

SunOS Reference Manual

Sun Microsystems, Inc.
2550 Garcia Avenue
Mountain View, CA 94043
U.S.A.



© 1994 Sun Microsystems, Inc. All rights reserved.
2550 Garcia Avenue, Mountain View, California 94043-1100 U.S.A.

This product and related documentation are protected by copyright and distributed under licenses restricting its use, copying, distribution, and decompilation. No part of this product or related documentation may be reproduced in any form by any means without prior written authorization of Sun and its licensors, if any.

Portions of this product may be derived from the UNIX® and Berkeley 4.3 BSD systems, licensed from UNIX Systems Laboratories, Inc., a wholly owned subsidiary of Novell, Inc., and the University of California, respectively. Third-party software, including font technology, in this product is protected by copyright and licensed from Sun's Suppliers.

RESTRICTED RIGHTS LEGEND: Use, duplication, or disclosure by the government is subject to restrictions as set forth in subparagraph (c)(1)(ii) of the Rights in Technical Data and Computer Software clause at DFARS 252.227-7013 and FAR 52.227-19.

This product or the products described herein may be protected by one or more U.S., foreign patents, or pending applications.

TRADEMARKS

Sun, Sun Microsystems, the Sun Logo, SunSoft, Sun Microsystems Computer Corporation and Solaris, are trademarks or registered trademarks of Sun Microsystems, Inc. in the U.S. and certain other countries. UNIX is a registered trademark of Novell, Inc., in the United States and other countries; X/Open Company, Ltd., is the exclusive licensor of such trademark. OPEN LOOK® is a registered trademark of Novell, Inc. All other product names mentioned herein are the trademarks of their respective owners.

All SPARC trademarks, including the SCD Compliant Logo, are trademarks or registered trademarks of SPARC International, Inc. SPARCstation, SPARCserver, SPARCengine, SPARCstorage, SPARCware, SPARCcenter, SPARCclassic, SPARCcluster, SPARCdesign, SPARC811, SPARCprinter, UltraSPARC, microSPARC, SPARCworks, and SPARCcompiler are licensed exclusively to Sun Microsystems, Inc. Products bearing SPARC trademarks are based upon an architecture developed by Sun Microsystems, Inc.

The OPEN LOOK® and Sun™ Graphical User Interfaces were 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.

X Window System is a product of the Massachusetts Institute of Technology.

THIS PUBLICATION IS PROVIDED "AS IS" WITHOUT WARRANTY OF ANY KIND, EITHER EXPRESS OR IMPLIED, INCLUDING, BUT NOT LIMITED TO, THE IMPLIED WARRANTIES OF MERCHANTABILITY, FITNESS FOR A PARTICULAR PURPOSE, OR NON-INFRINGEMENT.

THIS PUBLICATION COULD INCLUDE TECHNICAL INACCURACIES OR TYPOGRAPHICAL ERRORS. CHANGES ARE PERIODICALLY ADDED TO THE INFORMATION HEREIN; THESE CHANGES WILL BE INCORPORATED IN NEW EDITIONS OF THE PUBLICATION. SUN MICROSYSTEMS, INC. MAY MAKE IMPROVEMENTS AND/OR CHANGES IN THE PRODUCT(S) AND/OR THE PROGRAM(S) DESCRIBED IN THIS PUBLICATION AT ANY TIME.

Portions © AT&T 1983-1990 and reproduced with permission from AT&T.

Preface

OVERVIEW

A man page is provided for both the naive user, and sophisticated user who is familiar with the SunOS operating system and is in need of on-line information. A man page is intended to answer concisely the question “What does it do?” The man pages in general comprise a reference manual. They are not intended to be a tutorial.

The following contains a brief description of each section in the man pages and the information it references:

- Section 1 describes, in alphabetical order, commands available with the operating system.
- Section 1M describes, in alphabetical order, commands that are used chiefly for system maintenance and administration purposes.
- Section 2 describes all of the system calls. Most of these calls have one or more error returns. An error condition is indicated by an otherwise impossible returned value.
- Section 3 describes functions found in various libraries, other than those functions that directly invoke UNIX system primitives, which are described in Section 2 of this volume.

-
- Section 4 outlines the formats of various files. The C structure declarations for the file formats are given where applicable.
 - Section 5 contains miscellaneous documentation such as character set tables, etc.
 - Section 7 describes various special files that refer to specific hardware peripherals, and device drivers. STREAMS software drivers, modules and the STREAMS-generic set of system calls are also described.
 - Section 9 provides reference information needed to write device drivers in the kernel operating systems environment. It describes two device driver interface specifications: the Device Driver Interface (DDI) and the Driver–Kernel Interface (DKI).
 - Section 9E describes the DDI/DKI, DDI-only, and DKI-only entry-point routines a developer may include in a device driver.
 - Section 9F describes the kernel functions available for use by device drivers.
 - Section 9S describes the data structures used by drivers to share information between the driver and the kernel.

Below is a generic format for man pages. The man pages of each manual section generally follow this order, but include only needed headings. For example, if there are no bugs to report, there is no BUGS section. See the intro pages for more information and detail about each section, and **man(1)** for more information about man pages in general.

NAME

This section gives the names of the commands or functions documented, followed by a brief description of what they do.

SYNOPSIS

This section shows the syntax of commands or functions. When a command or file does not exist in the standard path, its full pathname is shown. Literal characters (commands and options) are in **bold** font and variables (arguments, parameters and substitution characters) are in *italic* font. Options and arguments are alphabetized, with single letter arguments first, and options with arguments next, unless a different argument order is required.

The following special characters are used in this section:

- [] The option or argument enclosed in these brackets is optional. If the brackets are omitted, the argument *must* be specified.
- ... Ellipses. Several values may be provided for the previous argument, or the previous argument can be specified multiple times, for example, *'filename ...'*.
- | Separator. Only one of the arguments separated by this character can be specified at time.

PROTOCOL

This section occurs only in subsection 3R to indicate the protocol description file. The protocol specification pathname is always listed in **bold** font.

AVAILABILITY

This section briefly states any limitations on the availability of the command. These limitations could be hardware or software specific.

A specification of a class of hardware platform, such as **x86** or **SPARC**, denotes that the command or interface is applicable for the hardware platform specified.

In Section 1 and Section 1M, **AVAILABILITY** indicates which package contains the command being described on the manual page. In order to use the command, the specified package must have been installed with the operating system. If the package was not installed, see **pkgadd(1)** for information on how to upgrade.

MT-LEVEL

This section lists the **MT-LEVEL** of the library functions described in the Section 3 manual pages. The **MT-LEVEL** defines the libraries' ability to support threads. See **Intro(3)** for more information.

DESCRIPTION

This section defines the functionality and behavior of the service. Thus it describes concisely what the command does. It does not discuss **OPTIONS** or cite **EXAMPLES**. Interactive commands, subcommands, requests, macros, functions and such, are described under **USAGE**.

IOCTLS

This section appears on pages in Section 7 only. Only the device class which supplies appropriate parameters to the **ioctl(2)** system call is called **ioctl** and generates its own heading. IOCTLS for a specific device are listed alphabetically (on the man page for that specific device). IOCTLS are used for a particular class of devices all which have an **io** ending, such as **mtio(7)**.

OPTIONS

This lists the command options with a concise summary of what each option does. The options are listed literally and in the order they appear in the SYNOPSIS section. Possible arguments to options are discussed under the option, and where appropriate, default values are supplied.

RETURN VALUES

If the man page documents functions that return values, this section lists these values and describes the conditions under which they are returned. If a function can return only constant values, such as 0 or -1, these values are listed in tagged paragraphs. Otherwise, a single paragraph describes the return values of each function. Functions declared as **void** do not return values, so they are not discussed in RETURN VALUES.

ERRORS

On failure, most functions place an error code in the global variable **errno** indicating why they failed. This section lists alphabetically all error codes a function can generate and describes the conditions that cause each error. When more than one condition can cause the same error, each condition is described in a separate paragraph under the error code.

USAGE

This section is provided as a *guidance* on use. This section lists special rules, features and commands that require in-depth explanations. The subsections listed below are used to explain built-in functionality:

- Commands**
- Modifiers**
- Variables**
- Expressions**
- Input Grammar**

EXAMPLES

This section provides examples of usage or of how to use a command or function. Wherever possible a complete example including command line entry and machine response is shown. Whenever an example is given, the prompt is shown as

example%

or if the user must be super-user,

example#

Examples are followed by explanations, variable substitution rules, or returned values. Most examples illustrate concepts from the SYNOPSIS, DESCRIPTION, OPTIONS and USAGE sections.

ENVIRONMENT

This section lists any environment variables that the command or function affects, followed by a brief description of the effect.

FILES

This section lists all filenames referred to by the man page, files of interest, and files created or required by commands. Each is followed by a descriptive summary or explanation.

SEE ALSO

This section lists references to other man pages, in-house documentation and outside publications.

DIAGNOSTICS

This section lists diagnostic messages with a brief explanation of the condition causing the error. Messages appear in **bold** font with the exception of variables, which are in *italic* font.

WARNINGS

This section lists warnings about special conditions which could seriously affect your working conditions — this is not a list of diagnostics.

