update with the `passwd` map.

The `Login incorrect` problem can also occur with older versions of NIS when the user name has more than eight characters. If this is the case, edit the NIS password file, change the user name to have eight or fewer characters, and then remake the NIS `passwd` map.

If you cannot log in to the system as root, despite knowing the proper password, it is possible that the `/etc/passwd` file is corrupted. Try to log in as a regular user and `su(1M)` to root.

If that does not work, see the message `su: No shell` and follow most of the instructions given there. Instead of changing the default shell, make the password field blank in `/etc/shadow`.

lp hang

Cause

On a print server, the queue continues to grow but nothing comes out of the printer. The printer daemon is hung.

Action

Below is a simple procedure for flushing a hung printing queue:

1. Login or switch user to `root`.
2. Issue the `reject(1M) printername` command to make sure no one sends any job to the printer.
3. Turn the power off to the printer.
4. If the active job appears to be causing the hang, remove it from the print queue with the `cancel(1) jobnumber` command and ask the owner to requeue that print job.
5. Shut down the print queue with the `/usr/lib/lpshut` command.
6. Remove the lock file `/var/spool/lp/SCHEDLOCK` and the temporary files `/var/spool/lp/tmp/*/*`.
7. Turn the printer back on.
8. Restart the print queue with the `/usr/lib/lpsched` command.

See Also

For more information on print queuing, see the *System Administration Guide, Volume 2*. If you are using AnswerBook online documentation, "print server" is a good search string.

"M"

Machine is not on the network

Cause

This error is remote file sharing (RFS) specific. It occurs when users try to advertise, unadvertise, mount, or unmount remote resources while the machine has not properly started a network connect.

Technical Notes

The symbolic name for this error is `ENONET`, `errno=64`.

Mail Tool is confused about the state of your Mail File.

Cause

This message appears in a pop-up dialog box whenever you ask *mailtool(1)* to access messages after another mail reader has modified your inbox. A request follows:
`Please Quit this Mail Tool.`

Action

Click continue to close the dialog box, then exit *mailtool(1)*. If you continue trying to read mail, messages deleted by the other mail reader will never appear, and *mailtool(1)* will fail to see any new messages.

mail: Your mailfile was found to be corrupted (Content-length mismatch).

Cause

This message comes from `mail(1)` or `mailx(1)` whenever it detects messages with a different content length than advertised. The `mail(1)` program tells you which message might be truncated or might have another message concatenated to it.

Two common causes of content length mismatches are the simultaneous use of different mail readers (such as `mail(1)` and `mailtool(1)`), or the use of a mail reading program (or an editor) that does not update the content-length field after altering a message.

Action

The `mailx(1)` program can usually recover from this error and delineate mail message boundaries correctly. Pay close attention to the message that might be truncated or combined with another message, and to all messages after that one. If a mail file becomes hopelessly corrupted, run it through a text editor to eliminate all Content-Length lines, and ensure that each message has a `From` (no colon) line for each message, preceded by a blank line.

To avoid mail file corruption, exit from `mailtool(1)` without saving changes when you are currently running `mail(1)` or `mailx(1)`.

mailtool: Can't create dead letter: Permission denied

Cause

An attempt was made to send a message with `mailtool(1)` from a directory where the user does not have write permission, and the user's home directory is currently unavailable.

Action

Change to another directory and start `mailtool(1)` again, or use `chmod(1)` to change permissions for the directory (if possible).

mailtool: Could not initialize the Classing Engine

Cause

When a user runs *mailtool(1)* on a remote machine, setting the DISPLAY environment back to the local machine, this message might appear inside a dialog box window. The message also indicates that the Classing Engine must be installed to use Attachments. This problem occurs because *rlogin(1)* does not propagate the user's environment.

Action

Exit *mailtool(1)* and set your OPENWINHOME environment variable to `/usr/openwin`. Then run *mailtool(1)* again. The error message does not appear, and you can now use Attachments.

Technical Notes

Classing Engine is a new name for Tool Talk. Earlier versions of *mailtool(1)* said Tool Talk: TT_ERR_NOMP instead of Classing Engine.

Management Server is VPN while client is NON-VPN

Cause

When the Windows GUI (*fwpolicy*) is started in Firewall-1 3.0 and the login process is initiated, the error message window pops up displaying this message.

Action

The Firewall-1 GUI packages *SUNWfwgui* and *SUNWfweui* were installed in the incorrect order. First, remove the packages using *pkgrm(1M)*. Next, install the *SUNWfwgui* and, then, the *SUNWfweui* *in that order* to resolve the error message.

file name may contain holes - can't swap on it.

Cause

A swap file was created with the following command:

```
# mkfile -nv 50m /ab/swap_50mb
```

When the user tried to add the file with

```
# swap -a /ab/swap_50mb
```

it failed with this message:

```
/ab/swap_50mb may contain holes - can't swap on it.  
/ab/swap_50mb: Error 0
```

Action

Starting with the Solaris 2.0 release, `-n` works only when the file is to be used by the NFS system. Local swap files cannot be created with the `-n` option.

mbuf map full

Cause

This error has to do with `mbuf` allocation.

Memory address alignment

Cause

This message can occur when printing large files on a SPARCprinter™ attached to a SPARCstation 2.

Action

Replace the SPARCstation 2 CPU with one that is at the most recent dash level.

memory leaks

Cause

An application uses up more and more memory, until all swap space is exhausted.

Action

Third-party software can help identify memory leaks in their applications. If you suspect that you have a memory leak, you can use `sar(1)` to check on the Kernel Memory Allocation (KMA). Any driver or module that uses KMA resources, but does not specifically return the resources before it exits, can create a memory leak.

See Also

For more information on memory leaks, see the section on monitoring system activity in the *System Administration Guide, Volume 2*. If you are using AnswerBook online documentation, "displaying disk usage" is a good search string. Also, see the section on system resource problems in the *NIS+ and FNS Administration Guide*.

Message too long

Cause

A message sent on a transport provider was larger than the internal message buffer or some other network limit.

Technical Notes

The symbolic name for this error is `EMSGSIZE`, `errno=97`.

mount: /dev/dsk/*string* is already mounted, /*string* is busy, or...

Cause

While trying to mount a file system, the `mount(1M)` command received a "Device busy" (`EBUSY`) error code. Several possible reasons are: this `/dev/dsk` file system is already mounted on a different directory, the busy path name is the working directory of an active process, or the system has exceeded its maximum number of mount points (unlikely).

Action

Run `/etc/mount` to see if the file system is already mounted. If not, check to see if any shells are active in the busy directory (did the user switch to the directory by using `cd(1)?`), or if any processes in the `ps(1)` listing are active in that directory. If the reason for the error message is not obvious, try using a different directory for the mount point.

mount: giving up on: /*string*

Cause

An existing server did not respond to an NFS mount request, so after retrying a number of times (default 1000), the `mount(1M)` command has ceased. Nonexistent servers or bad mount points produce different messages.

Action

If the `RPC: Program not registered` message precedes this one, the requested mount server probably did not share (export) any file systems, so it has no NFS daemons running. Have the superuser on the mount server run `share(1M)` on the file system, then run `/etc/init.d/nfs.server start` to begin NFS service.

If the requested mount server is down or slow to respond, check whether the machine needs repair or rebooting.

mount: mount-point /*string* does not exist.

Cause

Someone tried to mount a file system onto the specified directory, but there is no such directory.

Action

If this is the directory name you want, run `mkdir(1)` to create this directory as a mount point.

mount: the state of /dev/dsk/*string* is not okay

Cause

The system was unable to mount the file system that was specified because the super block indicates that the file system might be corrupted. This is not an impediment for read-only mounts.

Action

If you do not need to write on this file system, run `mount(1M)` on it using the `-o ro` option. Otherwise, do as one of the message continuation lines suggests and run `fsck(1M)` to correct the file system state and update the super-block.

See Also

For more information on using `fsck(1M)`, see the section on checking file system integrity in the *System Administration Guide, Volume 1*.

Multihop attempted

Cause

This error occurs when users try to access remote resources that are not directly accessible.

Technical Notes

The symbolic name for this error is EMULTIHOP, errno=74.

"N"

Name not unique on network

Cause

The given log name is not unique.

Technical Notes

The symbolic name for this error is ENOTUNIQ, errno=80.

named [pid]: *hostname.domainname* has CNAME and other data (illegal)

Cause

This error message is displayed on the DNS server.

Action

This error indicates that an alias (CNAME) is associated with another type of DNS record.

The DNS system allows you to set up an alias to a system using the CNAME record. See the following example:

```
alias1 IN CNAME host1.domain1.
```

The alias `alias1` cannot appear in any other type of record. Only the actual name of the host can be used. So, if you wanted to use this host as a mail exchanger, the record

```
alias1 IN MX 10 host2.domain1.
```

would be illegal and would produce the error.

Instead, you should use:

```
host1 IN MX 10 host2.domain1.
```

This remedy applies to all types of records, including HINFO and A records.

Also, this error might occur without explicitly setting the left side of a record. The DNS system defaults the left side to the last given left side. So you might have the following in a named database file:

```
host1 IN A 123.124.125.126
      IN HINFO Sun Solaris
alias1 IN CNAME host1.domain1.
      IN MX 10 host2.domain1.
```

In this fragment, an implied `alias1` is in the left side of the MX record. If the alias was added after the database was in use for a while, the error would suddenly occur. The MX record was legal until the CNAME was added in front of it. This example could be fixed either by reversing the order of the MX and CNAME records, or explicitly giving the `host1` in the left side of the MX record.

/net/string: No such file or directory

Cause

A user tried to change directory—for example with `cd(1)`—to a network partition on the system specified after `/net/`, but this host either does not exist or has not shared (exported) any file system.

Action

To gain access to files on this system, try `rlogin(1)`.

To export file systems from the remote system, become superuser on that system and run the `share(1M)` command with the appropriate options. If that system is sharing file systems for the first time, also run `/etc/init.d/nfs.server start` to begin the NFS service.

Network dropped connection because of reset

Cause

The host you were connected to crashed and rebooted.

Technical Notes

The symbolic name for this error is `ENETRESET`, `errno=129`.

Network is down

Cause

A transport connection failed because it encountered a dead network.

Action

Report this error to the system administrator for the network. If you are the person responsible for this network, check why the network is dead and what repairs are necessary.

Technical Notes

This error results from status information delivered by the underlying communication interface.

The symbolic name for this error is `ENETDOWN`, `errno=127`.

Network is unreachable

Cause

An operational error occurred either because there was no route to the network or because negative status information was returned by intermediate gateways or switching nodes.

The returned status is not always sufficient to distinguish between a network that is down and a host that is down. See the `No route to host` message.

Action

Check the network routers and switches to see if they are disallowing these packet transfers. If they are allowing all packet transfers, check network cabling and connections.

Technical Notes

The symbolic name for this error is `ENETUNREACH`, `errno=128`.

NFS getattr failed for server *string*: RPC: Timed out

Cause

This message appears on an NFS client that requested a service from an NFS server that has failing hardware. Often the message `NFS read failed` appears along with this message. If the server were merely down or slow to respond, the `NFS server not responding` message would appear instead. Data corruption on the server system is possible.

Action

Because this message usually indicates server hardware failure, initiate repair procedures as soon as possible. Check the memory modules, disk controllers, and CPU board.

See Also

For more information on NFS tuning, see the chapter on monitoring network performance in the *System Administration Guide, Volume 2*.

nfs mount: Couldn't bind to reserved port

Cause

This message appears when a client attempts to use NFS to mount a file system from a server that has more than one Ethernet interface configured on the same physical subnet.

Action

Always connect multiple Ethernet interfaces on one router system to different physical subnetworks.

nfs mount: mount: *string*: Device busy

Cause

This message appears when the superuser attempts to NFS mount on top of an active directory. The busy device is actually the working directory of a process.

Action

Determine which shell on the workstation is currently located below the mount point, and change that directory. Be wary of subshells (such as `su(1M)` shells) that could be in different working directories while the parents remain below the mount point.

NFS mount: /*string* mounted OK

Cause

While booting, the system failed to mount the directory specified after the first colon, probably because the NFS server involved was down or slow to respond. The mount ran in the background and successfully contacted the NFS server.

Action

This is strictly an informative message to notify you that the mount process has completed.

NFS mounted callog file Unsupported.

Cause

After installing the Solaris 2.6 software on a system, when users try to start their calendars either with CDE's calendar manager (`/usr/dt/bin/dtcm`) or OpenWindows' calendar (`/usr/openwin/bin/cm`), they see this dialog box:

```
Calendar :Informational - NFS mounted callog file Unsupported.
Your default startup Calendar file appears to be NFS mounted or
a symlink to the same. This is Not Supported.
Continue
```

The following error is displayed in the console window when the Continue button is clicked:

```
date time host rpc.cmsd[pid]: rpc.cmsd :
NFS mounted callog file Not Supported - user@host
date time host rpc.cmsd[pid]: rpc.cmsd :
NFS mounted callog file Not Supported - user@host
```

The calendars would have worked under the Solaris software versions including and prior to 2.5.1, however.

Action

It has long been known that NFS-mounted calendars are not supported. The calendar can be corrupted when more than one person uses the calendar at the same time. If two `rpc.cmsd` daemons write to the `callog` file at the same time, the file becomes corrupt. However, two `rpc.cmsd` daemons could be run simultaneously on the Solaris 2.5.1 release, even though this is not a supported configuration.

With the Solaris 2.6 release, this concurrency is no longer an option. `rpc.cmsd` does not allow the user to start a calendar that is NFS-mounted and produces the previous error message.

NFS read failed for server *string*

Cause

This message generally indicates a permissions problem. Perhaps a directory or file permission was changed while the client kept the file open. Perhaps the file system's share or netgroup permissions changed. If the server were down or the network overloaded, the `NFS server not responding` message would appear instead.

Action

Log in to the NFS server and check the permissions of directories leading to the file. Make certain that the file system is shared with (exported to) the client experiencing an NFS read failure.

See Also

For more information, see the chapter on NFS troubleshooting in the *System Administration Guide, Volume 3*.

nfs_server: bad getargs for *int/int*

Cause

This message comes from the NFS server when it receives a request with unrecognized or incorrect arguments. Typically, it means the request could not be XDR decoded properly. This error can result from corruption of the packet over the network, or from an implementation bug causing the NFS client to encode its arguments improperly.

Action

If this message originates from a single client, investigate that machine for NFS client software bugs. If this message appears throughout a network, especially accompanied by other networking errors, investigate the network cabling and connectors.

NFS server *string* not responding still trying

Cause

In most cases this common message indicates that the system has requested a service from an NFS server that is either down or extremely slow to respond. In some cases this message indicates that the network link to this NFS server is broken, although usually that condition generates other error messages as well. In a few cases this message indicates NFS client setup problems.

Action

Check the non-responding NFS server for the need for machine repair or rebooting. Encourage your user community to report such problems quickly but only once.

See Also

For more information, see the chapter on NFS troubleshooting in the *System Administration Guide, Volume 3*.

NFS server *string* ok

Cause

This message is the follow-up to the `NFS server not responding` error. It indicates that the NFS server is again operating.

Action

When an NFS server first starts, it is busy fulfilling client requests for a while. Be patient and wait for your client system to respond. Making many extraneous requests only further slows the NFS server response time.

NFS *string* failed for server *string*: error *int* (*string*)

Cause

The failed NFS operation could be any one of the following: `getattr`, `setattr`, `lookup`, `access`, `readlink`, `read`, `write`, `create`, `mkdir`, `symlink`, `mknod`, `remove`, `rmdir`, `rename`, `link`, `readdir`, `readdir+`, `fsstat`, `fsinfo`, `pathconf`, or `commit`.