NOTES

This section lists additional information that does not belong anywhere else on the page. It takes the form of an *aside* to the user, covering points of special interest. Critical information is never covered here.

BUGS

This section describes known bugs and wherever possible suggests workarounds.

NAME	Intro, intro – introduction to commands and application programs
AVAILABILITY	This section indicates which package contains the commands being described on this page. To be able to use the command, the indicated package must have been installed with the operating system. For information on how to add a package see pkgadd(1) .
DESCRIPTION	<p>This section describes, in alphabetical order, commands available with this operating system.</p> <p>Pages of special interest are categorized as follows:</p> <p>1B Commands found only in the <i>SunOS/BSD Compatibility Package</i>. Refer to the <i>Solaris Source Compatibility Guide</i> for more information.</p> <p>1C Commands for communicating with other systems.</p> <p>1F Commands associated with <i>Form and Menu Language Interpreter</i> (FMLI).</p> <p>1S Commands specific to the SunOS system.</p>
OTHER SECTIONS	<p>See these sections of the <i>man Pages(1M): System Administration Commands</i> for more information.</p> <ul style="list-style-type: none"> • Section 1M in this manual for system maintenance commands. • Section 4 of this manual for information on file formats. • Section 5 of this manual for descriptions of publicly available files and miscellaneous information pages. • Section 6 in this manual for computer demonstrations. <p>For tutorial information about these commands and procedures, see:</p> <ul style="list-style-type: none"> • <i>Solaris Advanced User's Guide</i> • <i>Programming Utilities Guide</i>
Manual Page Command Syntax	<p>Unless otherwise noted, commands described in the SYNOPSIS section of a manual page accept options and other arguments according to the following syntax and should be interpreted as explained below.</p> <p><i>name</i> [<i>-option...</i>] [<i>cmdarg...</i>]</p> <p>where:</p> <p>[] Surround an <i>option</i> or <i>cmdarg</i> that is not required.</p> <p>... Indicates multiple occurrences of the <i>option</i> or <i>cmdarg</i>.</p> <p><i>name</i> The name of an executable file.</p> <p><i>option</i> (Always preceded by a “-”.)</p> <p> <i>noargletter...</i> or,</p> <p> <i>argletter optarg[,...]</i></p>

<i>noargletter</i>	A single letter representing an option without an option-argument. Note that more than one <i>noargletter</i> option can be grouped after one “-” (Rule 5, below).
<i>argletter</i>	A single letter representing an option requiring an option-argument.
<i>optarg</i>	An option-argument (character string) satisfying a preceding <i>argletter</i> . Note that groups of <i>optargs</i> following an <i>argletter</i> must be separated by commas, or separated by a tab or space character and quoted (Rule 8, below).
<i>cmdarg</i>	Path name (or other command argument) <i>not</i> beginning with “-”, or “-” by itself indicating the standard input.

**Command Syntax
Standard: Rules**

These command syntax rules are not followed by all current commands, but all new commands will obey them. **getopts(1)** should be used by all shell procedures to parse positional parameters and to check for legal options. It supports Rules 3-10 below. The enforcement of the other rules must be done by the command itself.

1. Command names (*name* above) must be between two and nine characters long.
2. Command names must include only lower-case letters and digits.
3. Option names (*option* above) must be one character long.
4. All options must be preceded by “-”.
5. Options with no arguments may be grouped after a single “-”.
6. The first option-argument (*optarg* above) following an option must be preceded by a tab or space character.
7. Option-arguments cannot be optional.
8. Groups of option-arguments following an option must either be separated by commas or separated by tab or space character and quoted (**-o xxx,z,yy** or **-o "xxx z yy"**).
9. All options must precede operands (*cmdarg* above) on the command line.
10. “—” may be used to indicate the end of the options.
11. The order of the options relative to one another should not matter.
12. The relative order of the operands (*cmdarg* above) may affect their significance in ways determined by the command with which they appear.
13. “-” preceded and followed by a space character should only be used to mean standard input.

SEE ALSO

getopts(1), wait(1), exit(2), getopt(3C), wait(3B)

DIAGNOSTICS

Upon termination, each command returns two bytes of status, one supplied by the system and giving the cause for termination, and (in the case of “normal” termination) one supplied by the program [see **wait(3B)** and **exit(2)**]. The former byte is 0 for normal termination; the latter is customarily 0 for successful execution and non-zero to indicate troubles such as erroneous parameters, or bad or inaccessible data. It is called variously “exit code”, “exit status”, or “return code”, and is described only where special conventions are involved.

WARNINGS

Some commands produce unexpected results when processing files containing null characters. These commands often treat text input lines as strings and therefore become confused upon encountering a null character (the string terminator) within a line.

Name	Appears on Page	Description
acctcom	acctcom(1)	search and print process accounting files
adb	adb(1)	general-purpose debugger
addbib	addbib(1)	create or extend a bibliographic database
admin	sccs-admin(1)	create and administer SCCS history files
aedplot	plot(1B)	graphics filters for various plotters
alias	alias(1)	shell built-in functions to create your own pseudonym or shorthand for a command or series of commands
apropos	apropos(1)	locate commands by keyword lookup
ar	ar(1)	maintain portable archive or library
arch	arch(1B)	display the architecture of the current host
as	as(1)	assembler
at	at(1)	execute commands at a later time
atoplot	plot(1B)	graphics filters for various plotters
atq	atq(1)	display the jobs queued to run at specified times
atrm	atrm(1)	remove jobs spooled by at or batch
audioconvert	audioconvert(1)	convert audio file formats
audioplay	audioplay(1)	play audio files
audiorecord	audiorecord(1)	record an audio file
awk	awk(1)	pattern scanning and processing language
banner	banner(1)	make posters
basename	basename(1)	deliver portions of pathnames
basename	basename(1B)	display portions of pathnames
batch	at(1)	execute commands at a later time
bc	bc(1)	arbitrary precision arithmetic language
bdiff	bdiff(1)	big diff
bg	jobs(1)	shell built-in functions to control process execution
bgplot	plot(1B)	graphics filters for various plotters

biff	biff(1B)	give notice of incoming mail messages
break	break(1)	shell built-in functions to escape from or advance within a controlling while, for, foreach, or until loop
cal	cal(1)	display a calendar
calendar	calendar(1)	reminder service
cancel	lp(1)	send/cancel requests to an LP print service
case	case(1)	shell built-in functions to choose from among a list of actions
cat	cat(1)	concatenate and display files
cc	cc(1B)	C compiler
cd	cd(1)	shell built-in functions to change the current working directory
cdc	sccs-cdc(1)	change the delta commentary of an SCCS delta
chdir	cd(1)	shell built-in functions to change the current working directory
checkeq	eqn(1)	typeset mathematics test
checknr	checknr(1)	check nroff and troff input files; report possible errors
chgrp	chgrp(1)	change the group ownership of a file
chkey	chkey(1)	change user's secure RPC key pair
chmod	chmod(1)	change the permissions mode of a file
chown	chown(1)	change owner of file
chown	chown(1B)	change owner
ckdate	ckdate(1)	prompts for and validates a date
ckgid	ckgid(1)	prompts for and validates a group id
ckint	ckint(1)	display a prompt; verify and return an integer value
ckitem	ckitem(1)	build a menu; prompt for and return a menu item
ckkeywd	ckkeywd(1)	prompts for and validates a keyword
ckpath	ckpath(1)	display a prompt; verify and return a pathname
ckrange	ckrange(1)	prompts for and validates an integer
ckstr	ckstr(1)	display a prompt; verify and return a string answer
cktime	cktime(1)	display a prompt; verify and return a time of day
ckuid	ckuid(1)	prompts for and validates a user ID
ckyorn	ckyorn(1)	prompts for and validates yes/no
clear	clear(1)	clear the terminal screen
cmp	cmp(1)	compare two files
cocheck	coproc(1F)	communicate with a process
cocreate	coproc(1F)	communicate with a process
codestroy	coproc(1F)	communicate with a process
cof2elf	cof2elf(1)	COFF to ELF object file translation

col	col(1)	reverse line-feeds filter
comb	scs-comb(1)	combine SCCS deltas
comm	comm(1)	select or reject lines common to two sorted files
compress	compress(1)	compress, uncompress files or display expanded files
continue	break(1)	shell built-in functions to escape from or advance within a controlling while, for, foreach, or until loop
coproc	coproc(1F)	communicate with a process
coreceive	coproc(1F)	communicate with a process
cosend	coproc(1F)	communicate with a process
cp	cp(1)	copy files
cpio	cpio(1)	copy file archives in and out
cpp	cpp(1)	the C language preprocessor
crontab	crontab(1)	user crontab file
crtpplot	plot(1B)	graphics filters for various plotters
crypt	crypt(1)	encode or decode a file
csh	csh(1)	shell command interpreter with a C-like syntax
csplit	csplit(1)	split a file with respect to a given context
ct	ct(1C)	spawn login to a remote terminal
ctags	ctags(1)	create a tags file for use with ex and vi
cu	cu(1C)	call another UNIX system
cut	cut(1)	cut out selected fields of each line of a file
date	date(1)	print and set the date
dc	dc(1)	desk calculator
delta	scs-delta(1)	make a delta to an SCCS file
deroff	deroff(1)	remove nroff/troff, tbl, and eqn constructs
df	df(1B)	display status of disk space on file systems
diff3	diff3(1)	3-way differential file comparison
diff	diff(1)	display line-by-line differences between pairs of text files
diffmk	diffmk(1)	mark differences between versions of a troff input file
dircmp	dircmp(1)	directory comparison
dirname	basename(1)	deliver portions of pathnames
dirs	cd(1)	shell built-in functions to change the current working directory
dis	dis(1)	object code disassembler
disable	enable(1)	enable/disable LP printers
dispgid	dispgid(1)	displays a list of all valid group names
dispuid	dispuid(1)	displays a list of all valid user names
dos2unix	dos2unix(1)	convert text file from DOS format to ISO format
download	download(1)	host resident PostScript font downloader