See Also

For more information on NFS, see the *System Administration Guide, Volume 3*.

nfs umount: *string*: is busy

Cause

This message appears when the superuser attempts to unmount an active NFS file system. The busy point is the working directory of a process.

Action

Determine which shell (or process) on the workstation is currently located in the remotely mounted file system, and change— `cd(1)`—out of that directory. Be wary of subshells (such as `su(1M)` shells) that could be in different directories while the parent shells remain in the NFS file system.

NFS write error on host *string*: No space left on device.

Cause

This console message indicates that an NFS-mounted partition has filled up and cannot accept writing of new data. Unfortunately, software that attempts to overwrite existing files will usually zero-out all data in these files. This is particularly destructive on NFS-mounted `/home` partitions.

Action

Find the user or process that is filling up the file system, and stop the out-of-control process as soon as you can. Then delete files as necessary to create more space on the file system (large `core(4)` files are good candidates for deletion). Have users write any modified files to local disk if possible. If this error occurs often, redistribute directories to ease the demand on this partition.

See Also

For more information on disk usage, see the *System Administration Guide, Volume 2*. If you are using AnswerBook online documentation, "managing disk use" is a good search string.

NFS write failed for server *string*: RPC: Timed out

Cause

This error can occur when a file system is soft mounted, and server or network response time lags. Any data written to the server during this period could be corrupted.

Action

If you intend to write on a file system, never specify the soft-mount option. Use the default hard mount for all the file systems that are mounted read-write.

See Also

For more information, see the chapter on NFS troubleshooting in the *System Administration Guide, Volume 3*.

NIS+ authentication failure

Cause

This is a Federated Naming Service message. The operation could not be completed because the principal making the request could not be authenticated with the name service involved.

Action

Run the `nisdefaults(1)` command to verify that you are identified as the correct NIS+ principal. Also check that the system has specified the correct public key source.

See Also

For more information, see the authentication and authorization overview in the *NIS+ and FNS Administration Guide*.

nis_cachemgr: Error in reading NIS cold start file : '/var/nis/NIS_COLD_START'

Cause

After installing patches 104331-04 and 103612-33, `nis_cachemgr(1M)` failed to start. The symptoms are as follows during the reboot:

```
Sep 11 16:34:00 nis_cachemgr: Error in reading NIS cold start file :  
'/var/nis/NIS_COLD_START'
```

Additionally, `nis_cachemgr(1M)` is not running after login. Trussing `nis_cachemgr(1M)` showed that it is reading `/var/nis/NIS_COLD_START` and immediately reporting an error. Neither reinitializing the client nor copying `NIS_COLD_START` helps.

Action

This error is a timing problem. Put a `sleep(1)` before the NIS+ initialization in `/etc/init.d/rpc`, after `rpc.bind` has been started. `rpc.bind` is slow initializing and needs a few extra seconds before `nis_cachemgr(1M)` takes effect.

No buffer space available

Cause

An operation on a transport endpoint or pipe was not performed because the system lacked sufficient buffer space or because a queue was full. The target system

probably ran out of memory or swap space. Any data written during this condition is probably lost.

Action

To add more swap area, use the `swap -a` command on the target system. Alternatively, reconfigure the target system to have more swap space. As a general rule, swap space should be two to three times as large as physical memory.

Technical Notes

The symbolic name for this error is `ENOBUFFS`, `errno=132`.

No child processes

Cause

This message can appear when an application tries to communicate with a cooperating process that does not exist.

Action

Restart the parent process so it can create the child processes again. If that does not help, this error could be the result of a programming error; contact the vendor or author of the program for an update.

Technical Notes

A `wait(2)` system call was executed by a process that had no existing or unwaited-for child processes. The child processes could have exited prematurely, or might never have been created.

The symbolic name for this error is `ECHILD`, `errno=10`.

No default media available

Cause

The volume manager issues this message if a user makes an `eject(1)` request when the drives contain no diskette or CD-ROM to eject.

Action

Insert a diskette or CD-ROM. If the volume manager is confused and a diskette or CD-ROM is actually in a drive, run `volcheck(1)` to update the volume manager. If the system remains confused, try booting with the `-r` option to reconfigure devices.

No directory! Logging in with home=/

Cause

The `login(1)` program could not find the home directory listed in the password file or NIS `passwd(4)` map, so it deposited the user in the root directory.

Action

Check that the user's home directory is mounted and is owned by and accessible to that user. Perhaps the automounter tried to mount the home directory, but the NFS server did not respond quickly enough. Try listing the files in `/home/username`. If the NFS server responds to this request, have the user log out and log in again.

The automounter daemon might not be running. Run the `ps(1)` command to see if `automountd(1M)` is present. If not, run the second command; if it appears to be wedged, run both these commands:

```
# /etc/init.d/autofs stop
# /etc/init.d/autofs start
```

When the automounter daemon is running, verify that the `/etc/auto_master` file has a line like this:

```
/home auto_home
```

Verify that the `/etc/auto_home` file has a line like this:

```
+auto_home
```

These entries depend on the NIS `auto_home` map.

Also, the NFS server might not have shared (exported) this `/home` directory, or the NFS daemons on the server might have disappeared.

See Also

For more information on NFS, see the *System Administration Guide, Volume 3*.

No message of desired type

Cause

An attempt was made to receive a message of a type that does not exist on the specified message queue. See the `msgsnd(2)` and `msgrcv(2)` man pages for details.

Action

This message indicates an error in the System V IPC message facility. Generally the message queue is empty or devoid of the desired message type while `IPC_NOWAIT` is set.

Technical Notes

The symbolic name for this error is `ENOMSG`, `errno=35`.

No recipients specified

Cause

This message comes from the `mailx(1)` command whenever a user does not provide an address in the `To:` field.

Action

For details, refer to “Recipient names must be specified” on page 208.

No record locks available

Cause

No more record locks are available. The system lock table is full.

Action

Try again later, when more locks might be available.

Technical Notes

The symbolic name for this error is `ENOLCK`, `errno=46`.

Perhaps a process called `fcntl(2)` with the `F_SETLK` or `F_SETLKW` option, and the system maximum was exceeded. The system contains several different locking subsystems, including `fcntl(2)`, the NFS lock daemon, and mail locking. All subsystems can produce this error.

No route to host

Cause

An operational error occurred because there was no route to the destination host, or because of status information returned by intermediate gateways or switching nodes.

The returned status is not always sufficient to distinguish between a host that is down and a network that is down. Refer to “Network is unreachable” on page 163.

Action

Check that the network routers and switches are not disallowing these packet transfers. If they are allowing all packet transfers, check network cabling and connections.

Technical Notes

The symbolic name for this error is `EHOSTUNREACH`, `errno=148`.

No shell Connection closed

Cause

A user has attempted a remote login to the system, and has a valid account name and password, but the shell specified for the account is not available on that system.

Action

If you have a copy of the requested shell, become superuser and install the missing shell on that system. Otherwise, change the user's password file entry—perhaps only in the NIS+ or NIS `passwd(4)` map—to specify an available shell such as `/bin/csh` or `/bin/ksh`.

No space left on device

Cause

While writing an ordinary file or creating a directory entry, there was no free space left on the device. The disk, tape, or diskette is full of data. Any data written to that device during this condition can be lost.

Action

Remove unneeded files from the hard disk or diskette until there is space for all the data you are writing. You also might move some directories onto another file system and create symbolic links accordingly. When a tape is full, continue on another one, use a higher-density setting, or obtain a higher-capacity tape.

To create multi-volume tapes or diskettes, use the `pax(1)` or `cpio(1)` command; `tar(1)` is still limited to a single volume.

Technical Notes

The symbolic name for this error is `ENOSPC`, `errno=28`.

No such device

Cause

An attempt was made to apply an operation to an inappropriate device, such as writing to a nonexistent device.

Action

Check the `/devices` directory to find out why this device does not exist, or why the program expects it to exist. The similar `No such device or address` message tends to indicate I/O problems with an existing device, whereas this message tends to indicate a device that does not exist at all.

Technical Notes

The symbolic name for this error is `ENODEV`, `errno=19`.

No such device or address

Cause

This error can occur when a tape drive is offline or when a device has been powered off or removed from the system.

Action

For tape drives, make sure the device is connected, powered on, and toggled online (if applicable). For disk and CD-ROM drives, check that the device is connected and powered on.

With all SCSI devices, ensure that the target switch or dial is set to the number where the system originally mounted it. To inform the system of a change to the target device number, reboot using the `-r` (reconfigure) option.

Technical Notes

This message results from I/O to a special file's subdevice that either does not exist or that exists beyond the limit of the device.

The symbolic name for this error is `ENXIO`, `errno=6`.

No such file or directory

Cause

The specified file or directory does not exist. Either the file name or path name was entered incorrectly.

Action

Check the file name and path name for correctness and try again. If the specified file or directory is a symbolic link, it probably points to a nonexistent file or directory.

Technical Notes

The symbolic name for this error is `ENOENT`, `errno=2`.

no such map in server's domain

Cause

A user or an application tried to look up something using Network Information Services (NIS), but NIS has no corresponding database for this request.

Action

Check the following:

- Make sure the NIS map name is spelled correctly. To see a list of nicknames for the various NIS maps, run the `ypcat -x` command.
- To see a full list of the various NIS maps (databases), run the `ypwhich -m` command.
- If the NIS service was not running on the current machine, these commands would result in this message: “can’t communicate with ypbind” on page 65.

No such process

Cause

This process cannot be found. The process could have finished execution and disappeared, or it might still be in the system under a different numeric ID.

Action

Use the `ps(1)` command to check that the process ID you are supplying is correct.

Technical Notes

No process corresponds to the specified process ID (PID), lightweight process ID, or `thread_t`.

The symbolic name for this error is `ESRCH`, `errno=3`.

No such user as *string*— cron entries not created

Cause

A file exists in `/var/spool/cron/crontabs` for the specified user, but this user is not in `/etc/passwd` or the NIS `passwd(4)` map. The system cannot create `cron(1M)` entries for nonexistent users.

Action

To eliminate this message at boot time, remove the `cron` file for the nonexistent user, or rename it if the user's login name has changed. If this is a valid user, create an appropriate password entry for this name.

No utmpx entry

Cause

During login, file system full errors are seen and the login fails with the message `No utmpx entry`.

This error is caused by a full file system. The system has no space to write its `utmpx` (login information) entry.

Action

To correct this condition the system must be booted into single user mode. Then clear (do not delete) these files: `/var/adm/utmp` and `/var/adm/utmpx`. This can be done by typing:

```
#cat /dev/null > /var/adm/utmp
#cat /dev/null > /var/adm/utmpx
```

These commands zero-out the files but keep them with the correct permissions.

In some cases, after clearing these files, the `/var` file system might still be full. In this case type:

```
du -askd /var |sort -nr |more
```

This command gives you a listing of the files from largest to smallest in the `/var` file system. To create space you can zero these files: `/var/cron/log`, `/var/spool/lp/logs`, and `/var/adm/messages`. You can also check `/.wastebasket` for large files to delete.

no valid fm license

Cause

The firewall gives you this error when the proper module is not updated.

Action

When you run the VPN version, you need to use the module `fwmodvpn 5.x.o`. To make the update, you can follow these steps:

```
# fwstop
# cd $FWDIR/modules
# mv fwmod.5.x.o old.fwmod.5.x.o
# ln -s fwmodvpn.5.x.o fwmod.5.x.o
# fw putlic 0 0-0-0 0      # For Firewall-1 2.x)
# fw putlic -K           # 3.x Firewall)
```

(continued)

(Continuation)

```
# fwstart
```

no VTOC

Cause

In this case, the user installs the Solaris 2.6 IA software and receives this error when rebooting the system. Other error messages refer to not having a default boot device configured, but this is the usual error message. This error leaves the system unusable; the user cannot boot.

Action

The user needs to do the following:

1. Insert the Solaris 2.6 software CD in the drive.
2. Boot with the Device Configuration Assistant diskette.
3. Select the CD-ROM to boot when presented with the available devices.
4. Type `b -s` when asked to select either Interactive or Jumpstart to boot as a single user.
5. At the `#` prompt, type the following:

```
# mount /dev/dsk/cxdxpx /a (where "x" is information from your system)
```

```
# TERM=at386; export TERM
# cd /a/platform/i86pc/boot/solaris/devicedb
```

6. In this directory is a file called `master`. BEFORE EDITING this file, make a backup copy. After it is backed up, view the master file in `vi`. Look for the term `ata.bef` and replace it with the word `none`.
7. Run `touch /reconfigure` and then reboot the system. (The command `boot -r, reboot -- -r` also works.)

Not a data message

Cause

During a `read(2)`, `getmsg(2)`, or `ioctl(2)` `I_RECVFD` call to a STREAMS device, some data has come to the head of the queue that cannot be processed. That data depends on the call:

1. `read(2)` — Controls information or passes a file descriptor
2. `getmsg(2)` — Passes a file descriptor
3. `ioctl(2)` — Controls data information

Technical Notes

The symbolic name for this error is `EBADMSG`, `errno=77`.

Not a directory

Cause

A non-directory was specified where a directory is required, such as a path prefix or an argument to the `chdir(2)` call.

Action

Look at a listing of all the files in the current directory and try again, specifying a directory instead of a file.

Technical Notes

The symbolic name for this error is `ENOTDIR`, `errno=20`.

Not a stream device

Cause

A `putmsg(2)` or `getmsg(2)` system call was attempted on a file descriptor that is not a STREAMS device.

Technical Notes

The symbolic name for this error is `ENOSTR`, `errno=60`.

Not enough space

Cause

This message indicates that the system is running many large applications simultaneously and has run out of swap space (virtual memory). It could also indicate that applications failed without freeing pages from the swap area. Swap space is an area of disk set aside to store portions of applications and data not immediately required in memory. Any data written during this condition is probably lost.

Action

Reinstall or reconfigure the system to have more swap space. A general rule is that swap space should be two to three times as large as physical memory. Alternatively, use `mkfile(1M)` and `swap(1M)` to add more swap area. This example shows how to add 16 Mbytes of virtual memory in the `/usr/swap` file (any file system with enough free space would work):

```
# mkfile 16m /usr/swap
# swap -a /usr/swap
```

To make this reconfiguration automatic at boot time, add the following line to the `/etc/vfstab` file:

```
/usr/swap - - swap - no -
```

Technical Notes

When calling the `fork(2)`, `exec(2)`, `sbrk(2)`, or `malloc(3C)` routine, a program asked for more memory than the system could supply. This is not a temporary condition; swap space is a system parameter.

The symbolic name for this error is `ENOMEM`, `errno=12`.

not found

Cause

This message indicates that the Bourne shell could not find the program name given as a command.

Action

Check the form and spelling of the command line. If that data looks correct, do a `echo $PATH` to see if the user's search path is correct. When communications are garbled, it is possible to unset a search path to such an extent that only built-in shell commands are available. Below is a command to reset a basic search path:

```
$ PATH=/usr/bin:/usr/ccs/bin:/usr/openwin/bin:.
```

If the search path looks correct, check the directory contents along the search path for missing programs or directories that are not mounted.

Not login shell

Cause

This message results when a user tries to use the `logout(1)` command from a shell other than the one started at login time.

Action

To quit a non-login shell, use the `exit(1)` command. Continue doing so until you have logged out.

See Also

For more general information on the login shell, see the section on customizing your work environment in the *Solaris Advanced User's Guide*.

Not on system console

Cause

A user tried to use the `login(1)` command to a system as the superuser (`uid=0`, which is not necessarily `root`) from a terminal other than the console.

Action

Log in to that system as a normal user, then run `su(1M)` to become superuser. To allow superuser logins from any terminal, comment out the `CONSOLE` line in `/etc/default/login` (this is not recommended for security reasons).

Not owner

Cause

Either an ordinary user tried to do something reserved for the superuser, or the user tried to modify a file in a way restricted to the file's owner or to the superuser.

Action

Switch user to root and try again.

Technical Notes

The symbolic name for this error is `EPERM`, `errno=1`.

Not supported

Cause

This version of the system does not support the feature requested, although future versions of the system might provide support.

Action

This is generally not a system message from the kernel, but an error returned by an application. Contact the vendor or author of the application for an update.

Technical Notes

The symbolic name for this error is `ENOTSUP`, `errno=48`.

NOTICE: */string*: out of inodes

Cause

The file system specified after the first colon probably contains many small files, exceeding the per-file system limit for inodes (file information nodes).

Action

If many small files were created unintentionally, remove them to resolve the problem. Otherwise, follow these steps to increase file system capacity for small files:

1. Make several backup copies of the file system on different tapes (for safety).
2. Change the machine to single-user mode.
3. Use the `newfs(1M)` command with the `-i` option to increase inode density for this file system. The following is an example:

```
# newfs -i 1024 /dev/rdisk/partition
```

4. Restore the file system from a backup tape.

Note - Increasing the inode density slightly reduces the total file system capacity.

NOTICE: vxvm: unexpected status on close

Cause

Every time the system boots (or is shut down), the message is displayed on the console. Sometimes the following message is also displayed on the console and in the `/var/adm/messages` file:

```
WARNING:
/iommu@0,10000000/sbus@0,10001000/SUNW,soc@2,0/SUNW,pln@a0000000,74127a/ssd@4,2
(ssd22):
Error for Command: <undecoded cmd 0x35>      Error Level: Fatal
Requested Block: 0      Error Block: 0
Vendor: CONNER      Serial Number: 93081LPT
Sense Key: Aborted Command
ASC: 0xb3 (<vendor unique code 0xb3>), ASCQ: 0x0, FRU: 0x0
WARNING:
/iommu@0,10000000/sbus@0,10001000/SUNW,soc@2,0/SUNW,pln@a0000000,74127a/ssd@4,2
(ssd22): ssd_synchronize_cache failed (5)
```

Action

In a High Availability system with NVRAM, this error would be caused by unprocessed data in a NVRAM cache of the active logical host that has been down and started again later. Because of the possibility of error, NVRAM should not be used in an HA system. The problem can be solved in this case by removing the NVRAM on the HA system.

In a non-HA system, this error can also be caused by stale data in the NVRAM cache. (The example commands that follow assume the controller for the array is `c1`.)
To fix for a non-HA system:

1. Turn off all fast writes on this array and sync any remaining pending writes:

```
# ssaadm fast_write -d c1
# ssaadm sync_cache c1
```

2. When you sync the fast writes to the array, all pending writes are physically made to the disks. Anything that is left in the cache is stale; thus, it is safe to purge it.
Run this command:

```
# ssaadm purge c1
```

3. Turn on the fast writes for the disks. This command might be different on your system, depending on the disks where you want fast writes enabled and the types of fast writes you want:

```
# ssaadm fast_write -s -e cl
```

nsrck: SYSTEM error, more space needed to compress [*client*] index, 8.1 MB required

Cause

In `networker`, you cannot use the Remove Oldest Cycle feature because the `/nsr` file system is too full to perform a remove. An error message appears in the console window indicating that the file system is full.

Action

1. Stop the `networker` daemons so that some of the indexes can be moved. In the SunOS 5 system, use `/etc/init.d/networker stop`. In the SunOS 4 system, use `ps -ef | grep nsr` and `kill(1)` the processes.
2. Find a file system with enough space to move one of the client's indexes. Only one of the client's indexes should be moved, not the `networker` server's index. To find the size of a client's index, go to `/nsr/index/clientname/db` and list the contents using `ls -l`. The database file can be large (possibly over 500 Mbytes).
3. Move the contents of a client's index to the other file system and check that `/nsr` has freed the space to use. You might need to unmount and remount `/nsr`, or even to reboot to designate the space freed by the move, as available.
4. After the space is available, restart the daemons.
5. Open `nwadmin`. Under `Clients--Indexes`, select a client and use `Remove Oldest Cycle` to free more space.
6. Use `Reclaim Space` to reclaim the space from the removed cycles. After a few of the old cycles have been removed, enough space should be in the file system to move the removed client's index back.
7. Stop the daemons, and move the client's index back to `/nsr/index/clientname`.
8. Restart the daemons. Remove the oldest cycles for the client that was just moved.

Tweaking of the browse policy and retention policy might be necessary to prevent this situation from happening in the future.

Otherwise, as long-term solutions, add more hard disk and run `growfs`, or move `/nsr` to a drive with more space on it.

"O"

Object is remote

Cause

This error occurs when users try to share a resource that is not on the local machine, or try to mount/unmount a device or path name that is on a remote machine.

Technical Notes

The symbolic name for this error is `EREMOTE`, `errno=66`.

ok

Cause

This is the OpenBoot PROM monitor prompt. From this prompt, you can `boot` the system (from disk, CD-ROM, or net), or you can use the `go` command to continue where you left off.

Action

If you suddenly see this prompt, look at the messages above it to see if the system crashed. If no other messages appear, and you just typed `STOP-A` or plugged in a new keyboard, type `go` to continue. You might need to Refresh the window system from its Workspace Menu.

Technical Notes

Never invoke `sync` from the prompt without first running the `fsck(1M)` command, especially if the file system has changed.

open: no such device or address from FW-1

Cause

The FW-1 has been installed on a disk other than the default root disk. If the Default Filter option is set (allowing a default filter to be automatically installed during boot), FW-1 tries to load the default security policy from \$FWDIR, but the partition that contains \$FWDIR is not yet mounted. This mismatch causes this error.

Action

To work around this problem, follow these steps:

```
# cp /$FWDIR/modules/fwmod.5.x.0 /etc/fw.boot/  
# cp /$FWDIR/modules/fw.mkdev /etc/fw.boot/  
# cp /$FWDIR/modules/fw.conf /etc/fw.boot/
```

Go to /usr/kernel/drv and change the links as follows:

```
fw -> /etc/fw.boot/fwmod.5.x.0  
fw.conf -> /etc/fw.boot/fw.conf
```

Operation already in progress

Cause

An operation was attempted on a non-blocking object that already had an operation in progress.

Technical Notes

The symbolic name for this error is EALREADY, errno=149.

Operation canceled

Cause

The associated asynchronous operation was canceled before completion.

Technical Notes

The symbolic name for this error is `ECANCELED`, `errno=47`.

operation failed [error 185], unknown group error 0, *string*

Cause

When you use *admintool* to add a user to a newly created group, *admintool* issues this error.

Action

Apply patch 101384-05 to fix bug ID 1151837 and to provide a workaround for bug ID 1153087.

Operation not applicable

Cause

This error indicates that no system support exists for a function that the application requested.

Action

Ask the system vendor for an upgrade, or contact the vendor or author of the application for an update.

Technical Notes

This message indicates that no system support exists for an operation. Many modules set this error when a programming function is not yet implemented. If you are writing a program that produces this message, while calling a system library, find and use an alternative library function. Future versions of the system might support this operation; check system release notes for further information.

The symbolic name for this error is `ENOSYS`, `errno=89`.

Operation not supported on transport endpoint

Cause

As an example, this error could occur when trying to accept a connection on a datagram transport endpoint.

Technical Notes

The symbolic name for this error is `EOPNOTSUPP`, `errno=122`.

Operation now in progress

Cause

An operation that takes a long time to complete (such as a *connect*) was attempted on a non-blocking object.

Technical Notes

The symbolic name for this error is `EINPROGRESS`, `errno=150`.

`/opt/bin/jws: /solaris/bin/locate_dirs: not found`

Cause

This error message occurs if you try to start Java Workshop by linking from `/opt/bin/jws` to `/opt/SUNWjws/JWS/sparc-S2/bin/jws`. Typing the full path name works, but typing `jws` gives this error.

Action

This error occurs because `/opt/bin/jws` is not `/opt/SUNWjws/JWS/sparc-S2/bin/jws`, which is a script that runs another script: `$_SS_JWS_HOME/solaris/bin/locate_dirs`.

`/opt/bin/jws` is not setting `$_SS_JWS_HOME` correctly. Remove it from the path and replace it with `/opt/SUNWjws/JWS/sparc-S2/bin/jws`. Then, which `jws` can return `/opt/SUNWjws/JWS/sparc-S2/bin/jws`.

Option not supported by protocol

Cause

A bad option or level was specified when getting or setting options for a protocol.

Technical Notes

The symbolic name for this error is `ENOPROTOOPT`, `errno=99`.

out of memory

Cause

Hundreds of different programs can produce this message when the system is running many large applications simultaneously. This message usually means that the system has run out of swap space (virtual memory).

Action

For details, refer to “Not enough space” on page 183. Any data written during this condition is probably lost.

Out of stream resources

Cause

During a STREAMS open, either no STREAMS queues or no STREAMS head data structures were available. This is a temporary condition; you might recover from it if other processes release resources.

Technical Notes

The symbolic name for this error is `ENOSR`, `errno=63`.

overlapping swap volume

Cause

After creating volumes in `rootdg` to be used as additional swap and adding these to the `/etc/vfstab` file, an error message is displayed at boot time that indicates overlapping swap volumes.

Action

Change the names of these volumes to read `swap1`, `swap2`, and so forth.

If you still get this message after making the previous change, edit the `/sbin/swapadd` script. Find the line:

```
c=${SWAP} -l | grep -c '\\<${special}'>''
```

and change it to:

```
c=${SWAP} -l | grep -c ''${special}''
```

"P"

Package not installed

Cause

This error occurs when a user attempts to use a system call from a package that has not been installed.

Technical Notes

The symbolic name for this error is `ENOPKG`, `errno=65`.

page_create: invalid flag

Cause

This error occurs after a `vxvm` upgrade. In this case, the user had the drivers (`vxio` and `vxspec`) for the Solaris 2.5.1 software and not for the Solaris 2.6 software. This condition was verified by using `ls -l /kernel/drv/*vx*`.

Action

Execute a `pkgrm` or re-install `VXVM 2.4` and re-encapsulate the root.

Panic

Cause

A system panics and crashes when a program exercises an operating system bug. Although the crash might seem unfriendly to a user, the sudden stop actually safeguards the system and its data from further corruption.

In addition to stopping the operating system, the panic routine copies the memory contents in use to a dump device, recording critical information about the current state of the CPU from which the panic routine was called.

Because the primary swap device is usually the default dump device, the primary swap device should be large enough to hold a complete image of memory. The system tries to reboot after the memory image is saved.

If the system does not reboot successfully, consider these possibilities:

1. Catastrophic hardware failure, such as faulty memory or a crashed disk
2. Major kernel configuration faults, such as an unstable device driver
3. Major kernel-tuning errors, such as a too-large value for MAXUSERS
4. Data corruption, including corruption of the operating system files
5. Manual intervention needed, as when `fsck(1M)` expects answers to its queries

Action

To find out why a system crashed, you can look in the `/var/adm/message*` log files.

Of these methods, using `savecore(1M)` is the most informative. The `savecore(1M)` command transfers the system crash dump image generated by the panic routine from the dump device to a file system. The image can then be analyzed with a debugger, such as `adb(1)`.

See Also

Correctly setting up `savecore(1M)` and interpreting its results can be difficult. For more information about debugging system panics, refer to *Panic! UNIX System Crash Dump Analysis* by Chris Drake and Kimberley Brown (ISBN 0-13-149386-8).

panic -boot: Could not mount filesystem

Cause

The first problem comes from the following jumpstart error:

```
2ec00 RPC: Can't decode result.
whoami RPC call failed with rpc status: 2
panic - boot: Could not mount filesystem.
program terminated
ok
```

Normally, this error occurs when the boot process is unable to get to the install image.

Additionally, other users have the same error message, with an additional message:

```
'Timeout waiting for ARP/RARP packet...'
```

Action

To solve the first problem:

1. Check how the `dfstab(4)` (`/etc/dfs/dfstab` on the install image NFS server) looks:

```
share -F nfs -o ro,anon=0 /jumpstart-dir
```

2. Run `share(1M)` command on the installed image NFS server, to make sure it is shared properly.
3. Check `/etc/bootparams` file on the net install server. Look for entries with incorrect boot path.
4. Make sure that `/usr/sbin/rpc.bootparamd` is running on the boot server. If necessary, kill and restart it.
5. Check `/etc/ethers` on the boot server for duplicate or conflicting entries.
6. At the prompt, run `test net /test-net` and/or `watch net /watch-net` to test the network connectivity.

As a workaround for the second problem, check the `nsswitch.conf(4)` file. If some of the entries point to NIS, such as:

```
rpc nis files
hosts nis files
ethers nis files
bootparams files nis
```

change all of these entries to files first:

```
rpc files nis
hosts files nis
ethers files nis
bootparams files nis
```

Note - You might have to update these files manually if they do not contain information on the client machine you are trying to jumpstart.

Then, remove the client with `rm_install_client(1M)`, remove the contents of `tftpboot`, and again add the client:

```
add_install_client -c /jumpstart-dir/profiles 'client name' 'arch'
```

Panic on cpu 0: valloc'd past tmpptes

Cause

The machine is an SS20 with 256 Mbytes of RAM, an FDDI interface, and a single CPU. It is running Online DiskSuite for mirroring and striping. The following recommended kernel patches were installed:

```
102517--03
102436--02
102394--02
102516--06
```

After their installation, the machine was rebuilt to allow for the new patches to be implemented. However, the machine panicked immediately after loading the kernel with this error message.

Action

The kernel was rebuilt with a new `MAXUSERS` value of 96, and this kernel enabled the machine to boot properly.

Technical Notes

Information directly related to this situation was not available; however, there was a description of another type of panic that was related to `seg_u`. In that description, the `MAXUSERS` value was set too large, causing the kernel to overrun table space. Furthermore, the value of `MAXUSERS` varies among the different architectures and the different revisions of the OS and is directly related to the amount of physical RAM in the system in an inverse proportion. Further investigation revealed that the value of `MAXUSERS` was set to 128. Based on the related information, it seems that the panic was due to `valloc` attempting to define memory space in excess of the value of `tmpptes`.

PARTIALLY ALLOCATED INODE I=*int* CLEAR?

Cause

Probably the system crashed in the middle of a `sync(2)` or `write(2)` operation, and during phase 1, `fsck(1M)` found that the specified inode was neither allocated nor unallocated.

Action

If any directory entries point to this inode and you answer “yes” to this question, phase 2 might get `UNALLOCATED` messages. Carefully exit `fsck(1M)` and run `ncheck(1M)`—specifying the inode number after the `-i` option—to determine which file or directory is involved. You might be able to restore this file or directory from another system. `fsck(1M)` also might copy this file to the `lost+found` directory in a later phase.

See Also

For more information, see the chapter on checking file system integrity in the *System Administration Guide, Volume 1*.

passwd: Changing password for *string*

Cause

The following lines are put into `/etc/nsswitch.conf`:

```
passwd:      compat
passwd_compat: nis
```

Then, when `passwd` is run, it fails as follows:

```
server1% passwd
passwd: Changing password for khh
server1%
```

Note - `passwd` exits before a password is entered.

Action

In the man page for `passwd`, you see the following:

If all requirements are met, by default, the `passwd(1)` command consults `/etc/nsswitch.conf` to determine which repositories need a password update. It searches the `passwd(4)` and `passwd_compat` entries. The sources (repositories) associated with these entries are updated. However, the supported password update configurations are limited to the following five cases. Failure to comply with the configurations prevents users from logging in to the system.

```
passwd: files
passwd: files nis
passwd: files nisplus
passwd: compat (==> files nis)
passwd: compat (==> files nisplus)
passwd_compat: nisplus
```

Note - The `passwd(1)` man page does NOT say that you can use the line:
`passwd_compat: nis.` `passwd(1)` works exactly as described in the man page.