dpost	dpost(1)	troff postprocessor for PostScript printers
du	du(1B)	display the number of disk blocks used per directory or file
dumbplot	plot(1B)	graphics filters for various plotters
dump	dump(1)	dump selected parts of an object file
dumpcs	dumpcs(1)	show codeset table for the current locale
dumpkeys	loadkeys(1)	load and dump keyboard translation tables
echo	echo(1)	echo arguments
echo	echo(1B)	echo arguments to standard output
echo	echo(1F)	put string on virtual output
ed	ed(1)	text editor
edit	edit(1)	text editor (variant of ex for casual users)
egrep	egrep(1)	search a file for a pattern using full regular expressions
eject	eject(1)	eject media such as CD-ROM and floppy from drive
enable	enable(1)	enable/disable LP printers
env	env(1)	obtain or alter environment variables for command execution
eqn	eqn(1)	typeset mathematics test
errange	ckrange(1)	prompts for and validates an integer
errdate	ckdate(1)	prompts for and validates a date
errgid	ckgid(1)	prompts for and validates a group id
errint	ckint(1)	display a prompt; verify and return an integer value
erritem	ckitem(1)	build a menu; prompt for and return a menu item
error	error(1)	insert compiler error messages at right source lines
errpath	ckpath(1)	display a prompt; verify and return a pathname
errstr	ckstr(1)	display a prompt; verify and return a string answer
errtime	cktime(1)	display a prompt; verify and return a time of day
erruid	ckuid(1)	prompts for and validates a user ID
erryorn	ckyorn(1)	prompts for and validates yes/no
eval	exec(1)	shell built-in functions to execute other commands
ex	ex(1)	text editor
exec	exec(1)	shell built-in functions to execute other commands
exit	exit(1)	shell built-in functions to enable the execution of the shell to advance beyond its sequence of steps

expand	expand(1)	expand TAB characters to SPACE characters, and vice versa
export	set(1)	shell built-in functions to determine the characteristics for environmental variables of the current shell and its descendents
exportfs	exportfs(1B)	translates exportfs options to share/unshare commands
expr	expr(1)	evaluate arguments as an expression
expr	expr(1B)	evaluate arguments as a logical, arithmetic, or string expression
exstr	exstr(1)	extract strings from source files
face	face(1)	executable for the Framed Access Command Environment Interface
factor	factor(1)	obtain the prime factors of a number
false	true(1)	provide truth values
fastboot	fastboot(1B)	reboot/halt the system without checking the disks
fasthalt	fastboot(1B)	reboot/halt the system without checking the disks
fc	history(1)	shell built-in functions to re-use previous command-lines from the current shell
fdformat	fdformat(1)	format floppy diskette
fg	jobs(1)	shell built-in functions to control process execution
fgrep	fgrep(1)	search a file for a character string
file	file(1)	determine file type
file	file(1B)	determine the type of a file by examining its contents
find	find(1)	find files
finger	finger(1)	display information about local and remote users
fmlcut	fmlcut(1F)	cut out selected fields of each line of a file
fmlexpr	fmlexpr(1F)	evaluate arguments as an expression
fmlgrep	fmlgrep(1F)	search a file for a pattern
fmlj	fmlj(1)	invoke FMLI
fmlt	fmlt(1)	simple text formatters
fmlmsg	fmlmsg(1)	display a message on stderr or system console
fold	fold(1)	fold long lines
for	for(1)	shell built-in functions to repeatedly execute action(s) for a selected number of times
foreach	for(1)	shell built-in functions to repeatedly execute action(s) for a selected number of times
from	from(1B)	display the sender and date of newly-arrived mail messages
ftp	ftp(1)	file transfer program

function	function(1)	shell built-in command to define a function which is usable within this shell
gcore	gcore(1)	get core images of running processes
gencat	gencat(1)	generate a formatted message catalog
get	sccs-get(1)	retrieve a version of an SCCS file
getfrm	getfrm(1F)	returns the current frameID number
getitems	getitems(1F)	returns a list of currently marked menu items
getopt	getopt(1)	parse command options
getoptcv	getoptcv(1)	convert to getopt to parse command options
getopts	getopts(1)	shell built-in function to parse command-line options
gettext	gettext(1)	retrieve text string from message database
gettxt	gettxt(1)	retrieve a text string from a message database
gigiplot	plot(1B)	graphics filters for various plotters
glob	glob(1)	shell built-in function to expand a word list
goto	exit(1)	shell built-in functions to enable the execution of the shell to advance beyond its sequence of steps
gprof	gprof(1)	display call-graph profile data
graph	graph(1)	draw a graph
grep	grep(1)	search a file for a pattern
groups	groups(1)	print group membership of user
groups	groups(1B)	display a user's group memberships
grpck	grpck(1B)	check group database entries
hash	hash(1)	shell built-in functions to evaluate the internal hash table of the contents of directories
hashstat	hash(1)	shell built-in functions to evaluate the internal hash table of the contents of directories
head	head(1)	display first few lines of files
help	sccs-help(1)	ask for help regarding SCCS error or warning messages
helpdate	ckdate(1)	prompts for and validates a date
helpgid	ckgid(1)	prompts for and validates a group id
helpint	ckint(1)	display a prompt; verify and return an integer value
helpitem	ckitem(1)	build a menu; prompt for and return a menu item
helppath	ckpath(1)	display a prompt; verify and return a pathname
helpprange	ckrange(1)	prompts for and validates an integer
helpstr	ckstr(1)	display a prompt; verify and return a string answer

helptime	cktime(1)	display a prompt; verify and return a time of day
helpuid	ckuid(1)	prompts for and validates a user ID
helpyorn	ckyorn(1)	prompts for and validates yes/no
history	history(1)	shell built-in functions to re-use previous command-lines from the current shell
hostid	hostid(1B)	print the numeric identifier of the current host
hostname	hostname(1B)	set or print name of current host system
hp7221plot	plot(1B)	graphics filters for various plotters
hpplot	plot(1B)	graphics filters for various plotters
i286	machid(1)	get processor type truth value
i386	machid(1)	get processor type truth value
i486	machid(1)	get processor type truth value
i860	machid(1)	get processor type truth value
iAPX286	machid(1)	get processor type truth value
iconv	iconv(1)	code set conversion utility
if	if(1)	shell built-in functions to evaluate expression(s) or to make execution of actions dependent upon the evaluation of expression(s)
implot	plot(1B)	graphics filters for various plotters
indicator	indicator(1F)	display application specific alarms and/or the "working" indicator
indxbib	indxbib(1)	create an inverted index to a bibliographic database
install	install(1B)	install files
ipcrm	ipcrm(1)	remove a message queue, semaphore set, or shared memory ID
ipcs	ipcs(1)	report inter-process communication facilities status
jobs	jobs(1)	shell built-in functions to control process execution
join	join(1)	relational database operator
jsh	sh(1)	shell: the standard shell, and job control shell -- command interpreters
kbd	kbd(1)	manipulate the state of keyboard or display the type of keyboard
kdestroy	kdestroy(1)	destroy Kerberos tickets
kerberos	kerberos(1)	introduction to the Kerberos system
keylogin	keylogin(1)	decrypt and store secret key with keyserv
keylogout	keylogout(1)	delete stored secret key with keyserv
kill	kill(1)	terminate a process by default
kinit	kinit(1)	Kerberos login utility
klist	klist(1)	list currently held Kerberos tickets

ksh	ksh(1)	KornShell, a standard/restricted command and programming language
ksrvtgt	ksrvtgt(1)	fetch and store Kerberos ticket-granting ticket using a service key
last	last(1)	display login and logout information about users and terminals
lastcomm	lastcomm(1)	display the last commands executed, in reverse order
ld	ld(1)	link editor for object files
ld	ld(1B)	link editor, dynamic link editor
ldd	ldd(1)	list dynamic dependencies of executable files or shared objects
let	let(1)	shell built-in function to evaluate one or more arithmetic expressions
lex	lex(1)	lexical analysis program generator
limit	limit(1)	shell built-in functions to set/get limitations on the system resources available to the current shell and it's descendents.
line	line(1)	read one line
listusers	listusers(1)	list user login information
ln	ln(1)	make hard or symbolic links to files
ln	ln(1B)	make hard or symbolic links to files
loadfont	loadfont(1)	display or change font information in the RAM of the video card on an x86 system in text mode
loadkeys	loadkeys(1)	load and dump keyboard translation tables
logger	logger(1)	add entries to the system log
logger	logger(1B)	add entries to the system log
login	login(1)	sign on to the system
logname	logname(1)	get the name of the user running the process
logout	logout(1)	shell built-in function to exit from a login session
longline	readfile(1F)	reads file, gets longest line
look	look(1)	find words in the system dictionary or lines in a sorted list
lookbib	lookbib(1)	find references in a bibliographic database
lorder	lorder(1)	find ordering relation for an object or library archive
lp	lp(1)	send/cancel requests to an LP print service
lpc	lpc(1B)	line printer control program
lpq	lpq(1B)	display the queue of printer jobs
lpr	lpr(1B)	send a job to the printer
lprm	lprm(1B)	remove jobs from the printer queue

lpstat	lpstat(1)	print information about the status of the LP print service
lpctest	lpctest(1B)	generate lineprinter ripple pattern
ls	ls(1)	list contents of directory
ls	ls(1B)	list the contents of a directory
m4	m4(1)	macro processor
mach	mach(1B)	display the processor type of the current host
machid	machid(1)	get processor type truth value
Mail	mailx(1)	interactive message processing system
mail	mail(1)	read mail or send mail to users
mail	mailx(1)	interactive message processing system
mailstats	mailstats(1)	print statistics collected by sendmail
mailx	mailx(1)	interactive message processing system
make	make(1S)	maintain, update, and regenerate related programs and files
man	man(1)	find and display reference manual pages
mconnect	mconnect(1)	connect to SMTP mail server socket
mcs	mcs(1)	manipulate the comment section of an object file
mesg	mesg(1)	permit or deny messages
message	message(1F)	puts its arguments on FMLI message line
mkdir	mkdir(1)	make directories
mkmsgs	mkmsgs(1)	create message files for use by gettxt
mkstr	mkstr(1B)	create an error message file by massaging C source files
more	more(1)	browse or page through a text file
msgfmt	msgfmt(1)	create a message object from a message file
mt	mt(1)	magnetic tape control
mv	mv(1)	move files
nawk	nawk(1)	pattern scanning and processing language
neqn	eqn(1)	typeset mathematics test
newaliases	newaliases(1)	rebuild the data base for the mail aliases file
newform	newform(1)	change the format of a text file
newgrp	newgrp(1)	shell built-in function to allow new group permissions to the user
news	news(1)	print news items
nice	nice(1)	run a command at low priority
NIS+	nis+(1)	a new version of the network information name service
nis+	nis+(1)	a new version of the network information name service
nis	nis+(1)	a new version of the network information name service
niscat	niscat(1)	display NIS+ tables and objects

nischgrp	nischgrp(1)	change the group owner of a NIS+ object
nischmod	nischmod(1)	change access rights on a NIS+ object
nischown	nischown(1)	change the owner of a NIS+ object
nischttl	nischttl(1)	change the time to live value of a NIS+ object
nisdefaults	nisdefaults(1)	display NIS+ default values
niserror	niserror(1)	display NIS+ error messages
nisgrep	nismatch(1)	utilities for searching NIS+ tables
nisgrpadm	nisgrpadm(1)	NIS+ group administration command
nisln	nisln(1)	symbolically link NIS+ objects
nisls	nisls(1)	list the contents of a NIS+ directory
nismatch	nismatch(1)	utilities for searching NIS+ tables
nismkdir	nismkdir(1)	create NIS+ directories
nispasswd	nispasswd(1)	change NIS+ password information
nismrm	nismrm(1)	remove NIS+ objects from the namespace
nismrmdir	nismrmdir(1)	remove NIS+ directories
nistbladm	nistbladm(1)	NIS+ table administration command
nistest	nistest(1)	return the state of the NIS+ namespace using a conditional expression
nl	nl(1)	line numbering filter
nm	nm(1)	print name list of an object file
nohup	nohup(1)	run a command immune to hangups and quits
notify	jobs(1)	shell built-in functions to control process execution
nroff	nroff(1)	format documents for display or line-printer
od	od(1)	octal dump
on	on(1)	execute a command on a remote system, but with the local environment
onintr	trap(1)	shell built-in functions to respond to (hardware) signals
pack	pack(1)	compress and expand files
page	more(1)	browse or page through a text file
pagesize	pagesize(1B)	display the size of a page of memory
passwd	passwd(1)	change login password and password attributes
paste	paste(1)	merge same lines of several files or subsequent lines of one file
pathconv	pathconv(1F)	search FMLI criteria for filename
pcat	pack(1)	compress and expand files
pcmapkeys	pcmapkeys(1)	set keyboard extended map and scancode translation for the PC console in text mode
pdp11	machid(1)	get processor type truth value
pg	pg(1)	files perusal filter for CRTs
pkginfo	pkginfo(1)	display software package information
pkgmk	pkgmk(1)	produce an installable package
pkgparam	pkgparam(1)	displays package parameter values

pkgproto	pkgproto(1)	generate prototype file entries for input to pkgmk command
pkgtrans	pkgtrans(1)	translate package format
plot	plot(1B)	graphics filters for various plotters
plottoa	plot(1B)	graphics filters for various plotters
popd	cd(1)	shell built-in functions to change the current working directory
postdaisy	postdaisy(1)	PostScript translator for Diablo 630 daisy-wheel files
postdmd	postdmd(1)	PostScript translator for DMD bitmap files
postio	postio(1)	serial interface for PostScript printers
postmd	postmd(1)	matrix display program for PostScript printers
postplot	postplot(1)	PostScript translator for plot(4) graphics files
postprint	postprint(1)	PostScript translator for text files
postreverse	postreverse(1)	reverse the page order in a PostScript file
posttek	posttek(1)	PostScript translator for Tektronix 4014 files
pr	pr(1)	print files
print	print(1)	shell built-in function to output characters to the screen or window
printenv	printenv(1B)	display environment variables currently set
printf	printf(1)	print formatted output
priocntl	priocntl(1)	display or set scheduling parameters of specified process(es)
prof	prof(1)	display profile data
prs	sccs-prs(1)	display selected portions of an SCCS history
prt	sccs-prt(1)	display delta table information from an SCCS file
ps	ps(1)	report process status
ps	ps(1B)	display the status of current processes
pushd	cd(1)	shell built-in functions to change the current working directory
pwd	pwd(1)	working directory name
rcp	rcp(1)	remote file copy
rdist	rdist(1)	remote file distribution program
read	read(1)	shell built-in function to receive from standard input (keyboard)
readfile	readfile(1F)	reads file, gets longest line
readonly	readonly(1)	shell built-in function to protect the value of the given variable from reassignment
red	ed(1)	text editor
refer	refer(1)	expand and insert references from a bibliographic database
regcmp	regcmp(1)	regular expression compile
regex	regex(1F)	match patterns against a string

rehash	hash(1)	shell built-in functions to evaluate the internal hash table of the contents of directories
reinit	reinit(1F)	runs an initialization file
remote_shell	rsh(1)	remote shell
remsh	rsh(1)	remote shell
renice	renice(1B)	alter priority of running processes
repeat	repeat(1)	shell built-in function to execute a command more than once
reset	reset(1F)	reset the current form field to its default values
reset	tset(1B)	establish or restore terminal characteristics
return	exit(1)	shell built-in functions to enable the execution of the shell to advance beyond its sequence of steps
rksh	ksh(1)	KornShell, a standard/restricted command and programming language
rlogin	rlogin(1)	remote login
rm	rm(1)	remove files or directories
rmail	mail(1)	read mail or send mail to users
rm del	sccs-rmdel(1)	remove a delta from an SCCS file
rmdir	rm(1)	remove files or directories
roffbib	roffbib(1)	format and print a bibliographic database
rpcgen	rpcgen(1)	an RPC protocol compiler
rsh	rsh(1)	remote shell
run	run(1F)	run an executable
rup	rup(1)	show host status of remote machines (RPC version)
rup	rup(1C)	show host status of remote machines (RPC version)
runtime	runtime(1)	show host status of local machines
rusage	rusage(1B)	print resource usage for a command
rusers	rusers(1)	who's logged in on remote machines
rwho	rwho(1)	who's logged in on local machines
sact	sccs-sact(1)	show editing activity status of an SCCS file
sag	sag(1)	system activity graph
sar	sar(1)	system activity reporter
sccs-admin	sccs-admin(1)	create and administer SCCS history files
sccs-cdc	sccs-cdc(1)	change the delta commentary of an SCCS delta
sccs-comb	sccs-comb(1)	combine SCCS deltas
sccs-delta	sccs-delta(1)	make a delta to an SCCS file
sccs-get	sccs-get(1)	retrieve a version of an SCCS file
sccs-help	sccs-help(1)	ask for help regarding SCCS error or warning messages
sccs-prs	sccs-prs(1)	display selected portions of an SCCS history