`passwd (SYSTEM): System error: repository out of range`

Cause

When trying to lock a user account and using `nispaswd` with the `-l` option in the Solaris 2.6 release, you get this error: `passwd (SYSTEM): System error: repository out of range.`

Action

Use `passwd -r nisplus -l username` instead.

passwd.org_dir: NIS+ servers unreachable

Cause

This is the first of three messages that an NIS+ client prints when it cannot locate an NIS+ server on the network.

Action

For details, refer to “hosts.org_dir: NIS+ servers unreachable” on page 123.

Password does not decrypt secret key for *unix.uid@string*

Cause

This message appears at login when a user's password is not identical to the user's `keylogin(1)` network password. When a system is running NIS+, the login program first performs UNIX authentication, and then attempts a `keylogin(1)` for secure RPC authentication.

Action

To gain credentials for secure RPC, users can run `keylogin(1)` (after login) and type their secret key. To stop this message from appearing at time of login, users can run the `chkey -p` command and set their network password to be the same as their NIS+ password. If a user does not remember the network password, the system administrator should delete and re-create the user's credentials table entry so the user can establish a new network password with `chkey(1)`.

password file busy - try again later.

Cause

On a SunOS system running NIS (YP), the user runs `yppasswdd(1M)` and the system reports this error. On the NIS Master server, this error is in the messages file from `rpc.yppasswdd: password file busy - try again`. This error is

caused superficially by the existence of a lock file, `/var/yp/passwd.ptmp`. Removing this file allows `yppasswdd` to run to completion, but subsequent invocations still fail with the same error message. The root cause is that `yppasswdd` has the `-m` option, which says to run `make` to push the maps out to the slave servers. In this situation, a problem occurred in pushing the maps to a slave server; the push would hang. Thus, the push was never completed, and the lock file was never removed. This was tested by doing the following:

```
#cd /var/yp
#make passwd
passwd is up to date
#touch passwd
#make passwd
```

From here, the `make` remakes the map, but then hangs on the `push` to the slave.

Action

To fix the root cause, find out why the map does not push. In this situation, it was a routing issue; however, the remedy could lie elsewhere.

`pdbadm` start node fails cluster_establish join not allowed

Cause

The user created a disk group, but forgot to make it shared. After it was made a shared disk group, the user attempted to start the second node (which had not been rebooted). `pdbadm start node` on second `pdb` node failed with this repeated message until it finally timed out:

```
return from cluster_establish is join not allowed now
retrying cluster_establish
```

Action

You can either reboot the second node or run `vxdctl enable`.

`pdadmin start node now works.`

Permission denied

Cause

An attempt was made to access a file in a way forbidden by the protection system.

Action

Check the ownership and protection mode of the file (with a long listing from the `ls -l` command) to see who is allowed access to the file. Then change the file or directory permissions, as needed.

Technical Notes

The symbolic name for this error is `EACCES`, `errno=13`.

Please specify a recipient.

Cause

With *mailto(1)*, this message comes up in a dialog box whenever a user tries to deliver a message with no address in the `TO:` field.

Action

For details, refer to “Recipient names must be specified” on page 208.

Protocol error

Cause

A protocol error occurred. This error is device specific, but is generally not related to a hardware failure.

Technical Notes

The symbolic name for this error is `EPROTO`, `errno=71`.

protocol error, *string* closed connection

Cause

`rlogin(1)` fails on a machine with the SunOS system installed.

Action

1. Check the permissions in `in.rlogind` on the machine you are trying to connect to. The permissions should look like this:

```
-rwxr-xr-x 1 root  staff  16384 Jan 20 1994 /usr/sbin/in.rlogind
```

2. Check the login line in the `/etc/inetd.conf` file. It should look like the following:

```
login stream tcp nowait root /usr/sbin/in.rlogind in.rlogind
```

3. Check `/etc/passwd` to see if an invalid login shell has been substituted in the entry for the login ID.

Protocol family not supported

Cause

The protocol family has not been configured into the system or no implementation for it exists. This is used for the Internet protocols.

Technical Notes

The symbolic name for this error is `EPFNOSUPPORT`, `errno=123`.

Protocol not supported

Cause

The requested networking protocol has not been configured into the system, or no implementation for it exists. (A protocol is a formal description of the messages to be exchanged and the rules to be followed when systems exchange information.)

Action

Verify that the protocol is in the `/etc/inet/protocols` file and in the NIS protocols map, if applicable. If the protocol is not listed, and you want to permit its use, configure the protocol as documented or as required.

Technical Notes

The symbolic name for this error is `EPROTONOSUPPORT`, `errno=120`.

Protocol wrong type for socket

Cause

This message indicates either an application programming error, or badly configured protocols.

Action

Make sure that the `/etc/protocols` file corresponds number-for-number with the NIS `protocols(4)` map. If it does, ask the vendor or author of the application for an update.

Technical Notes

A protocol was specified that does not support the semantics of the socket type requested. This protocol amounts to a request for an unsupported type of socket. Look at the source code that made this socket request and check that it requested one of the types specified in `/usr/include/sys/socket.h`.

The symbolic name for this error is `EPROTOTYPE`, `errno=98`.

"Q"

quotactl: open Is a directory

Cause

When using `edquota` to set user limits, the command displays this error. `edquota` updates all quota files that are on a mounted file system. A directory named `quotas` causes it to fail.

Action

In one of the mounted file systems is a directory named `quotas`. To fix the problem, move the directory from the mounted file system and rename or delete it.

For example: If you have `/usr/quotas/old_info`, the directory `/usr/quotas` will cause `edquota` to fail. Either move `/usr/quotas` to `/usr/old_quotas` or delete the directory.

"R"

Read error from network: Connection reset by peer

Cause

This message appears when a user logs in remotely to a machine that crashes or is rebooted during the `rlogin(1)` or `rsh(1)` session. Any data changes that were not saved are probably lost. Sometimes this message appears only when the user types some data, even though the system failed hours before.

Action

Try to `rlogin(1)` again, perhaps after waiting a few minutes for the system to reboot.

Reading configuration data

Cause

In this situation, the user loaded SunPC 4.1 on a SPARCstation 5 machine. The Solaris 2.5 operating environment is patched to the Solaris 2.5.1. The user also has a SunPC accelerator card installed. When starting SunPC, the user gets this error message on the SunPC splash screen. If the user clicks anywhere in the screen, the whole console locks. The user has to move to another machine and use `rlogin` and then `kill` the SunPC process. In an effort to resolve the problem, the user had installed and removed SunPC and the 102924-25 patch with the same results. The user also removed the accelerator card, performed a `boot -r` and still SunPC 4.1 hung at the splash screen. The following error was found in the `/var/adm` file:

```
modrput() sdos_mbsigolint failed -1
```

Action

In this situation, the user had wiped the operating system off the SPARCstation 5 machine and, at that point, was not sure which patches had been applied. The user installed a copy of the Solaris 2.5.1 software and, then, performed the SunPC installation. That solved the problem. SunPC worked without the Accelerator card. The user added the Accelerator card, performed a `boot -r`, and ran SunPC with no problems.

Read-only file system

Cause

Files and directories on file systems that are mounted read-only cannot be changed.

Action

If you only modify these files and directories occasionally, use `rlogin(1)` to log in to the servers of the mounted file systems and change the files or directories from there.

If you change these files and directories frequently, use `mount(1M)` to make the file systems read-write.

Technical Notes

The symbolic name for this error is `EROFS`, `errno=30`.

rebooting...

Cause

This message appears on the console to indicate that the machine is booting, either after the superuser issued a `reboot(1M)` command, or after a system panic, if the EEPROM's `watchdog-reboot?` variable is set to true.

Action

Allow the machine to boot itself. In case of a system panic, look above this message for other indications of what went wrong.

Recipient names must be specified

Cause

Someone sent mail without a valid recipient in the `To:` field. Thus, `sendmail(1M)` could not deliver the mail message. Using `mail(1)`, the recipient's address might have been specified using spaces or non-alphanumeric characters. The *mailtool(1)* and `mailx(1)` commands try to prevent such problems by issuing `Please specify a recipient` or `No recipients specified` messages instead. If at least one valid recipient exists, each invalid recipient address will generate a `User unknown` message.

Action

Look in the sender's `dead.letter` file for the automatically saved message, and have the originator send it again; this time the sender specifies a recipient.

See Also

For more information about `sendmail(1M)`, see the *System Administration Guide, Volume 3*.

refused connect from *hostIP* to callit(ypserv)

Refer to “connect from *hostIP* to callit(ypserv): request from unauthorized host” on page 71.

Reset tty pgrp from *int* to *int*

Cause

The C shell sometimes issues this message when it clears away the window process group after the user exits the window system. This clearing can happen when the window system does not clean up after itself.

Action

Proceed with your work. This message is only informational.

Resource temporarily unavailable

Cause

This error indicates that the `fork(2)` system call failed because the system's process table is full, or that a system call failed because of insufficient memory or swap space. Also, a user might not be allowed to create more processes.

Action

Simply waiting often gives the system time to free resources. However, if this message occurs often on a system, reconfigure the kernel and allow more processes. To increase the size of the process table, increase the value of `MAXUSERS` in the `/etc/system` file. The default `MAXUSERS` value is the amount of main memory in Mbytes, minus 2.

If one user is not allowed to create any more processes, that user has probably exceeded the memory size limit; see the `limit(1)` man page for details.

Technical Notes

The symbolic name for this error is `EAGAIN`, `errno=11`.

Restartable system call

Action

Restart the interrupted system call.

Technical Notes

The symbolic name for this error is `ESTART`, `errno=91`.

Result too large

Cause

This is a programming error or a data input error.

Action

Ask the program's author to fix this condition.

Technical Notes

This error indicates an attempt to evaluate a mathematical programming function at a point where its value would overflow or underflow. The value of a programming function in the math package (3M) is not representable within machine precision. This error could occur after floating point overflow or underflow (either single or double precision), or after total loss of numeric significance in Bessel functions.

This message can indicate `Result too small` in the case of floating point underflow.

To help pinpoint a program's math errors, use the `matherr(3M)` facility.

The symbolic name for this error is `ERANGE`, `errno=34`.

rlogin: no directory! connection closed

Cause

When a user tries to remotely log in to a machine, the user gets this error.

The machine that the user was trying to access with `rlogin(1)` had permissions of 700 on its `root` directory. The permissions on `root` should be 755.

After the permissions on the `root` file system were changed to 755, the user was able to proceed farther when attempting to execute an `rlogin`, but it still failed with the following:

```
Last login: Fri Aug 29 10:24:43 from machinename
no shell
connection closed
```

Action

The machine that the user was trying to access with `rlogin` had the permissions set to 700 on both the `root` and `/usr/bin` directories. For both directories, the permissions should be 775. Once the user changed the permissions to 775, `rlogin(1)` was successful.

Also, check the user's `passwd(1)` entry in the NIS/NIS+ map. A login shell such as `/usr/dist/exe/tcsh` or `/net/lab/.../csh` could cause the failure because of NFS mount permission.

rmdir: *string*: Directory not empty

Cause

The `rmdir(1)` command can only remove empty directories. The directory with the name appearing after the first colon in the message still contains some files or directories.

Action

Use `rm(1)` instead of `rmdir(1)`. To remove this directory and everything underneath it, use the `rm -ir` command to descend the directory recursively, and respond to

requests to delete each element. To remove the directory and all its contents without prompts for approval, use the `rm -r` command.

ROOT LOGIN /dev/console

Cause

This syslog message indicates that someone has logged in as `root` on the system console.

Action

If you have just logged in as `root`, take no action. If you are not `root`, consider the possibility of a security breach. The best site-wide policy is for all system administrators to use `su(1M)` instead of logging in as `root`.

ROOT LOGIN /dev/pts/*int* FROM *string*

Cause

This syslog message indicates that someone has logged in remotely as `root` on a pseudo-terminal from the system specified after the `FROM` keyword.

Action

For security reasons, it is a bad practice to allow `root` logins from anywhere other than the console. To restrict superuser logins to the console, remove the comment from the `CONSOLE` line in `/etc/default/login`.

route: socket: Protocol not supported

Cause

During a boot, this error is displayed and the multicast is not configured.

Action

An `inittab(4)` from a previous release of the operating environment was used. Thus, the following entry, which is required for the `route` command in the Solaris 2.6 release, was missing from `/etc/inittab`.

```
ap::sysinit:/sbin/soconfig -f /etc/sock2path
```

By default, this is the second entry in the file. After this entry was added, the multicast configured at boot time without error.

RPC: Program not registered

Cause

Check the `rpc.bynumber` NIS map.

rx framing error

Cause

Usually this error indicates a hardware problem.

Action

Check the Ethernet cabling and connectors to locate a problem.

Technical Notes

A framing error occurs when the Ethernet I/O driver receives a non-integral unit of octets, such as 63 bytes and then 3 bits. (Ethernet specifies the use of octets.) Framing errors are caused by corruption of the starting or ending frame delimiters. These delimiters can be corrupted by some violation of the encoding scheme.

Framing errors are a subset of CRC errors, which are usually caused by anomalies on the physical media. An alignment/framing error is a type of CRC error where octet boundaries do not align.

"S"

save: SYSTEM error, Arg list too long

Cause

The save fails with this error because the database (index) file for the client is greater than 2 Gbytes. With the Solaris 2.6 release and SBU 5.0.1 this is no longer a problem.

Action

However, with earlier versions of the Solaris software you need to open `nwadmin` -> indexes -> select appropriate client -> select appropriate fs -> remove oldes cycle -> reclaim space.

You might have to repeat a few times to reclaim enough space. The indexes can be re-created later, if necessary, by using a scanner.

SCSI bus DATA IN phase parity error

Cause

The most common cause of this problem is unapproved hardware. Some SCSI devices for the PC market do not meet the high I/O speed requirements for the UNIX market. Other possible causes of this problem are improper cabling or termination, and power fluctuations. Data corruption is possible, but unlikely to occur, because this parity error prevents data transfer.

Action

Check that all SCSI devices on the bus are Sun-approved hardware. Then verify that all cables measure no longer than six meters total and that all SCSI connections are properly terminated. If power fluctuations are occurring, invest in an uninterruptible power supply.

SCSI transport failed: reason 'reset'

Cause

This message indicates that the system sent data over the SCSI bus, but the data never reached its destination because of a SCSI bus reset. The most common cause of this condition is conflicting SCSI targets. Data corruption is possible, but unlikely to occur, because this failure prevents data transfer.

Action

Verify that all cables measure no longer than six meters total and that all SCSI connections are properly terminated. If power surges are a problem, acquire a surge suppressor or an uninterruptible power supply.

A machine's internal disk drive is usually SCSI target 3. Make sure that external and secondary disk drives are targeted to 1, 2, or 0, and do not conflict with each other. Also, make sure that tape drives are targeted to 4 or 5, and CD drives to 6, avoiding any conflict with each other or with disk drives. If the targeting of the internal disk drive is in question, power off the machine, remove all external drives, turn on the power, and from the PROM monitor run the `probe-scsi-all` or `probe-scsi` command.

If SCSI device targeting is acceptable, memory configuration could be the problem. Ensure that high-capacity memory chips (such as 4-Mbyte SIMMs) are in lower banks, while lower-capacity memory chips (such as 1-Mbyte SIMMs) are in the upper banks.

SPARC systems do not always support third-party CD-ROM drives, and can generate a similar `unknown vendor` error message. Check with the CD-ROM vendor for specific configuration requirements.

Some third-party disk drives have a read-ahead cache that interferes with the Solaris device drivers. Make sure that any existing read-ahead cache facility is turned off.

See Also

For more information on SCSI targets, see the section on device naming conventions in the *Solaris Transition Guide*. If you are using AnswerBook online documentation, "SCSI targets" is a good search string.

Security exception on host *string*. USER ACCESS DENIED.

Cause

When trying to create a user with `Adminsuite` by placing the home directory on a system remote from the NIS+ server, the user gets this error message:

```
Security exception on host hostname. USER ACCESS DENIED.  
The user identity (555)username was received, but that user  
is not authorized to execute the requested functionality  
on this system. Is this user a member of an appropriate  
security group on this system ?  
(Function: class directory method create_dir)
```

The user can use `rsh(1)` to access the remote machine and create a home directory on the system.

Action

The user was not in the system administration group NIS+ tables.

```
# niscat group.org_dir | grep sysadmin  
sysadmin::14:
```

Add the user name to the system administration group.

Segmentation Fault

Cause

Segmentation faults usually come from a programming error. This message is usually accompanied by a core dump, except on read-only file systems.

Action

To see which program produced a `core(4)` file, run either the `file(1)` command or the `adb(1)` command. The following examples show the output of the `file(1)` and `adb(1)` commands on a core file from the `dtmail` program.

```
$ file core
core: ELF 32-bit MSB core file SPARC Version 1, from 'dtmail'
```

```
$ adb core
core file = core -- program 'dtmail'
SIGSEGV 11: segmentation violation
^D      (use Control-d to quit the adb rogram)
```

Ask the vendor or author of this program for a debugged version.

Technical Notes

A process has received a signal indicating that it attempted to access an area of memory that is protected or that does not exist. The two most common causes of segmentation faults are attempting to dereference a null pointer or indexing past the bounds of an array.

**sendmail[]: can't lookup data via name server
"dns" or sendmail[]: can't lookup data via name
server "nis"**

Cause

The following entry in the `/etc/nsswitch.conf` file, `sendmailvars: dns nis files`, causes the messages to appear in the console window.

Action

The `sendmailvars` database can be used only with local files and/or NIS+. If you do not have this database setup, the default `sendmailvars` entry should look as follows in the `/etc/nsswitch.conf` file:

```
sendmailvars: files
```

sendmail[init]: NOQUEUE: SYSERR(root): Cannot bind to domain <domain>: no such map in server's domain: Bad file number

Cause

The user is running NIS and receives this error on several NIS machines.

Action

Check the following:

1. For the system(s) not working, make sure there is a `/var/yp/nicknames` file. Also, make sure that this file contains this entry: `aliases mail.aliases`
2. On one of the systems not working, execute the following:

```
ypcat aliases
```

You will probably get this message: `no such map in servers domain`. Do a `ypwhich` to see which NIS server the system is bound to. Next, go to that server and verify that the `mail.aliases` map is missing from `/var/yp/domainname`. This map must either be created or copied over from one of the NIS servers that contains the map.

sendmail[int]: NOQUEUE: SYSERR: net hang reading from *string*

Cause

This is a `sendmail(1M)` message that appears on the console and in the log file `/var/adm/messages`. If this message occurs once for a particular user, a mail message from this user might end with a partial line (having no terminating newline character). If this message appears frequently or at busy times, especially along with other networking errors, it could indicate network problems.

Action

Check the user's mail spool file to see if a message ends without a newline character. If so, talk with the user and determine how to prevent the problem from occurring again. If these messages are the result of network problems, you could try moving the mail spool directory to another machine with a faster network interface.

Technical Notes

During the SMTP receipt of DATA phase, a message-terminating period on a line of its own never arrived. `sendmail(1M)` timed out and produced this error.

Service wouldn't let us acquire selection

Cause

This message indicates that the OpenWindows selection service failed to seize the requested selection from `/tmp/winselection`.

Consider the following diagnostics: the requested selection could be 0 for unknown, 1 for caret, 2 for primary, 3 for secondary, or 4 for clipboard. The result could be 0 for failure, 2 for nonexistent, 3 for did not have, 4 for wrong rank, 5 for continued, 6 for cancelled, or 7 for unrecognized.

setmnt: Cannot open `/etc/mnttab` for writing

Cause

The system is having problems writing to `/etc/mnttab`. The file system containing `/etc` might be mounted read-only, or not mounted at all.

Action

Check that this file exists and is writable by root. If so, ensure that the `/etc` file system has been mounted, and is mounted read-write, rather than read-only.

share_nfs: `/home`: Operation not applicable

Cause

This message usually indicates that the system has a local file system mounted on `/home`, which is where the automounter usually mounts users' home directories.

Action

When a system is running the automounter, do not mount local file systems on the `/home` directory. Mount them on another directory, such as `/disk2`, which on most systems you have to create. You could also change the automounter `auto_home` entry, but that is a more difficult solution.

Signal 8 error

Cause

In this case, the user gets a Signal 8 error during installation—right after starting Openwindows—and installation stops.

Action

Shut down the system “gracefully,” and, as it is rebooting, place a ZIP drive cartridge (blank or used) in the ZIP drive. Begin the normal installation of the Solaris IA software. It is not possible to continue the existing installation of the Solaris software by putting a cartridge in the ZIP drive after receiving this error. When the Solaris software checks all of your hardware, it thinks the ZIP drive is just another hard drive and attempts to read from it. If there is no cartridge in the drive, then you receive the `signal 8 error`. If the Solaris software installation “sees” a cartridge in the ZIP drive, it reads from it, even if there is no data on the cartridge, and then continues.

SIMS license error: licenses invalid

Cause

This is a license internet mail server problem. The user is installing a departmental version of SIMS 3.1 on a Pentium 2 PC that is running the Solaris 2.6 IA release. The system is using a Java™ interface and keeps getting the above error. The two license files from the license center are:

```
SERVER server
DAEMON lic.SUNW /etc/opt/licenses/lic.SUNW
INCREMENT SLAPD.1 lic.SUNW 1.000 08-Mar-1998 1

SERVER nwlab4 727a2b6a 7588
```

(continued)

(Continuation)

```
DAEMON suntechd /etc/opt/licenses/suntechd /etc/opt/licenses/daemon_options
INCREMENT sun.mail.mbox suntechd 3.100 08-Mar-1998 100
```

Action

Merge the two license files together and delete the extra SERVER line.

Slice c0t1d0s0 is too small to contain 1 replicas

Cause

When trying to add a state replica using metatool to cylinder 0 of a disk, the following error message appears:

```
Your attempt to attach metastate database
replicas on slice "c?t?d?s?" failed for the
following reason: Slice c?t?d?s? is too small
to contain 1 replicas.
```

This is because metatool masks out the very first cylinder to protect the disk label. On disksuite v4.1, metatool does allow adding the databases to cylinder 0 on 2.1Gbyte disks or larger.

Action

As a workaround, start at cylinder 1 (not cylinder 0) or use the command line (metadb -a).

snmpdx: bind() failed on udp on 161 [errno: address already in use] 125 snmpdx dmid: unable to connect to snmpdx

Cause

The user is running the Solaris 2.6 release with a Cisco FDDI card and is receiving the above error.

Action

In the Solaris 2.6 software a startup script is included in `/etc/rc3.d` that starts `snmpdx` (which uses port 161). You receive the error message because the FDDI SNMP agent is running, and it has already claimed port 161. Two solutions are:

1. Move the `snmpdx` start-up script

```
mv /etc/rc3.d/S76snmpdx /etc/rc3.d/s76snmpdx
```

so that `snmpdx` does not start.

2. Check if the FDDI can use a different port, other than 161.

Socket type not supported

Cause

The support for the socket type has not been configured into the system or no implementation for it exists.

Technical Notes

The symbolic name for this error is `ESOCKTNOSUPPORT`, `errno=121`.

Soft error rate (*int%*) during writing was too high

Cause

This message from the SCSI tape drive appears when Exabyte or DAT tapes generate too many soft (recoverable) errors. It is followed by the advisory `Please, replace`

tape cartridge message. Soft errors are an indication that hard errors could soon occur, causing data corruption.

Action

First, clean the tape head with a cleaning tape, as recommended by the manufacturer. If that remedy does not work, replace the tape cartridge. If the problem persists, you might need to replace the tape drive with new tape cartridges.

Software caused connection abort

Cause

A connection abort occurred internally to your host machine.

Technical Notes

The symbolic name for this error is `ECONNABORTED`, `errno=130`.

Srmount error

Cause

This error is RFS specific. It occurs when an attempt is made to stop RFS while resources are still mounted by remote machines, or when a resource is readvertised with a client list that does not include a remote machine with the resource currently mounted.

Technical Notes

The symbolic name for this error is `ESRMNT`, `errno=69`.

Stale NFS file handle

Cause

A file or directory that was opened by an NFS client was either removed or replaced on the server.

Action

If you were editing this file, write it to a local file system instead. Try remounting the file system on top of itself or shutting down any client processes that refer to stale file handles. If neither of these solutions works, reboot the system.

Technical Notes

The original `vnode` is no longer valid. The only way to remove this error is to force the NFS server and client to renegotiate file handles.

The symbolic name for this error is `ESTALE`, `errno=151`.

start up failure no such file or directory

Refer to “late initialization error” on page 143.

statd: cannot talk to statd at *string*

Cause

This message comes from the NFS status monitor daemon `statd(1M)`, which provides crash recovery services for the NFS lock daemon `lockd(1M)`. The message indicates that `statd(1M)` has left old references in the `/var/statmon/sm` and `/var/statmon/sm.bak` directories. After a user has removed or modified a host in the `hosts` database, `statd(1M)` might not properly purge files in these directories, which results in its trying to communicate with a nonexistent host.

Action

Remove the file named *variable* (where *variable* is the host name) from both the `/var/statmon/sm` and `/var/statmon/sm.bak` directories. Then kill the

`statd(1M)` daemon and restart it. If that does not get rid of the message, kill and restart `lockd(1M)` as well. If that remedy does not work, reboot the machine at your convenience.

stty: TCGETS: Operation not supported on socket

Cause

This message occurs when a user tries to use remote copy with `rcp(1)` or remote shell with `rsh(1)` from one machine to another, but has an `stty(1)` command in the remote `.cshrc` file. This error creates failure for the `rcp(1)` or `rsh(1)` command.

Action

The solution is to move the invocation of the `stty(1)` command to the user's `.login` (or equivalent) file. Alternatively, execute the `stty(1)` command in `.cshrc` only when the shell is interactive. You could perform the following test:

```
if ($?prompt) stty ...
```

Technical Notes

The `rcp(1)` and `rsh(1)` commands make a connection using sockets, which do not support `stty(1)`'s TCGETS *ioctl*.

su: No shell

Cause

This message indicates that someone changed the default login shell for `root` to a program that is missing from the system. For example, the final colon-separated field in `/etc/passwd` could have been changed from `/sbin/sh` to `/usr/bin/bash`, which does not exist in that location. Possibly an extra space was appended at the end of the line. The outcome is that you cannot login as `root` or switch user to `root`, and, thus, cannot directly fix this problem.

Action

The only solution is to reboot the system from another source, then edit the password file to correct this problem. Invoke `sync(1M)` several times, then halt the machine by typing `Stop-A` or by pressing the reset button. Reboot as single-user from CD-ROM, the net, or diskette, such as by typing `boot cdrom -s` at the prompt.

After the system starts and gives you a `#` prompt, mount the device corresponding to the original `root` partition somewhere, such as with a `mount(1M)` command similar to the one that follows. Then run an editor on the newly mounted system password file (use `ed(1)` if terminal support is lacking):

```
# mount /dev/dsk/c0t3d0s0 /mnt
# ed /mnt/etc/passwd
```

Use the editor to change the password file's `root` entry to call an existing shell, such as `/usr/bin/csh` or `/usr/bin/ksh`.

Technical Notes

To keep the `No shell` problem from happening, habitually use *admintool* or `/usr/ucb/vipw` to edit the password file. These tools make it difficult to change password entries in ways that make the system unusable.

su: 'su root' failed for *login* on `/dev/pts/int`

Cause

The user specified by *login* tried to become superuser, but typed the wrong password.

Action

If the user is supposed to know the `root` password, wait to see if the correct password is supplied. If the user is not supposed to know the `root` password, ask why he or she is attempting to become superuser.

su: 'su root' succeeded for *login* on */dev/pts/int*

Cause

The user specified by *login* just became superuser by typing the `root` password.

Action

If the user is supposed to know the `root` password, this message is only informational. If the user is not supposed to know the `root` password, change this password immediately and ask how the user learned it.

SunPC may NOT run correctly as root

Cause

With SunPC 4.1 and the 102924 jumbo patch installed, a user (who is not `root`) attempts to run SunPC and receives the following error message:

```
SunPC may NOT run correctly as root.  
Please run in user mode.  
SunPC script is exiting
```

The user's primary group ID is probably `root`. For example:

```
$ /usr/bin/id  
uid=33650(gruff) gid=0(root)
```

Action

Change the user's primary group to another group, such as `10`, and, because the user still needs to be in the `root` group, add the `root` group to the user's secondary group list.

syncing file systems...

Cause

This message indicates that the kernel is updating the super-blocks before taking the system down to ensure file system integrity. This message appears after a `halt(1M)` or `reboot(1M)` command. It can also appear after a system panic, in which case the system might contain corrupted data.

Action

If you just halted or rebooted the machine, take no action. This message is normal. In case of a system panic, look up the panic messages. Your system vendor might be able to help diagnose the problem. So that you can describe the panic to the vendor, either leave your system in its panicked state or be sure that you can reproduce the problem.

Technical Notes

Numbers that sometimes display after the three dots in the message show the count of dirty pages that are being written out. Numbers in brackets show an estimate of the number of busy buffers in the system.

syslog service starting.

Cause

During system reboot, this message might appear and the boot seemingly hangs. After starting `syslogd(1M)` service, the system runs `/etc/rc2.d/S75cron`, which in turn calls `ps(1)`. Sometimes after an abrupt system crash `/dev/bd.off` becomes a link to nowhere, causing the `ps(1)` command to hang indefinitely.

Action

Reboot as a single user (for example with `boot -s`) and run `ls -l /dev/bd*` to see if this is the problem. If so, remove `/dev/bd.off`, then run `bdconfig off` or reboot with the `-r` (reconfigure) option.

This is the most commonly reported situation that causes `ps(1)` to hang.

System booting after fatal error FATAL

Cause

The system reboots automatically. Afterward, the messages file contains `System booting after fatal error FATAL`.

The message is issued during a reboot after the system detects a hardware error. The following can cause this response: UPA address parity error, Master queue overflows, DTAG parity errors, E-Cache tag parity errors, and Coherence errors.

Action

Use `prtdiag(1M)` to help identify failed hardware components. The errors indicate that you either have a bad CPU module or a bad system board.

SYSTEM error, Arg list too long

Cause

When trying to back up a client with networker, the following error occurs:

```
* heaven.com:/export/heaven2 save: SYSTEM error, Arg list too long
* heaven.com:/export/heaven2 save: Cannot open save session with heaven.com
* heaven.com:/export/heaven3 1 retry attempted
* heaven.com:/export/heaven3 save: SYSTEM error, Arg list too long
* heaven.com:/export/heaven3 save: Cannot open save session with heaven.com
```

Action

An error like this is due to an index file (`/nsr/index/clientname`) that is greater than 2 Gbytes in Solstice backup revisions less than 5.0.1. In 5.0.1 the indexes are segmented so this error should no longer be a problem. In any revision of Solstice backup this error can also be due to a corrupt client index. If so, running the following command might resolve the problem:

```
# nsrck -F clientname
```

If this remedy does not fix the problem, shut down the networker daemons, remove the client index, and restart the daemons. The backup should then run fine.

system hang

Cause

4.1.3C Sbus cards suffered a system freeze.