sccs-prt	sccs-prt(1)	display delta table information from an SCCS file
sccs-rmdel	sccs-rmdel(1)	remove a delta from an SCCS file
sccs-sact	sccs-sact(1)	show editing activity status of an SCCS file
sccs-sccsdiff	sccs-sccsdiff(1)	compare two versions of an SCCS file
sccs-unget	sccs-unget(1)	undo a previous get of an SCCS file
sccs-val	sccs-val(1)	validate an SCCS file
sccs	sccs(1)	front end for the Source Code Control System (SCCS)
sccsdiff	sccs-sccsdiff(1)	compare two versions of an SCCS file
script	script(1)	make record of a terminal session
sdiff	sdiff(1)	print differences between two files side-by-side
sed	sed(1)	stream editor
select	case(1)	shell built-in functions to choose from among a list of actions
set	set(1)	shell built-in functions to determine the characteristics for environmental variables of the current shell and its descendents
set	set(1F)	set and unset local or global environment variables
setcolor	setcolor(1F)	redefine or create a color
setenv	set(1)	shell built-in functions to determine the characteristics for environmental variables of the current shell and its descendents
sh	sh(1)	shell: the standard shell, and job control shell -- command interpreters
shell	shell(1F)	run a command using shell
shell_builtins	shell_builtins(1)	shell command interpreter built-in functions
shift	shift(1)	shell built-in function to traverse either a shell's argument list or a list of field-separated words
shutdown	shutdown(1B)	close down the system at a given time
size	size(1)	print section sizes in bytes of object files
sleep	sleep(1)	suspend execution for an interval
soelim	soelim(1)	resolve and eliminate .so requests from nroff or troff input
sort	sort(1)	sort and/or merge files
sortbib	sortbib(1)	sort a bibliographic database
source	exec(1)	shell built-in functions to execute other commands
sparc	machid(1)	get processor type truth value
spell	spell(1)	find spelling errors
spline	spline(1)	interpolate smooth curve
split	split(1)	split a file into pieces

srchtxt	srchtxt(1)	display contents of, or search for a text string in, message data bases
stop	jobs(1)	shell built-in functions to control process execution
strchg	strchg(1)	change or query stream configuration
strconf	strchg(1)	change or query stream configuration
strings	strings(1)	find printable strings in an object or binary file
strip	strip(1)	strip symbol table, debugging and line number information from an object file
stty	stty(1)	set the options for a terminal
stty	stty(1B)	set the options for a terminal
sum	sum(1)	print checksum and block count for a file
sum	sum(1B)	calculate a checksum for a file
sun	machid(1)	get processor type truth value
suspend	suspend(1)	shell built-in function to halt the current shell
switch	case(1)	shell built-in functions to choose from among a list of actions
symorder	symorder(1)	rearrange a list of symbols
sysV-make	sysV-make(1)	maintain, update, and regenerate groups of programs
t300	plot(1B)	graphics filters for various plotters
t300s	plot(1B)	graphics filters for various plotters
t4013	plot(1B)	graphics filters for various plotters
t450	plot(1B)	graphics filters for various plotters
tabs	tabs(1)	set tabs on a terminal
tail	tail(1)	deliver the last part of a file
talk	talk(1)	talk to another user
tar	tar(1)	create tape archives, and add or extract files
tbl	tbl(1)	format tables for nroff or troff
tcopy	tcopy(1)	copy a magnetic tape
tee	tee(1)	replicate the standard output
tek	plot(1B)	graphics filters for various plotters
telnet	telnet(1)	user interface to a remote system using the TELNET protocol
test	if(1)	shell built-in functions to evaluate expression(s) or to make execution of actions dependent upon the evaluation of expression(s)
test	test(1B)	condition evaluation command
test	test(1F)	condition evaluation command
tftp	tftp(1)	trivial file transfer program
time	time(1)	time a command

times	times(1)	shell built-in function to report time usages of the current shell
timex	timex(1)	time a command; report process data and system activity
tip	tip(1)	connect to remote system
touch	touch(1)	update access time and/or modification time of a file
tput	tput(1)	initialize a terminal or query terminfo database
tr	tr(1)	translate characters
tr	tr(1B)	translate characters
trap	trap(1)	shell built-in functions to respond to (hardware) signals
troff	troff(1)	typeset or format documents
true	true(1)	provide truth values
truss	truss(1)	trace system calls and signals
tset	tset(1B)	establish or restore terminal characteristics
tsort	tsort(1)	topological sort
tty	tty(1)	get the name of the terminal
type	typeset(1)	shell built-in functions to set/get attributes and values for shell variables and functions
typeset	typeset(1)	shell built-in functions to set/get attributes and values for shell variables and functions
u370	machid(1)	get processor type truth value
u3b15	machid(1)	get processor type truth value
u3b2	machid(1)	get processor type truth value
u3b5	machid(1)	get processor type truth value
u3b	machid(1)	get processor type truth value
ucblinks	ucblinks(1B)	adds /dev entries to give SunOS 4.1 compatible names to SunOS 5.x devices
ul	ul(1)	do underlining
ulimit	limit(1)	shell built-in functions to set/get limitations on the system resources available to the current shell and it's descendents.
umask	umask(1)	shell built-in function to restrict read/write/execute permissions
unalias	alias(1)	shell built-in functions to create your own pseudonym or shorthand for a command or series of commands
uname	uname(1)	print name of current system
uncompress	compress(1)	compress, uncompress files or display expanded files
unexpand	expand(1)	expand TAB characters to SPACE characters, and vice versa
unget	scs-unget(1)	undo a previous get of an SCCS file

unhash	hash(1)	shell built-in functions to evaluate the internal hash table of the contents of directories
unifdef	unifdef(1)	resolve and remove ifdef'ed lines from C program source
uniq	uniq(1)	report repeated lines in a file
units	units(1)	converts quantities expressed in standard scales to other scales
unix2dos	unix2dos(1)	convert text file from ISO format to DOS format
unlimit	limit(1)	shell built-in functions to set/get limitations on the system resources available to the current shell and it's descendents.
unpack	pack(1)	compress and expand files
unset	set(1)	shell built-in functions to determine the characteristics for environmental variables of the current shell and its descendents
unset	set(1F)	set and unset local or global environment variables
unsetenv	set(1)	shell built-in functions to determine the characteristics for environmental variables of the current shell and its descendents
until	while(1)	shell built-in functions to repetitively execute a set of actions while/until conditions are evaluated TRUE
uptime	uptime(1)	show how long the system has been up
users	users(1B)	display a compact list of users logged in
uucp	uucp(1C)	UNIX-to-UNIX system copy
uudecode	uuencode(1C)	encode a binary file, or decode its ASCII representation
uuencode	uuencode(1C)	encode a binary file, or decode its ASCII representation
uuglist	uuglist(1C)	print the list of service grades that are available on this UNIX system
uulog	uucp(1C)	UNIX-to-UNIX system copy
uuname	uucp(1C)	UNIX-to-UNIX system copy
uupick	uuto(1C)	public UNIX-to-UNIX system file copy
uustat	uustat(1C)	uucp status inquiry and job control
uuto	uuto(1C)	public UNIX-to-UNIX system file copy
uux	uux(1C)	UNIX-to-UNIX system command execution
vacation	vacation(1)	reply to mail automatically
val	sccs-val(1)	validate an SCCS file
valdate	ckdate(1)	prompts for and validates a date
valgid	ckgid(1)	prompts for and validates a group id

valint	ckint(1)	display a prompt; verify and return an integer value
valpath	ckpath(1)	display a prompt; verify and return a pathname
valrange	ckrange(1)	prompts for and validates an integer
valstr	ckstr(1)	display a prompt; verify and return a string answer
valtime	cktime(1)	display a prompt; verify and return a time of day
valuid	ckuid(1)	prompts for and validates a user ID
valyorn	ckyorn(1)	prompts for and validates yes/no
vax	machid(1)	get processor type truth value
vc	vc(1)	version control
vedit	vi(1)	screen-oriented (visual) display editor based on ex
vgrind	vgrind(1)	grind nice program listings
vi	vi(1)	screen-oriented (visual) display editor based on ex
view	vi(1)	screen-oriented (visual) display editor based on ex
vipw	vipw(1B)	edit the password file
volcancel	volcancel(1)	cancel user's request for removable media that is not currently in drive
volcheck	volcheck(1)	check for media in a drive. Default checks all floppy media.
volmissing	volmissing(1)	notify user that volume requested is not in the CD-ROM or floppy drive
vplot	plot(1B)	graphics filters for various plotters
vsig	vsig(1F)	synchronize a co-process with the controlling FMLI application
w	w(1)	who is logged in, and what are they doing
wait	wait(1)	shell built-in function to wait for other jobs or processes
wc	wc(1)	display a count of lines, words and characters in a file
what	what(1)	extract SCCS version information from a file
whatis	whatis(1)	display a one-line summary about a keyword
whence	typeset(1)	shell built-in functions to set/get attributes and values for shell variables and functions
whereis	whereis(1B)	locate the binary, source, and manual page files for a command
which	which(1)	locate a command; display its pathname or alias

while	while(1)	shell built-in functions to repetitively execute a set of actions while/until conditions are evaluated TRUE
who	who(1)	who is on the system
whoami	whoami(1B)	display the effective current username
whois	whois(1)	Internet user name directory service
write	write(1)	write to another user
xargs	xargs(1)	construct argument list(s) and execute command
xgettext	xgettext(1)	extract gettext call strings from C programs
xstr	xstr(1)	extract strings from C programs to implement shared strings
yacc	yacc(1)	yet another compiler-compiler
ypcat	ypcat(1)	print values in a NIS database
ypmatch	ypmatch(1)	print the value of one or more keys from a NIS map
yppasswd	yppasswd(1)	change your network password in the NIS database
ypwhich	ypwhich(1)	return name of NIS server or map master
zcat	compress(1)	compress, uncompress files or display expanded files