SYSTEM HANGS DURING BOOT

Cause

When the user boots a system, it hangs after the following boot messages: `root on`, `swap on`, and `dump on`. After the system displays these messages, the LEDs flash and the system hangs.

This is due to an earlier `fsck` that deleted devices under the `/dev` directory. Check for the `/dev/console` device and, if it is missing, create one.

system will not connect to port 80

Refer to “late initialization error” on page 143.

"T"

tar: /dev/rmt/0: No such file or directory

Cause

The default tape device `/dev/rmt/0` or possibly the device specified by the `TAPE` environment variable is not currently connected to the system, is not configured, or its hardware symbolic link is broken.

Action

List the files in the `/dev/rmt` directory to see which tape devices are currently configured. If none are configured, ensure that a tape device is correctly attached to the system, and reboot with the `-r` option to reconfigure devices.

If tape devices other than `/dev/rmt/0` are configured, you could specify one of them after the `-f` option of `tar(1)`.

tar: directory checksum error

Cause

This error message from `tar(1)` indicates that the checksum of the directory and the files it has read from tape does not match the checksum advertised in the header block. Usually this message indicates the wrong blocking factor, although it could indicate corrupt data on tape.

Action

To resolve this problem, make certain that the blocking factor you specify on the command line (after `-b`) matches the blocking factor originally specified. If in doubt, leave out the block size and let `tar(1)` determine it automatically. If that remedy does not help, the tape data could be corrupted.

tar: tape write error

Cause

A physical write error has occurred on the `tar(1)` output file, which is usually a tape, although it could be a diskette or disk file. Look on the system console, where the device driver should provide the actual error condition. The condition might be a write-protected tape, a physical I/O error, an end-of-tape condition, or a file-too-large limitation.

Action

In the case of write-protected tapes, enable the write switch. For physical I/O errors, replace the tape with a new one. For end-of-tape conditions, try using a higher density, if the device supports one, or use `cpio(1)` or `pax(1)` for their multi-volume

support. When encountering the file-too-large limitations, use the parent shell's `limit(1)` or `ulimit(1)` facility to increase the maximum file size.

See Also

For more information on tar tapes, see the section on copying UFS files in the *System Administration Guide, Volume 1*.

Text file busy

Cause

This error can occur when an attempt was made to execute a pure-procedure program that is currently open for writing. It also occurs when attempting to open for writing or to remove a pure-procedure program being executed. (This message is obsolete.)

Technical Notes

The symbolic name for this error is `ETXTBSY`, `errno=26`.

Text is lost because the maximum edit log size has been exceeded.

Cause

This message appears at the beginning of a `cmdtool(1)` session after 100,000 characters have scrolled by. Clicking the top rectangle of the scrollbar might display this message. No data were lost, but the user cannot scroll back before this wraparound point.

Action

To increase the maximum size of the Command Tool log file, use `cmdtool -M`, specifying more than 100,000 bytes.

tftpd: nak: Transport endpoint is already connected

Cause

After configuring an Autoclient (Autoclient 2.1 - Solstice AdminSuite 2.3), particularly on a Solaris 2.6 environment, you might get a similar error message on your Server from `/dev/console` and/or from `/var/adm/messages`:

```
tftpd: nak: Transport endpoint is already connected
```

A subsequent boot net by the Autoclient hangs. For example:

```
Boot Device:...  
File and Args...
```

Note - This error message is difficult to decipher. Also, at this early point in the autoclient's boot, there is a minimum record of the event. To troubleshoot this problem, a snoop of the client, run from another system on the client's subnet, is necessary.

Action

A change was made in the Solaris 2.6 `in.tftpd` to use `sendto()`, instead of `send()`. Because the Solaris 2.5.1 environment uses `send()` as opposed to `sendto()`, one workaround would be to copy `in.tftpd` from a Solaris 2.5.1 to the Solaris 2.6 environment. Another workaround would be to troubleshoot from the server the nonexistent file that it is trying to receive by doing a `snoop` of the client.

For example (assuming you are using an onboard Ethernet interface):

```
# snoop autoclient_name
```

or

```
# snoop ethernet_address_of_autoclient_name
```

In this case, you might get a Trivial File Transfer Protocol (TFTP) read similar to the following:

```
81911ED4.SUN4C
TFTP Error: access violation
```

The error tells you that something is wrong within your `/tftpboot` directory.

For an AUTOCLIENT: The problem lies in the `/tftpboot` directory of the boot server. Confirm that the `HOSTID` and `HOSTID.ARCH` files are linked to the correct `inetboot` file for your architecture. This is a correct entry for a sun4m system:

```
81971904 -> inetboot.sun4m.Solaris_2.4
81971904.SUN4M -> inetboot.sun4m.Solaris_2.4
```

This is an incorrect entry for a sun4m system:

```
C753002F -> inetboot.axil4m.Solaris_2.5.1
C753002F.AXIL4M -> inetboot.axil4m.Solaris_2.5.1
```

If they are not correct, remove the entry for that particular client in this directory and again add the client with the `add_install_client` script or through the Solstice tool.

For a JUMPSTART client: The `Error: access violation` from the server to the client might be an indication that the wrong kernel architecture has been specified in the `add_install_client` command line. On the server, type these commands:

```
# cd /cdrom/cdrom0/s0
# ./add_install_client host_name correct_architecture
```

The `add_install_client` script cleans out the incorrect architecture and sets up the install server with the correct architecture to boot the client. If a problem arises using `add_install_client`, use `./rm_install_client` and `./add_install_client` with the correct architecture.

All other follow the same path of checking the `/tftpboot` directory.

THE FOLLOWING FILE SYSTEM(S) HAD AN UNEXPECTED INCONSISTENCY:

Cause

At boot time the `/etc/rcS` script runs the `fsck(1M)` command to check the integrity of file systems marked `fsck` in `/etc/vfstab`. If `fsck(1M)` cannot repair a file system automatically, it interrupts the boot procedure and produces this message. When `fsck(1M)` gets into this state, it cannot repair file systems without losing one or more files, so it defers this responsibility to you, the administrator. Data corruption has probably already occurred.

Action

First run `fsck -n` on the file system to see how many and what type of problems exist. Then run `fsck(1M)` again to repair the file system. If you have a backup of the file system, you can generally answer "y" to all the `fsck(1M)` questions. It is a good practice to keep a record of all problematic files and inode numbers for later reference. To run `fsck(1M)` yourself, specify options as recommended by the boot script. For example:

```
# fsck /dev/rdisk/c0t4d0s0
```

Usually, files lost during `fsck(1M)` repair were created just before a crash or power outage, and cannot be recovered. If important files are lost, you can recover them from backup tapes.

If you do not have a backup, ask an expert to run `fsck(1M)` for you.

See Also

For more information, see the section on checking file system integrity in the *System Administration Guide, Volume 1*.

The SCSI bus is hung. Perhaps an external device is turned off.

Cause

This message appears near the beginning of rebooting, immediately after a `Boot device: ...` message. Then, the system hangs. The problem is conflicting SCSI

targets for a non-boot device. Having an external device turned off is unlikely to cause this problem.

Action

For a solution, refer to “Boot device: /iommu/sbus/directory/directory/sd@3,0” on page 56.

See Also

For more information, see the section on halting and booting in the *System Administration Guide, Volume 1*.

THE SYSTEM IS BEING SHUT DOWN NOW !!!

Cause

This message means the system is going down immediately, and it is too late to save any changes.

Action

This message is often preceded by messages telling you that the system is going down in 15 minutes, 10 minutes, and so on. When you see these initial broadcast shutdown messages, save all your work, send any email you are working on, and close your files. Fortunately, `vi(1)` sessions are automatically saved for later recovery, but many other applications have no crash protection mechanism. Data loss is likely.

See Also

For more information on shutting down the system, see the *System Administration Guide, Volume 1*. If you are using AnswerBook online documentation, “halting the system” is a good search string.

The system will be shut down in *int* minutes

Cause

This message from the system `shutdown(1M)` script informs you that the superuser is taking down the system.

Action

Save all changes now or your work will be lost. Write out any files you were changing, send any email messages you were composing, and close your files.

See Also

For more information on shutting down the system, see the *System Administration Guide, Volume 1*. If you are using AnswerBook online documentation, "halting the system" is a good search string.

This gateway does not support Unix Password.

Cause

While using Firewall v2.0, the following sequence happens:

```
# telnet firewall-machine
Trying 192.29.174.60 ...
Connected to firewall-machine
Escape character is '^]'.
CheckPoint FireWall-1 authenticated Telnet server running on
firewall-machine
Login: testuser
This gateway does not support Unix Password.
```

Action

Under Network Objects, edit the Gateway object Host Properties Auth Schemes and select UNIX Password. UNIX Password is not checked by default as it is considered an unsecure method of authentication.

This mail file has been changed by another mail reader.

Cause

This message appears in a pop-up dialog box whenever you start *mailtool(1)* while another mail reader has the inbox locked. A question follows: Do you wish to ask that mail reader to save the changes? You are given three choices.

Action

If you choose `Save Changes`, *mailtool(1)* requests the other mail reader to relinquish its lock and write out any changes it has made to your inbox. If you choose `Ignore`, *mailtool(1)* reads your inbox without locking it. If you choose `Cancel`, *mailtool(1)* exits.

Timeout waiting for ARP/RARP packet

Cause

This problem can occur while booting from the net, and indicates a network connection problem.

Action

Make sure the Ethernet cable is connected to the network. Check that this system has an entry in the NIS `ethers(4)` map or locally on the boot server. Then check the IP address of the server and the client to make sure they are on the same subnet. Local `/etc/hosts` files must agree with one another and with the NIS `hosts(4)` map.

If those conditions are not causing the problem, go to the system's PROM monitor `ok` prompt and run `test net` to test the network connection. (On older PROM monitors, use `test-net` instead.) If the network test fails, check the Ethernet port, card, fuse, and cable, replacing them if necessary. Also check the twisted pair port to make sure it is patched to the correct subnet.

See Also

For more information on packets, see *SPARC: Installing Solaris Software*. If you are using AnswerBook online documentation, "ARP/RARP" is a good search string.

Timer expired

Cause

The timer set for a STREAMS `ioctl` call has expired. The cause of this error is device specific and could indicate either a hardware or software failure, or perhaps a time-out value that is too short for the specific operation. The status of the `ioctl(2)` operation is indeterminate. This is also returned in the case of `_lwp_cond_timedwait(2)` or `cond_timedwait(3THR)`.

Technical Notes

The symbolic name for this error is `ETIME`, `errno=62`.

token ring hangs

Cause

4.1.3C Sbus cards suffered a system freeze.

Too many links

Cause

An attempt was made to create more than the maximum number of hard links (`LINK_MAX`, by default 32767) to a file. Because each subdirectory is a link to its parent directory, the same error results from trying to create too many subdirectories.

Action

Check why the file has so many links to it. To get more than the maximum number of hard links, use symbolic links instead.

Technical Notes

The symbolic name for this error is `EMLINK`, `errno=31`.

Too many open files

Cause

A process has too many files open at once. The system imposes a per-process soft limit on open files, `OPEN_MAX` (usually 64), which can be increased, and a per-process hard limit (usually 1024), which cannot be increased.

Action

You can control the soft limit from the shell. In the C shell, use the `limit(1)` command to increase the number of descriptors. In the Bourne or Korn shells, use the `ulimit -n` command to increase the number of file descriptors.

If the window system refuses to start new applications because of this error, increase the open-file limit in your login shell before starting the window system.

Technical Notes

The symbolic name for this error is `EMFILE`, `errno=24`.

Transport endpoint is already connected

Cause

A connect request was made on an already connected transport endpoint; or, a `sendto(3XNET)` or `sendmsg(3XNET)` transport endpoint specified a destination when already connected.

Technical Notes

The symbolic name for this error is `EISCONN`, `errno=133`.

Transport endpoint is not connected

Cause

A request to send or receive data was disallowed because the transport endpoint is not connected and (when sending a datagram) no address was supplied.

Technical Notes

The symbolic name for this error is `ENOTCONN`, `errno=134`.

TRAP 3E

Cause

The Ultra system fails to boot with `TRAP 3E`. The system sometimes also displays bad magic number errors.

This error is caused by a bad super block on the boot disk. Which, in turn, could have been caused by a SCSI configuration problem.

Action

To fix:

1. Check the SCSI bus for illegal configuration, bad cables, and duplicate SCSI addresses.
2. Boot from CD-ROM as single user.

```
OK boot cdrom -sw
```

3. Attempt to `fsck(1M)` boot disk. This could fail with a super block error.

```
# fsck /dev/rdisk/device
```

4. Find the locations of alternate super blocks. BE SURE TO USE AN UPPERCASE -N. For example:

```
# newfs -N /dev/rdisk/c0t0d0s0  
/dev/rdisk/c0t0d0s0: 2048960 sectors in 1348 cylinders of 19 tracks,
```

(continued)

(Continuation)

```
80 sectors 1000.5MB in 85 cyl groups (16 c/g, 11.88MB/g, 5696 i/g)
super-block backups (for fsck -F ufs -o b=#) at:
32, 24432, 48832, 73232, 97632, 122032, 146432, 170832, 195232, 219632,
244032, 268432, 292832, 317232, 341632, 366032, 390432, 414832, 439232,
463632, 488032, 512432, 536832, 561232, 585632, 610032, 634432, 658832,
683232, 707632, 732032, 756432, 778272, 802672, 827072, 851472, 875872,
900272, 924672, 949072, 973472, 997872, 1022272, 1290672, ...
```

5. Using an alternate super block, run `fsck(1M)` on the disk. You might have to try more than one alternate super block to make this to work. Pick a couple from the beginning, the middle, and the end.

```
# fsck -o b=<altblk> /dev/rdisk/c0t0d0s0
```

6. The boot block is probably bad too. Restore it while you are booted from the CD-ROM.

```
# /usr/sbin/installboot /usr/platform/architecture/lib/fs/ufs/bootblk
/dev/rdisk/c0t0d0s0
```

7. Reboot the operating environment.

```
# reboot
```

"U"

ufsdump 4mm commands

Cause

Dump syntax was used with autoloader.

umount: warning: */string* not in mnttab

Cause

This message occurs when the superuser attempts to unmount a file system that is not mounted. Subdirectories of file systems, such as `/var`, cannot be unmounted.

Action

Run the `mount(1M)` or `df(1M)` command to see which file systems are mounted. If you really want to unmount one of them, specify the existing mount point.

Unable to connect to license server. Inconsistent encryption code.

Cause

The user receives this error message, and only the IP address of the machine has changed.

Action

The IP address defined with `ifconfig(1M)` must match that in `/etc/hosts`. That is, if you change the machine's IP address with `ifconfig(1M)`, you must also change the machine's entry in the `/etc/hosts` file.

For machines with multiple interfaces, you must check and possibly update `/etc/hostname.*`.

unable to get pty!

Cause

When trying to open a Terminal window (`dtterm`) in CDE, a pop-up window appears stating, `Unable to get pty!`

`dtterm` is not able to open `/dev/pts/int` (where `int` is an integer). The user cannot open this file because `grantpt(3C)` failed to change the permissions on the file.

grantpt(3C) failed because the binary /usr/lib/pt_chmod is not setuid root.
The permissions on /usr/lib/pt_chmod must be 4111.

Action

To restore the correct permissions to pt_chmod, use the following command (as root):

```
# chmod 4111 /usr/lib/pt_chmod
```

Unable to install/attach driver '*string*'

Cause

These messages appear in /var/adm/messages at boot time, when the system tries to load drivers for devices the machine does not have.

Action

This message is strictly informational. You probably do not want all these device drivers because they make your system kernel larger, requiring more memory.

Unable to open nwrecover, Error: nwrecover: NSR: please start a server on *client_name*

Cause

While trying to open the graphical recovery interface by running nwrecover from the client, this error was displayed.