NAME	acctcom – search and print process accounting files
SYNOPSIS	<pre>acctcom [-abfhikmqrtv] [-C sec] [-e time] [-E time] [-g group] [-H factor] [-I chars] [-l line] [-n pattern] [-o output-file] [-O sec] [-s time] [-S time] [-u user] [filename ...]</pre>
AVAILABILITY	SUNWaccu
DESCRIPTION	<p>acctcom reads <i>filenames</i>, the standard input, or /var/adm/pacct, in the form described by acct(4) and writes selected records to standard output. Each record represents the execution of one process. The output shows the COMMAND NAME, USER, TTYNAME, START TIME, END TIME, REAL (SEC), CPU (SEC), MEAN SIZE (K), and optionally, F (the fork()/exec() flag: 1 for fork() without exec()), STAT (the system exit status), HOG FACTOR, KCORE MIN, CPU FACTOR, CHARS TRNSFD, and BLOCKS READ (total blocks read and written).</p> <p>A '#' is prepended to the command name if the command was executed with super-user privileges. If a process is not associated with a known terminal, a '?' is printed in the TTYNAME field.</p> <p>If no <i>filename</i> is specified, and if the standard input is associated with a terminal or /dev/null (as is the case when using '&' in the shell), /var/adm/pacct is read; otherwise, the standard input is read.</p> <p>If any <i>filename</i> arguments are given, they are read in their respective order. Each file is normally read forward, that is, in chronological order by process completion time. The file /var/adm/pacct is usually the current file to be examined; a busy system may need several such files of which all but the current file are found in /var/adm/pacctincr.</p>
OPTIONS	<p>-a Show some average statistics about the processes selected. The statistics will be printed after the output records.</p> <p>-b Read backwards, showing latest commands first. This option has no effect when standard input is read.</p> <p>-f Print the fork()/exec() flag and system exit status columns in the output. The numeric output for this option will be in octal.</p> <p>-h Instead of mean memory size, show the fraction of total available CPU time consumed by the process during its execution. This "hog factor" is computed as (total CPU time)/(elapsed time).</p> <p>-i Print columns containing the I/O counts in the output.</p> <p>-k Instead of memory size, show total kcore-minutes.</p> <p>-m Show mean core size (the default).</p> <p>-q Do not print any output records, just print the average statistics as with the -a option.</p> <p>-r Show CPU factor (user-time/(system-time + user-time)).</p>

-t	Show separate system and user CPU times.
-v	Exclude column headings from the output.
-C sec	Show only processes with total CPU time (system-time + user-time) exceeding <i>sec</i> seconds.
-e time	Select processes existing at or before <i>time</i> .
-E time	Select processes ending at or before <i>time</i> . Using the same <i>time</i> for both -S and -E shows the processes that existed at <i>time</i> .
-g group	Show only processes belonging to <i>group</i> . The <i>group</i> may be designated by either the group ID or group name.
-H factor	Show only processes that exceed <i>factor</i> , where <i>factor</i> is the “hog factor” as explained in option -h above.
-I chars	Show only processes transferring more characters than the cutoff number given by <i>chars</i> .
-l line	Show only processes belonging to terminal /dev/term/line .
-n pattern	Show only commands matching <i>pattern</i> that may be a regular expression as in regcmp (3G), except + means one or more occurrences.
-o output-file	Copy selected process records in the input data format to <i>output-file</i> ; suppress printing to standard output.
-O sec	Show only processes with CPU system time exceeding <i>sec</i> seconds.
-s time	Select processes existing at or after <i>time</i> , given in the format <i>hr</i> [: <i>min</i> [: <i>sec</i>]].
-S time	Select processes starting at or after <i>time</i> .
-u user	Show only processes belonging to <i>user</i> . The user may be specified by a user ID, a login name that is then converted to a user ID, ‘ # ’ (which designates only those processes executed with superuser privileges), or ‘ ? ’ (which designates only those processes associated with unknown user IDs).

FILES /etc/group
 /etc/passwd
 /var/adm/pacctincr

SEE ALSO **ps**(1), **acct**(1M), **acctcms**(1M), **acctcon**(1M), **acctmerg**(1M), **acctprc**(1M), **acctsh**(1M), **fwtmp**(1M), **runacct**(1M), **su**(1M), **acct**(2), **regcmp**(3G), **acct**(4), **utmp**(4)

Security, Performance, and Accounting Administration

NOTES **acctcom** reports only on processes that have terminated; use **ps**(1) for active processes. If *time* exceeds the present time, then *time* is interpreted as occurring on the previous day.

NAME	adb – general-purpose debugger
SYNOPSIS	adb [-w] [-k] [-I dir] [-P prompt] [<i>objectfile</i> [<i>corefile</i> [<i>swapfile</i>]]]
AVAILABILITY	SUNWtoo
DESCRIPTION	<p>adb is an interactive, general-purpose debugger. It can be used to examine files and provides a controlled environment for the execution of programs.</p> <p><i>objectfile</i> is normally an executable program file, preferably containing a symbol table. If the file does not contain a symbol table, it can still be examined, but the symbolic features of adb cannot be used. The default for <i>objectfile</i> is a.out. <i>corefile</i> is assumed to be a core image file produced after executing <i>objectfile</i>. The default for <i>corefile</i> is core. <i>swapfile</i> is the image of the swap device used. It is valid only when used with the -k option.</p>
OPTIONS	<p>-w Create both <i>objectfile</i> and <i>corefile</i>, if necessary, and open them for reading and writing so that they can be modified using adb.</p> <p>-k Perform kernel memory mapping; use when <i>corefile</i> is a system crash dump or /dev/mem, or when using a <i>swapfile</i>.</p> <p>-I dir Specify a directory where files to be read with \$< or \$<< (see below) will be sought; the default is /usr/kvm/lib/adb.</p> <p>-P prompt Specify the adb prompt string.</p>
USAGE	<p>adb reads commands from the standard input and displays responses on the standard output. It does not supply a prompt by default. It ignores the QUIT signal. INTERRUPT invokes the next adb command. adb generally recognizes command input of the form:</p> <p style="padding-left: 40px;">[<i>address</i>] [, <i>count</i>] [<i>command</i>] [;]</p> <p><i>address</i> and <i>count</i> (if supplied) are expressions that result, respectively, in a new current address, and a repetition count. <i>command</i> is composed of a verb followed by a modifier or list of modifiers.</p> <p>The symbol '.' represents the current location. It is initially zero. The default <i>count</i> is '1'.</p>
Expressions	<p>. The value of <i>dot</i>.</p> <p>+ The value of <i>dot</i> incremented by the current increment.</p> <p>^ The value of <i>dot</i> decremented by the current increment.</p> <p>& The last <i>address</i> typed. (In older versions of adb, "" was used.)</p> <p><i>integer</i> A number. The prefixes 0o and 0O indicate octal; 0t and 0T, decimal; 0x and 0X, hexadecimal (the default).</p> <p><i>int.frac</i> A floating-point number.</p> <p>'cccc' ASCII value of up to 4 characters.</p> <p><name The value of <i>name</i>, which is either a variable name or a register name.</p> <p><i>symbol</i> A symbol in the symbol table.</p> <p>(exp) The value of <i>exp</i>.</p>

<i>Unary Operators</i>	<i>*exp</i>	The contents of location <i>exp</i> in <i>corefile</i> .
	<i>%exp</i>	The contents of location <i>exp</i> in <i>objectfile</i> (In older versions of adb , '@' was used).
	<i>-exp</i>	Integer negation.
	<i>~exp</i>	Bitwise complement.
	<i>#exp</i>	Logical negation.
<i>Binary Operators</i>		Binary operators are left associative and have lower precedence than unary operators.
	<i>+</i>	Integer addition.
	<i>-</i>	Integer subtraction.
	<i>*</i>	Integer multiplication.
	<i>%</i>	Integer division.
	<i>&</i>	Bitwise conjunction ("AND").
	<i> </i>	Bitwise disjunction ("OR").
	<i>#</i>	<i>lhs</i> rounded up to the next multiple of <i>rhs</i> .
Variables		Named variables are set initially by adb but are not used subsequently.
	0	The last value printed.
	1	The last offset part of an instruction source.
	2	The previous value of variable 1.
	9	The count on the last \$< or \$<< command.
		On entry the following are set from the system header in the <i>corefile</i> or <i>objectfile</i> as appropriate.
	b	The base address of the data segment.
	d	The data segment size.
	e	The entry point.
	m	The 'magic' number
	t	The text segment size.
Commands		Commands to adb consist of a <i>verb</i> followed by a <i>modifier</i> or list of modifiers.
<i>Verbs</i>	?	Print locations starting at <i>address</i> in <i>objectfile</i> .
	/	Print locations starting at <i>address</i> in <i>corefile</i> .
	=	Print the value of <i>address</i> itself.
	:	Manage a subprocess.
	>	Assign a value to a variable or register.
	RETURN	Repeat the previous command with a <i>count</i> of 1. Increment '.'.
	!	Shell escape.

?, /, and = Modifiers

The following format modifiers apply to the commands ?, /, and =. To specify a format, follow the command with an optional repeat count, and the desired format letter or letters:

{ ?/,= } [[*r*] *f*...]

where *r* is a repeat count, and *f* is one of the format letters listed below:

o	(<i>'</i> increment: 2) Print 2 bytes in octal.
O	(4) Print 4 bytes in octal.
q	(2) Print in signed octal.
Q	(4) Print long signed octal.
d	(2) Print in decimal.
D	(4) Print long decimal.
x	(2) Print 2 bytes in hexadecimal.
X	(4) Print 4 bytes in hexadecimal.
u	(2) Print as an unsigned decimal number.
U	(4) Print long unsigned decimal.
f	(4) Print a single-precision floating-point number.
F	(8) Print a double-precision floating-point number.
b	(1) Print the addressed byte in octal.
c	(1) Print the addressed character.
C	(1) Print the addressed character using \wedge escape convention.
s	(<i>n</i>) Print the addressed string.
S	(<i>n</i>) Print a string using the \wedge escape convention.
Y	(4) Print 4 bytes in date format.
i	(4 on SPARC; <i>n</i> on x86) Print as machine instructions.
a	(0) Print the value of ' <i>'</i> in symbolic form.
p	(4) Print the addressed value in symbolic form.
t	(0) Tab to the next appropriate TAB stop.
r	(0) Print a SPACE.
n	(0) Print a NEWLINE.
"...."	(0) Print the enclosed string.
^	(0) Decrement ' <i>'</i> .
+	(0) Increment ' <i>'</i> .
-	(0) Decrement ' <i>'</i> by 1.
? and / Modifiers	
I value mask	Apply <i>mask</i> and compare for <i>value</i> ; move ' <i>'</i> to matching location.
L value mask	Apply <i>mask</i> and compare for 4-byte <i>value</i> ; move ' <i>'</i> to matching location.
w value	Write the 2-byte <i>value</i> to address.
W value	Write the 4-byte <i>value</i> to address.
m b1 e1 f1[?]	Map new values for <i>b1</i> , <i>e1</i> , <i>f1</i> . If the ? or / is followed by * then the second segment (<i>b2</i> , <i>e2</i> , <i>f2</i>) of the address mapping is changed.
v	Like w , but writes only bytes at a time.

: <i>Modifiers</i>	b commands Set breakpoint, execute <i>commands</i> when reached. r Run <i>objectfile</i> as a subprocess. d Delete breakpoint at <i>address</i> . z Delete all breakpoints. cs x86: The subprocess is continued with signal <i>s</i> . ss Single-step the subprocess with signal <i>s</i> . i Add the signal specified by <i>address</i> to the list of signals passed directly to the subprocess. t Remove the signal specified by <i>address</i> from the list implicitly passed to the subprocess. k Terminate (kill) the current subprocess, if any. A Attach adb to an existing process id. (For example, 0t1234:A would attach adb to decimal process number 1234 .) R Release the previously attached process.
\$ <i>Modifiers</i>	<filename Read commands from the file <i>filename</i> . <<filename Similar to < , but can be used in a file of commands without closing the file. >filename Append output to <i>filename</i> , which is created if it does not exist. l x86: Show the current lightweight process (lwp) ID. L x86: Show all the lwp IDs. P Specify the adb prompt string. ? Print process ID, the signal which stopped the subprocess, and the registers. r Print the names and contents of the general CPU registers, and the instruction addressed by pc . x or X x86: Print the contents of floating point registers. \$x and \$X accept a "count" which determines the precision in which the floating point registers will be printed; the default is 25 . Using \$X will produce more verbose output than using \$x . x SPARC: Print the names and contents of floating-point registers 0 through 15. X SPARC: Print the names and contents of floating-point registers 16 through 31. b Print all breakpoints and their associated counts and commands. c C stack backtrace. On Sun-4 systems, it is impossible for adb to determine how many parameters were passed to a function. The default that adb chooses in a \$c command is to show the six parameter registers. This can be overridden by appending a hexadecimal number to the \$c command, specifying how many parameters to display. For example, the \$cf command will print 15 parameters for each function in the stack trace. C x86: Same as \$c , but in addition it displays the frame pointer values. d Set the default radix to <i>address</i> and report the new value. Note: <i>address</i> is interpreted in the (old) current radix. Thus ' 10\$d ' never changes the default radix.

e	Print the names and values of external variables.
w	Set the page width for output to <i>address</i> (default 80).
s	Set the limit for symbol matches to <i>address</i> (default 255).
o	All integers input are regarded as octal.
q	Exit from adb .
v	Print all non zero variables in octal.
m	Print the address map.
f	Print a list of known source filenames.
p	(<i>Kernel debugging</i>) Change the current kernel memory mapping to map the designated user structure to the address given by <i>u</i> ; this is the address of the user's proc structure.
i	Show which signals are passed to the subprocess with the minimum of adb interference.
W	Reopen <i>objectfile</i> and <i>corefile</i> for writing, as though the -w command-line argument had been given.

EXAMPLES To start **adb** on the running kernel, use (as **root**):

```
example# adb -k /dev/ksyms /dev/mem
```

/dev/ksyms is a special driver that provides an image of the kernel's symbol table. This can be used to examine kernel state and debug device drivers. Refer to the Debugging chapter in *Writing Device Drivers* for more information.

FILES	/usr/kvm/lib/adb	default directory in which files are to be read with \$< and \$<<.
	a.out	
	core	
	/dev/ksyms	special driver to provide an image of the kernel's symbolic table.

SEE ALSO **ptrace(2)**, **a.out(4)**, **core(4)**, **proc(4)**, **ksyms(7)**
Writing Device Drivers

DIAGNOSTICS **adb**, when there is no current command or format, comments about inaccessible files, syntax errors, abnormal termination of commands, etc. Exit status is **0**, unless last command failed or returned nonzero status.

NOTES **adb** should be changed to use the new format symbolic information generated by **-g**. **adb** is platform and release dependent. Kernel core dumps should be examined on the same platform they were created on.

BUGS On SPARC systems, there does not seem to be any way to clear all breakpoints. Since no shell is invoked to interpret the arguments of the **:r** command, the customary wild-card and variable expansions cannot occur. Since there is little type-checking on addresses, using a sourcefile address in an inappropriate context may lead to unexpected results.

The `$parameter-count` command is a workaround.

NAME	addbib – create or extend a bibliographic database
SYNOPSIS	addbib [-a] [-p <i>promptfile</i>] <i>database</i>
AVAILABILITY	SUNWdoc
DESCRIPTION	When addbib starts up, answering y to the initial Instructions? prompt yields directions; typing n or RETURN skips them. addbib then prompts for various bibliographic fields, reads responses from the terminal, and sends output records to <i>database</i> . A null response (just RETURN) means to leave out that field. A ‘-’ (minus sign) means to go back to the previous field. A trailing backslash allows a field to be continued on the next line. The repeating Continue? prompt allows the user either to resume by typing y or RETURN, to quit the current session by typing n or q , or to edit <i>database</i> with any system editor (see vi(1) , ex(1) , ed(1)).
OPTIONS	<p>-a Suppress prompting for an abstract; asking for an abstract is the default. Abstracts are ended with a CTRL-D.</p> <p>-p <i>promptfile</i> Use a new prompting skeleton, defined in <i>promptfile</i>. This file should contain prompt strings, a TAB, and the key-letters to be written to the <i>database</i>.</p>
USAGE Bibliography Key Letters	<p>The most common key-letters and their meanings are given below. addbib insulates you from these key-letters, since it gives you prompts in English, but if you edit the bibliography file later on, you will need to know this information.</p> <p>%A Author’s name</p> <p>%B Book containing article referenced</p> <p>%C City (place of publication)</p> <p>%D Date of publication</p> <p>%E Editor of book containing article referenced</p> <p>%F Footnote number or label (supplied by refer)</p> <p>%G Government order number</p> <p>%H Header commentary, printed before reference</p> <p>%I Issuer (publisher)</p> <p>%J Journal containing article</p> <p>%K Keywords to use in locating reference</p> <p>%L Label field used by -k option of refer</p> <p>%M Bell Labs Memorandum (undefined)</p> <p>%N Number within volume</p> <p>%O Other commentary, printed at end of reference</p>

%P Page number(s)
%Q Corporate or Foreign Author (unreversed)
%R Report, paper, or thesis (unpublished)
%S Series title
%T Title of article or book
%V Volume number
%X Abstract — used by **roffbib**, not by **refer**
%Y,Z Ignored by **refer**

EXAMPLES

Except for A, each field should be given just once. Only relevant fields should be supplied.