Action

In this case, multiple networker servers existed and nwrecover could not determine which network server to use for the client.

The server can be specified to the nwrecover command with the -s option.

```
nwrecover -c client_name -s server_name
```

`-s server_name` sets the NetWorker server, and `-c client_name` sets the NetWorker client index.

uname: error writing name when booting

Cause

The system cannot bootstrap.

Action

Boot from the CD-ROM and check `/etc/nodename`. The file must contain exactly one line with the name of the system. No blank or other lines are allowed.

undefined control

Cause

This message, prefaced by the file name and line number involved, is from the C preprocessor `/usr/ccs/lib/cpp` and indicates a line starting with a pound-sign (#) but not followed by a valid keyword (such as `define` or `include`).

Action

A piece of software might be running the C preprocessor on an initialization file that you thought was interpreted by a shell. In most shells, the sharp (#) indicates a comment. The C preprocessor considers comments to be anything between `/*` and `*/` delimiters.

unknown host exception: unknown host

Cause

The user tries to install Sun Directory Services 1.0 using the Java front end. During the installation, an error occurs: `unknown host exception: unknown host`. Then the Services displays the host name with domain name appended twice.

Action

The user had the following line in `/etc/nsswitch.conf`: `hosts: dns files`.

By changing the line to point first to files and then to DNS, `hosts: files dns`, the problem was resolved.

Other considerations: This error could also happen if you are using a fully qualified host name. Make sure your host name does not have the domain appended. If you use a fully qualified host name, the domain is appended twice. Also, verify that the domain name specified in `/etc/resolv.conf` is a reachable domain.

Unmatched ‘

Cause

This message from the C shell `csh(1)` indicates that a user typed a command containing a backquote symbol (```) without a closing backquote. Similar messages occur from an unmatched single quote (`'`) or an unmatched double quote (`"`). Other shells generally give a continuation prompt when a command line contains an unmatched quote symbol.

Action

Correct the command line and try again. To continue typing on another line, give the C shell a backslash right before the newline.

UNREF FILE I=i OWNER=o MODE=m SIZE=s MTIME=t CLEAR?

Cause

During phase 4, `fsck(1M)` discovered that the specified file was orphaned because the inode had no record of its path name. In other words, the file was not connected with any directory.

Action

Answer "yes" to reconnect the file into the `lost+found` directory. Then contact the file's owner to ask if you should send it back, and where to place it.

See Also

For more information, see the chapter on checking file system integrity in the *System Administration Guide, Volume 1*.

UnsatisfiedLinkError

Cause

A user was able to use a demo version only when dialed-in to an Internet provider. The user further noted that this Java error message occurred when trying to load library pages without a connection.

Action

The Java WorkShop package relies on the Java Development Kit to provide networking services. There could be two possible problems:

1. The JDK/VM tries to load `net.dll`, which then loads `wsock32.dll` as its socket services. The winsocket program might have done something with the system socket DLLs and might have broken the JDK `net.dll`, which could explain the `UnsatisfiedLinkError`.
2. When JDK creates a `ServerSocket` or `Socket` object, it tries to resolve the local host name by calling `gethostbyaddr()`, which eventually queries the DNS on the Win95/NT, if the user has a DNS entry configured for the TCP/IP. (This normally results in a "Dialup dialog" coming up.)

For the first problem: If the winsocket program renames/moves the `wsock32.dll` or `wsock.dll`, the resolution includes modifying the JDK.

For the second problem: To avoid the DNS query, add an entry to your `%WinDir%\HOSTS` file. Refer to the Java WorkShop release notes for more details.

Use "logout" to logout.

Cause

This C shell message might come as a surprise to Bourne or Korn shell users accustomed to logging out with a Control-D.

Action

When `ignoreeof` is set, the C shell requires users to log out by typing `logout(1)` or `exit(1)`. Write any modified files to disk before exiting.

user unknown

Cause

When trying to mail to a user, the error `Username... User unknown` is displayed. The user is on the same system.

Action

Check for a typographical error in the entered email address. Otherwise, the user could be aliased to a nonexistent email address in `/etc/mail/aliases` or the user's `.mailrc` file.

You cannot mail to a user that has capital letters in its name. `sendmail(1M)` converts all the capital letters to lowercase before attempting to find the user. Because UNIX is case sensitive, it finds no user name on the system with all lowercase letters, so it displays the `User unknown` message.

As a workaround, make sure all user names are composed of *only* lowercase letters.

`/usr/dt/bin/rpc.ttdbserverd:Child Status'` changed

Cause

While running CDE, the error in the console or `/var/adm/messages` file was as follows:

```
Oct 19 04:41:00 darkcastle last message repeated 393 times
Oct 19 04:41:01 darkcastle inetd[120]: /usr/dt/bin/
rpc.ttdbserverd:Child Status Changed
```

Action

Create the following soft links:

```
ln -s /usr/openwin/bin/rpc.ttdbserver /usr/dt/bin/rpc.ttdbserver
ln -s /usr/openwin/bin/rpc.ttdbserverd /usr/dt/bin/rpc.ttdbserverd
```

`/usr/openwin/bin/xinit: connection to X server lost`

Cause

This error means that the *xinit(1)* program, which sets up X11 resources and starts a window manager, failed to locate the X server process. Perhaps the user interrupted window system startup, or exited abnormally from OpenWindows (for example, by killing processes or by rebooting). The X server might have crashed. Data loss is possible in some cases. Depending on the process timing, this message might be normal when the OpenWindows environment exits during a system reboot.

Action

The only solution is to exit and restart the OpenWindows environment. You do not need to reboot the system unless it hangs and fails to give you a console prompt. To exit the OpenWindows environment, select Workspace->Exit. To restart the OpenWindows environment, type *openwin(1)* at the system prompt.

`/usr/ucb/cc: language optional software package not installed`

Cause

While compiling some code for BSD compatibility, the error occurred after invoking `usr/ucb/cc`. The unbundled compiler SPARCworks Professional C product was installed in `/opt`.

`/usr/ucb/cc` is a script that checks for the file `/usr/ccs/bin/ucbcc` and, if found, invokes it with appropriate library flags for BSD-compatibility compilation.

`/usr/ucb/cc` is part of the package `SUNWscpu`. `/usr/ccs/bin/ucbcc` is supposed to be a symbolic link to `/opt/SUNWspro/bin/acc`, which is created during the installation of the unbundled C compiler, `SPROcc`.

Action

Verify that you have the essential OS-bundled Developer packages, `SUNWscpu`, `SUNWbtool`, and the unbundled C compiler, `SPROcc`. However, in this case, `/usr/ccs/bin/ucbcc` was missing on the user's system. Evidently, somehow this link was removed.

Solve the problem by creating a new symbolic link:

```
# ln -s /opt/SUNWspro/bin/acc /usr/ccs/bin/ucbcc
```

Invoke `usr/ucb/cc` to verify this remedy worked.

The following commands are used to identify which packages contain the particular components involved:

```
craterlake% grep ucb/cc /var/sadm/install/contents
/usr/ucb/cc f none 0555 bin bin 3084 50323 814621113 *SUNWscpu
craterlake% ls -l /usr/ucb/cc
-r-xr-xr-x  1 bin      bin           3084 Oct 25  1995 /usr/ucb/cc
craterlake% file !$
file /usr/ucb/cc
/usr/ucb/cc:  executable /usr/bin/sh script
craterlake% grep ucbbc /var/sadm/install/contents
/usr/ccs/bin/ucbcc=/opt1/40/SUNWspro/SC4.0/bin/acc s none SPROcc SPROcc.2 SPROcc.5
craterlake% file /usr/ccs/bin/ucbcc
/usr/ccs/bin/
ucbcc:  ELF 32-bit MSB executable SPARC Version 1, dynamically linked, stripped
craterlake% ls -l /usr/ccs/bin/ucbcc
lrwxrwxrwx  1 root    other           31 Aug 23  1996 /usr/ccs/bin/ucbcc
              -> /opt1/40/SUNWspro/SC4.0/bin/acc
```

UX: userdel: error: Cannot update system files login cannot be deleted

Cause

This error is displayed when using `userdel` to delete a user,

```
userdel -r userid
```

and the `root (/)` file system is full.

Action

Free up some space on the `root (/)` file system.

"V"

Value too large for defined data type

Cause

The user ID or group ID of an IPC object or file system object was too large to be stored in an appropriate member of the caller-provided structure.

Action

Run the application on a newer system, or ask the program's author to fix this condition.

Technical Notes

This error occurs only on systems that support a larger range of user or group ID values than a declared member structure can support. This condition usually occurs because the IPC or file system object resides on a remote machine with a larger value of type `uid_t`, `off_t`, or `gid_t` than that of the local system.

The symbolic name for this error is `E_OVERFLOW`, `errno=79`.

Volume Manager reports error: Configuration daemon can't speak protocol version

Cause

While attempting to run `vxva` (the volume manager GUI) with an upgrade from VXVM 2.0 or 2.1 to VXVM 2.3, you receive this message:

```
Volume Manager reports error:
Configuration daemon can't speak protocol version
```

This message indicates that there is a version mismatch between the version of the volume manager daemon, `vxconfigd`, and the GUI, `vxva`, that you are trying to run. For example, you are running the 2.3 version of `vxconfigd`, and trying to run an old (2.1) version of `vxva`.

Most likely you are using the wrong path for `vxva`. For versions 2.1 and below of `vxva`, the binary can be found in `/opt/vxva/bin`; but starting with 2.1.1, the location was changed to `/opt/SUNWvxva/bin`.

If you did not remove the old `SUNWvxva` package before installing the new 2.3 version (which is normal, since you do not NEED to remove the old package), you probably still have the old `/opt/vxva/bin` in your `$PATH`, and, thus, you are attempting to run the older version of `vxva`.

Action

Run the newer `vxva` program: `/opt/SUNWvxva/bin/vxva`. If that remedy does work and you do not get the error message, remove `/opt/vxva/bin/vxva` from your path statement or remove the old version of `vxva` and create a symbolic link to the new version with the following two commands:

```
# rm /opt/vxva/bin/vxva
# ln -s /opt/SUNWvxva/bin/vxva /opt/vxva/bin/vxva
```

Volume too large for defined data type

Cause

This error occurred when trying to open a database file that was greater than 2 Gbytes in size. You should be able to do this, because the Solaris 2.6 release supports file sizes greater than 2 Gbytes.

Action

It is true that the Solaris 2.6 software supports file sizes greater than 2 Gbytes, but to open a file of that size, you must use a new version of the standard calls. There are 64-bit versions of most system calls and `libc` functions. For example: `open64` instead of `open`.

See Also

Refer to the `lf64(5)` man page.

vxconfigd error: segmentation fault

Cause

When the system boots, the `vxconfigd` fails to start. It fails with a segmentation fault (core dump).

```
vxconfigd error: segmentation fault
[ vxvm warning: _illegal vminor encountered ]
```

Action

Check the date on the system using `date(1)` (`/bin/date` or `/usr/bin/date`). If the date on the system is old (like 1970) or far out in the future (like 2010), `vxconfigd` core dumps.

Change the date on the system using `/bin/date` or `/usr/bin/date` and `vxconfigd` starts without problems.

vxfs filesystems not mounting

Cause

In this case, the user was unable to mount and was getting uncorrectable error messages from mountall. Below is the individual mount report:

```
mount: You don't have a license to run this program
```

However, vxserial -p showed the following:

```
Feature name: CURRSET [95]
Number of licenses: 1 (non-floating)
Expiration date: Sun Jan 18 03:00:00 1998 (22.8 days from now)
Release Level: 20
Machine Class: All

Feature name: RAID [96]
Number of licenses: 1 (non-floating)
Expiration date: Sun Jan 18 03:00:00 1998 (22.8 days from now)
Release Level: 20
Machine Class: All
```

Action

Use vxfsserial -p to see the state of the vxfs license. In this case, it had expired. Unexpired vxfsserial -p output looks similar to the following:

```
Feature name: VXFS [80]
Number of licenses: 1 (non-floating)
Expiration date: No expiration date
Release Level: 22
Machine Class: 934986342
```

vxvm:vx slicer:ERROR unsupported disk layout

Cause

When trying to encapsulate a disk you receive this error.

Action

You must meet the minimum requirements to encapsulate a disk:

1. You must have two free, zero-length, slices on the disk (no cylinders should be assigned to these slices).
2. You must have two free cylinders on the disk. These two cylinders must not be in use by any slice other than slice two.
3. The two free cylinders must be located at the beginning or end of the drive.

"W"

WARNING: add_spec: No major number for sf

Cause

The system prints the following warning message while booting:

```
SunOS Release 5.5.1 Version Generic_103640-03 [UNIX(R)
System V Release 4.0]
Copyright (c) 1983-1996, Sun Microsystems, Inc.
WARNING: add_spec: No major number for sf
```

The `sf(7D)` driver is specific for a Sun Enterprise Network Array (SENA), also known as a "photon."

Action

If no SENA is attached to the system, the message can be safely ignored. To stop seeing the message, comment out the last line in `/kernel/drv/ssd.conf` that references `sf(7D)`.

If you do this, and then later attach a SENA to your system, remember to uncomment this line again.

warning:cachefs:invalid cache version

Cause

While running the Solaris 2.5.1 release and using Adminsuite2.3/Autoclient2.1, the user added 5 autoclients. During startup of the clients, the user received this error message.

Action

The `/kernel/fs/cachefs` files between server and client are different versions. Cachefs versions on the server and the client should be the same as shown in the following:

On the server:

```
# cd /kernel/fs
# ls -al cachefs
-rwxr-xr-x  1 root      sys          229396 Jul 15  1997 cachefs*
```

On the client:

```
# cd /export/root/clientname/kernel/fs
# ls -al cachefs
-rwxr-xr-x  1 root      sys          229396 Jul 15  1997 cachefs*
solution: load patch 104849--02 or higher
```

To solve the problem, load patch 104849--02 or higher.

WARNING: Clock gained *int* days– CHECK AND RESET THE DATE!

Cause

Each workstation contains an internal clock powered by a rechargeable battery. After the system is halted and turned off, the internal clock continues to keep time. When the system is powered on and reboots, the system notices that the internal clock has gained time since the workstation was halted.

Action

In most cases, especially if the power has been off for less than a month, the internal clock keeps the correct time, and you do not have to reset the date. Use the `date(1)` command to check the date and time on your system. If the date or time is wrong, become superuser and use the `date(1)` command to reset them.

Warning: Could not find matching rule in rules.ok

Cause

After an upgrade to the Solaris 2.5.1 release, jumpstart fails with this message:

```
Checking rules.ok file...
Warning: Could not find matching rule in rules.ok
```

This message can occur even if the rules file is known to work, or, after review, it appears to be fine, and the check script has been run.

Action

Remove the rule keyword, `network`, from the rule file and run the check, again. Jumpstart should run without error.

WARNING: FAN FAILURE check if fans are still spinning

Cause

A SPARCcenter™ 2000/2000E might get one of these error messages, `WARNING: FAN FAILURE check if fans are still spinning` or `WARNING: FAN FAILURE still sensed`, displayed on the console screen at any time, with a record of the event in `/var/adm/messages`.

Action

The error itself is descriptive and self-explanatory, and you might suspect that a hardware problem occurred with the system's blower or fan assembly located at the top-most rear of the system cabinet.

Upon further investigation you note that the blower is indeed spinning at a good rate. Given that, you should then check to see if the "AC Dist to Blower to Filter to Keyswitch Harness" plug/adaptor is plugged in correctly. Two cable assemblies connect the blower assembly to the unit's power supply. One is the "power supply" cable and the other is the "AC Dist to Blower to Filter to Keyswitch Harness."

Once the harness is securely connected, you see another message, NOTICE: FAN RECOVERED, logged on the system's console screen, or, if missed, it is in `/var/adm/messages`.

WARNING: FAN FAILURE still sensed

Refer to "WARNING: FAN FAILURE check if fans are still spinning" on page 257.

WARNING: No network locking on *string*: contact admin to install server change

Cause

The `mount(1M)` command issues this message whenever it mounts a file system that does not have NFS locking, such as a standard SunOS 4.1 exported file system. Data loss is possible in applications that depend on locking.

Action

On the remote SunOS 4.1 system, install the appropriate `rpc.lockd` jumbo patch to implement NFS locking. For the SunOS 4.1.4 system, install patch #102264; for the SunOS 4.1.3 system, install patch #100075; for earlier 4.1 releases, install patch #101817.

WARNING: processor level 4 interrupt not serviced

Cause

This message is basically a diagnostic from the SCSI driver. It can appear on the console every 10 minutes or so.

Action

To reduce the frequency of this message, add this line near the bottom of the `/etc/system` file and reboot:

```
set esp:esp_use_poll_loop=0
```

Technical Notes

You might also see this message repeatedly after manually removing a CD when it was busy. Do not do this! To return the system to normal, reboot the system with the `-r` (reconfigure) option.

WARNING: /tmp: File system full, swap space limit exceeded

Cause

The system swap area (virtual memory) has filled up. You need to reduce swap space consumption by killing some processes or possibly by rebooting the system.

Action

For information about increasing swap space, refer to “Not enough space” on page 183.

WARNING: TOD clock not initialized– CHECK AND RESET THE DATE!

Cause

This message indicates that the Time Of Day (TOD) clock reads zero, so its time is the beginning of the UNIX epoch: midnight, 31 December 1969. On a brand-new system, the manufacturer might have neglected to initialize the system clock. On older systems it is more likely that the rechargeable battery has run out and requires replacement.

Action

First replace the battery according to the manufacturer's instructions. Then become superuser and use the `date(1)` command to set the time and date. On some systems the clock is powered by the same battery as the NVRAM, so a dead battery also causes loss of the machine's Ethernet address and host ID, which are more serious problems for networked systems.

WARNING: Unable to repair the / filesystem. Run fsck

Cause

This message comes at boot time from the `/etc/rcS` script whenever it gets a bad return code from `fsck(1M)` after checking a file system. The message recommends an `fsck(1M)` command line, and instructs you to exit the shell when done to continue booting. Then the script places the system in single-user mode so `fsck(1M)` can be run effectively.

Action

For information about repairing UFS file systems, refer to “`/dev/rdisk/string: UNEXPECTED INCONSISTENCY; RUN fsck MANUALLY.`” on page 88.

For information about repairing non-UFS file systems, refer to “THE FOLLOWING FILE SYSTEM(S) HAD AN UNEXPECTED INCONSISTENCY:” on page 235.

WARNING: vxvm:vxio: Illegal vminor encountered

Cause

In this case, the message occurred during booting. The system was sharing an SSA1XX with an identical system. The user was also getting an error in disk group configuration copies during booting. The identical system was booting up fine—without error messages. vxconfigd died. A vxprivutil scan of one of the disks indicated the following:

```
diskid: 880409237.1043.system_that_comes_up
hostid: none
```

Action

The user quickly applied a vxinstall on both systems: first, on the system that did not successfully boot, and then on the system that did. The user had to run a custom vxinstall, selecting only the disks desired for each system.

Technical Notes

Note - The following attempt to resolve the problem failed.

```
vxiod set 10
vxconfigd -m disable
vxdctl init hostname
vxdctl enable
```

Watchdog Reset

Cause

This fatal error usually indicates some kind of hardware problem. Data corruption on the system is possible.

Action

Look for some other message that might help diagnose the problem. By itself, a watchdog reset does not provide enough information; because traps are disabled, all information has been lost. If all that appears on the console is an `ok` prompt, issue the following PROM command to view the final messages that occurred just before system failure:

```
ok f8002010 wector p
```

Yes, that word is `wector`, not `vector`.

The result is a display of messages similar to those produced by the `dmesg(1M)` command. These messages can be useful in finding the cause of system failure.

Technical Notes

This message does not come from the kernel, but from the OpenBoot PROM monitor, a piece of Forth software that gives you the `ok` prompt before you boot UNIX. If the CPU detects a trap when traps are disabled (an unrecoverable error), it signals a watchdog. The OpenBoot PROM monitor detects the watchdog, issues this message, and shuts down the system.

Who are you?

Cause

Many networking programs can print this message, including `from(1B)`, `lpr(1B)`, `lprm(1B)`, `mailx(1)`, `rdist(1)`, `sendmail(1M)`, `talk(1)`, and `rsh(1)`. The command prints this message when it cannot locate a password file entry for the current user. This error might occur if a user logged in just before the superuser deleted that user's password entry, or if the network naming service fails for a user who has no entry in the local password file.

Action

If a user's password file entry was accidentally deleted, restore it from backups or from another password file. If a user's login name or user ID was changed, ask that user to log out and log in again. If the network naming service failed, check the NIS server(s) and repair or reboot as necessary.

Technical Notes

A known problem exists with starting hundreds of `rsh(1)` processes on another machine. This message appears because `rsh(1)` hangs while binding to a reserved port and responds too slowly to interact with the network naming service.

Window Underflow

Cause

This message often occurs at boot time, sometimes along with a `Watchdog Reset` error. It comes from the OpenBoot PROM monitor, which was passed a processor trap from the hardware. This error indicates that some program tried to access a register window that was not accessible from the processor.

Action

On some system architectures the problem could be that different capacity memory chips are mixed together. Someone might have placed 1-Mbyte SIMMs in the same bank with 4-Mbyte SIMMs. If this is so, rearrange the memory chips. Make sure to put higher-capacity SIMMs in the first bank(s), and lower-capacity SIMMs in the remaining bank(s); never mix different capacity SIMMs in the same bank.

The problem could also be that cache memory on the motherboard has gone bad and needs replacement. If main memory is installed correctly, try swapping the motherboard.

Technical Notes

The best way to isolate the problem is to look at the `%PC` register to see where it got its arguments, and why the arguments were bad. If you can reproduce the condition causing this message, your system vendor might be able to help diagnose the problem.

"X"

X connection to *string:0.0* broken (explicit kill or server shutdown).

Cause

This error means that the client has lost its connection to the X server. The "0.0" represents the display device, which is usually the console. This message can appear when a user is running an X application on a remote system with the DISPLAY set back to the original system and the remote system's X server disappears, perhaps because someone exited X windows or rebooted the machine. It sometimes appears locally when a user exits the window system. Data loss is possible if applications were killed before saving files.

Action

Try to run the application again in a few minutes after the system has rebooted and the window system is running.

xinit: not found

Cause

The OpenWindows environment was probably not installed properly, and the *openwin(1)* program could not find *xinit(1)* to start the X windows system. If the user is running another version of X windows, such as the MIT X11 distribution, the *startx* program serves the same function as *xinit(1)*.

Action

Check the PATH environment variable to make sure it contains the appropriate X windows install directory. Verify that *xinit(1)* is in this directory as an executable program.

XIO: fatal IO error 32 (Broken pipe) on X server "string:0.0"

Cause

This error means that I/O with the X server has been broken. The 0.0 represents the display device, which is usually the console. This message can appear when a user is running Display PostScript applications and the X server disappears or the client is shut down. Data loss is possible, if applications disappeared before saving files.

Action

Try to run the application again in a few minutes after the system has rebooted and the window system is running.

Xlib: connection to "string:0.0" refused by server

Cause

This message is immediately followed by the Xlib: Client is not authorized to connect to Server message. These messages indicate that an X windows application tried to run on the X server specified inside double quotes, which did not allow the request. The 0.0 represents the display device, which is usually the console. If no servername appears, the superuser probably tried to run an X application on the current machine in an X session that was owned by somebody else.

Action

To allow this client to connect to the X server, run *xhost(1) +clientname* on the X server system. Only the owner of the current X session (who is not necessarily the superuser) is allowed to run the *xhost(1)* command. If somebody else is running X windows on the server, ask them to log out and then start your own X session on that server; remote X connections are usually allowed for the same user ID.

Xlib: extension "GLX" missing on display "0.0"

Cause

Install the OpenGL[®] 1.0 software and test the configuration by running `/usr/openwin/demo/GL/ogl_install_check`, which provides the following results:

```
# ./ogl_install_check
Xlib: extension "GLX" missing on display "0.0".
Xlib: extension "GLX" missing on display "0.0".
Xlib: extension "GLX" missing on display "0.0".
can't find visual
```

Action

First check that the installation has worked correctly by running the package check utility on the runtime package: `pkgchk SUNWglrt`. This should result in an error message like this:

```
ERROR: /usr/openwin/server/etc/OWconfig
file size <187> expected <5423> actual
file cksum <14394> expected <27045> actual
```

(The numbers might be different, but there should be only one file.) If other errors occur, re-install OpenGL, especially the `SUNWglrt` package.

Assuming that is fine, look at the process owner for the Xsun process using the following:

```
# ps -aef | grep Xsun | grep -v grep
nobody 20022 225 0 11:36:22 ? 0:34 /usr/openwin/bin/
Xsun :0 -nobanner
```

If the owner is not `root`, that is most likely the problem. There is a permissions issue loading the graphic pipelines.

If you are using CDE, ensure that the Xservers file has this form:

```
:0 Local local_uid@console root /usr/openwin/bin/Xsun :0 -nobanner
```

The Xservers file can be found in `/usr/dt/config`, if you have not done any customization. Otherwise, it can more than likely be found in `/etc/dt/config/`. Additional arguments after the `-nobanner` option are acceptable.

Another way of proving this is to run the OpenWindows environment from the command line as root. It ensures that the Xsun process is owned by root.

Another possibility is that the system is NOT a Creator 3D. You can only run OpenGL 1.0 on an Ultra machine with a Creator 3D graphics card. If you install this application on an Ultra machine with a Creator framebuffer and NOT a Creator 3D, you see these same error messages.

xntpd: clnt_dg_create: out of memory

Cause

At boot time, the error occurs after configuring NTP. Except for the error, everything seems to be working properly.

Action

As a workaround, move the script for `xntpd` from `S74xntpd` to `S77xntpd`, so it starts after `S76nsd`.

xterm: fatal IO error 32 (Broken Pipe) or KillClient on X server "*string:0.0*"

Cause

This error means that `xterm(1)` has lost its connection to the X server. The `0.0` represents the display device, which is usually the console. This message can appear when a user is running `xterm` and the X server disappears or the client is shut down. Data loss is possible if applications were killed before saving files.

Action

Try to run the terminal emulator again in a few minutes after the system has rebooted and the window system is running.

XView warning: Cannot load font set '*string*' (Font Package)

Cause

This message from the XView library warns that a requested font is not installed on the X server. Often multiple warnings are displayed for the same font. The set of available fonts can vary from release to release.

Action

To see which fonts are available on the X server, run the *xlsfonts(1)* program. Then specify another font name that you see in the output of *xlsfonts(1)*. Sometimes you can locate a similar font from a different vendor.

Technical Notes

Two packages of X windows fonts are: the common but not required fonts (*SUNWxwcfnt*), and the optional fonts (*SUNWxwoft*). Run *pkginfo(1)* to see if both packages are installed, and add them to the system as you desire.

"Y"

yp_all RPC clnt_call (transport level) failure

Cause

At random times, a slave NIS server has a problem that causes *ypbind(1M)* to report *ypserver* not responding, and the machine must be rebooted. The *syslog* file contains the following:

```
Dec 14 07:11:03 rahab syslog: yp_all -  
RPC clnt_call (transport level) failure:  
RPC: Unable to receive; An event requires attention
```

Action

As a workaround, increase the file descriptor limit in the `yp` startup script, `/etc/rc2.d/S71rpc`. Add this command to the script before `ypserv` is started:

```
ulimit -n 256
```

`ypbind[int]: NIS server for domain "string" OK`

Cause

This message appears after an NIS server not responding message to indicate that `ypbind(1M)` is able to communicate with an NIS server again.

Action

Proceed with your work. This message is strictly informational.

`ypbind[int]: NIS server not responding for domain "string"; still trying`

Cause

This means that the NIS client daemon `ypbind(1M)` cannot communicate with an NIS server for the specified domain. This message appears when a workstation running the NIS naming service has become disconnected from the network, or when NIS servers are down or extremely slow to respond.

Action

If other NIS clients are behaving normally, check the Ethernet cabling on the workstation that is getting this message. Note the following differences between architectures:

- On SPARC machines, disconnected network cabling also produces a series of `no carrier` messages.
- On IA machines, the NIS+ messages might be the only indication that network cabling is disconnected.

If many NIS clients on the network are giving this message, go to the NIS server in question and reboot or repair as necessary. To locate the NIS server for a domain, run the `ypwhich(1)` command. When the server machine returns to operation, NIS clients give an NIS server for domain OK message.

See Also

For more information about `ypbind(1M)`, see the section on administering secure NFS in the *System Administration Guide, Volume 3*.

`ypserv[int]: restarting resolv server. old one not responding`

Cause

In this instance, the NIS Server, which had been upgraded from version 2.5.1 to version 2.6, was repeating this error message every ten minutes. Also, the Server was less frequently repeating the following message:

```
rpc.nisd_resolv[7472]: svc_getreqset: no transport handle for fd2
```

The SUNWypu and SUNWypv packages had been installed.

Action

Install Patch-ID# 105552-01. Also, set `B=` in the Makefile. Run `make` again to recreate the maps on the following:

```
#B=-b  
B=
```

You might also need to remove the `-d` option from the `ypserv` command in the `/usr/lib/netsvc/yp/ypstart` script. Then, you must reboot the machine.

ypwhich: can't communicate with ypbind

Cause

This message from the `ypwhich(1)` command indicates that the NIS binder process `ypbind(1M)` is not running on the local machine.

Action

If the system is not configured to use NIS, this message is normal and expected. Configure the system to use NIS if necessary.

If the system is configured to use NIS, but the `ypbind(1M)` process is not running, invoke the following command to start it up:

```
# /usr/lib/netsvc/yp/ypbind -broadcast
```

"Z"

zsint: silo overflow

Cause

This message means that the Zilog 8530 character input silo (or serial port FIFO) overflowed before it could be serviced. The `zs(7D)` driver, which talks to a Zilog Z8530 chip, is reporting that the FIFO (holding about two characters) has been overrun. The number after `zs(7D)` shows which serial port experienced an overflow:

```
zs0 - tty serial port 0 (/dev/ttya)
zs1 - tty serial port 1 (/dev/ttyb)
zs2 - keyboard port (/dev/kbd)
zs3 - mouse port (/dev/mouse)
```

Action

Silo overflows indicate that data in the respective serial port FIFO have been lost. However, the consequences of silo overflows might be negligible if the overflows occur infrequently, if data loss is not catastrophic, or if data can be recovered or reproduced. For example, although a silo overflow on the mouse driver (`zs3`) indicates that the system could not process mouse events quickly enough, the user can perform mouse motions again. Similarly, lost data from a silo overflow on a serial port with a modem connection transferring data using `uucp(1C)` is recovered when `uucp(1C)` discovers the loss of data and requests retransmission of the corrupted packet.

Frequent silo overflow messages can indicate a `zs(7D)` hardware FIFO problem, a serial driver software problem, or abnormal data or system activity. For example, the system ignores interrupts during system panics, so mouse and keyboard activity result in silo overflows.

If the serial ports experiencing silo overflows are not being used, a silo overflow could indicate the onset of a hardware problem.

Technical Notes

Another type of silo overflow is one that occurs during reboot, when an HDLC line is connected to any of the terminal ports. For example, an X.25 network could be sending frames before the kernel has been told to expect them. Such overflow messages can be ignored.