%A **Mark Twain**
%T **Life on the Mississippi**
%I **Penguin Books**
%C **New York**
%D **1978**

SEE ALSO

ed(1), **ex(1)**, **indxbib(1)**, **lookbib(1)**, **refer(1)**, **roffbib(1)**, **sortbib(1)**, **vi(1)**

NAME	alias, unalias – shell built-in functions to create your own pseudonym or shorthand for a command or series of commands
SYNOPSIS	
cs h	alias [<i>name</i> [<i>def</i>]] unalias <i>pattern</i>
k sh	†† alias [-t x] [<i>name</i> [= <i>value</i>]] ... unalias <i>name</i> ...
DESCRIPTION	
cs h	alias assigns <i>def</i> to the alias <i>name</i> . <i>def</i> is a list of words that may contain escaped history-substitution metasyntax. <i>name</i> is not allowed to be alias or unalias . If <i>def</i> is omitted, the alias <i>name</i> is displayed along with its current definition. If both <i>name</i> and <i>def</i> are omitted, all aliases are displayed. unalias discards aliases that match (filename substitution) <i>pattern</i> . All aliases may be removed by ' unalias * '.
k sh	alias with no arguments prints the list of aliases in the form <i>name=value</i> on standard output. An <i>alias</i> is defined for each name whose <i>value</i> is given. A trailing space in <i>value</i> causes the next word to be checked for alias substitution. The -t flag is used to set and list tracked aliases. The value of a tracked alias is the full pathname corresponding to the given <i>name</i> . The value becomes undefined when the value of PATH is reset but the aliases remained tracked. Without the -t flag, for each <i>name</i> in the argument list for which no <i>value</i> is given, the name and value of the alias is printed. The -x flag is used to set or print <i>exported aliases</i> . An <i>exported alias</i> is defined for scripts invoked by name. The exit status is non-zero if a <i>name</i> is given, but no value, and no alias has been defined for the <i>name</i> . The <i>aliases</i> given by the list of <i>names</i> may be removed from the <i>alias</i> list with unalias . On this man page, ksh (1) commands that are preceded by one or two † (daggers) are treated specially in the following ways: 1. Variable assignment lists preceding the command remain in effect when the command completes. 2. I/O redirections are processed after variable assignments. 3. Errors cause a script that contains them to abort. 4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.
SEE ALSO	cs h(1), k sh(1), shell_builtins (1),

NAME	apropos – locate commands by keyword lookup
SYNOPSIS	apropos <i>keyword</i> ...
AVAILABILITY	SUNWdoc
DESCRIPTION	<p>apropos displays the man page name, section number, and a short description for each man page whose NAME line contains <i>keyword</i>. This information is contained in the <code>/usr/share/man/windex</code> database created by catman(1M). If catman(1M) was not run, or was run with the <code>-n</code> option, apropos fails. Each word is considered separately and the case of letters is ignored. Words which are part of other words are considered; for example, when looking for ‘compile’, apropos finds all instances of ‘compiler’ also.</p> <p>apropos is actually just the <code>-k</code> option to the man(1) command.</p> <p>Try</p> <p style="padding-left: 40px;">example% apropos password</p> <p>and</p> <p style="padding-left: 40px;">example% apropos editor</p> <p>If the line starts ‘<i>filename(section) ...</i>’ you can do ‘man -s section filename’ to display the man page for <i>filename</i>. Try</p> <p style="padding-left: 40px;">example% apropos format</p> <p>and then</p> <p style="padding-left: 40px;">example% man -s 3s printf</p> <p>to get the manual page on the subroutine printf().</p>
FILES	<code>/usr/share/man/windex</code> table of contents and keyword database
SEE ALSO	man (1), whatis (1), catman (1M)
DIAGNOSTICS	<p><code>/usr/share/man/windex: No such file or directory</code> This database does not exist. catman(1M) must be run to create it.</p>

NAME	ar – maintain portable archive or library
SYNOPSIS	ar [-V] – key [arg] [posname] afile [name. . .]
DESCRIPTION	<p>The ar command maintains groups of files combined into a single archive file. Its main use is to create and update library files. However, it can be used for any similar purpose. The magic string and the file headers used by ar consist of printable ASCII characters. If an archive is composed of printable files, the entire archive is printable.</p> <p>When ar creates an archive, it creates headers in a format that is portable across all machines. The portable archive format and structure are described in detail in ar(4). The archive symbol table (described in ar(4)) is used by the link editor ld to effect multiple passes over libraries of object files in an efficient manner. An archive symbol table is only created and maintained by ar when there is at least one object file in the archive. The archive symbol table is in a specially named file that is always the first file in the archive. This file is never mentioned or accessible to the user. Whenever the ar command is used to create or update the contents of such an archive, the symbol table is rebuilt. The s option described below will force the symbol table to be rebuilt.</p>
OPTIONS	<p>-V ar prints its version number on standard error.</p> <p>Unlike command options, the <i>key</i> is a required part of the ar command line. The <i>key</i> is formed with one of the following letters: drqtpmx. Arguments to the <i>key</i>, alternatively, are made with one or more of the following set: vuaibcs. <i>posname</i> is an archive member name used as a reference point in positioning other files in the archive. <i>afile</i> is the archive file. The <i>names</i> are constituent files in the archive file. The meanings of the <i>key</i> characters are as follows:</p> <p>d Delete the named files from the archive file.</p> <p>r Replace the named files in the archive file. If the optional character u is used with r, then only those files with dates of modification later than the archive files are replaced. If an optional positioning character from the set abi is used, then the <i>posname</i> argument must be present and specifies that new files are to be placed after (a) or before (b or i) <i>posname</i>. Otherwise new files are placed at the end.</p> <p>q Quickly append the named files to the end of the archive file. Optional positioning characters are invalid. The command does not check whether the added members are already in the archive. This option is useful to avoid quadratic behavior when creating a large archive piece-by-piece.</p> <p>t Print a table of contents of the archive file. If no names are given, all files in the archive are listed. If names are given, only those files are listed.</p> <p>p Print the named files in the archive.</p> <p>m Move the named files to the end of the archive. If a positioning character is present, then the <i>posname</i> argument must be present and, as in r, specifies where the files are to be moved.</p>

x Extract the named files. If no names are given, all files in the archive are extracted. In neither case does **x** alter the archive file.

The meanings of the other key arguments are as follows:

v Give a verbose file-by-file description of the making of a new archive file from the old archive and the constituent files. When used with **t**, give a long listing of all information about the files. When used with **x**, print the filename preceding each extraction.

c Suppress the message that is produced by default when *afile* is created.

s Force the regeneration of the archive symbol table even if **ar(1)** is not invoked with a command which will modify the archive contents. This command is useful to restore the archive symbol table after the **strip(1)** command has been used on the archive.

SEE ALSO **ld(1)**, **lorder(1)**, **strip(1)**, **a.out(4)**, **ar(4)**

NOTES If the same file is mentioned twice in an argument list, it may be put in the archive twice. By convention, archives are suffixed with the characters **.a**.

NAME	arch – display the architecture of the current host
SYNOPSIS	/usr/ucb/arch /usr/ucb/arch -k /usr/ucb/arch <i>archname</i>
DESCRIPTION	arch displays the application architecture of the current host system. Systems can be broadly classified by their <i>architectures</i> , which define what executables will run on which machines. A distinction can be made between <i>kernel architecture</i> and <i>application architecture</i> (or, commonly, just “architecture”). Machines that run different kernels due to underlying hardware differences may be able to run the same application programs.
OPTIONS	-k Display the kernel architecture, such as sun4 , sun4c , etc. This defines which specific SunOS kernel will run on the machine, and has implications only for programs that depend on the kernel explicitly (for example, ps(1)). <i>archname</i> Return “true” (exit status 0) if <i>application</i> binaries for <i>archname</i> can run on the current host system, otherwise, return “false” (exit status 1). This is the preferred method for installation scripts to determine the environment of the host machine; that is, which architecture of a multi-architecture release to install on this machine. <i>archname</i> must be a valid application architecture.
SEE ALSO	mach(1B) , uname(1)

NAME	as – assembler
SPARC SYNOPSIS	as [-b] [-K PIC] [-L] [-m] [-n] [-o outfile] [-P [[-Dname] [-Dname=def] [-Ipath] [-Uname]]...] [-q] [-Qy / n] [-s] [-S[a / C]] [-T] [-V] [-xF] <i>filename</i> ...
x86 SYNOPSIS	as [-m] [-n] [-o outfile] [-P [[-Dname] [-Dname=def] [-Ipath] [-Uname]]...] [-Qy / n] [-s] [-V] <i>filename</i> ...
AVAILABILITY	SUNWspot
DESCRIPTION	The as command creates object files from assembly language source <i>files</i> .
OPTIONS	The following flags may be specified in any order: <ul style="list-style-type: none"> -Dname -Dname=def When the -P option is in effect, these options are passed to the cpp(1) preprocessor without interpretation by the as command; otherwise, they are ignored. -Ipath When the -P option is in effect, this option is passed to the cpp(1) preprocessor without interpretation by the as command; otherwise, it is ignored. -m Run the m4(1) macro processor on the input to the assembler. -n Suppress all the warnings while assembling. -o outfile Put the output of the assembly in <i>outfile</i>. By default, the output file name is formed by removing the .s suffix, if there is one, from the input file name and appending a .o suffix. -P Run cpp(1), the C preprocessor, on the files being assembled. The preprocessor is run separately on each input file, not on their concatenation. The preprocessor output is passed to the assembler. -Qy n Produce the "assembler version" information in the comment section of the output object file if the <i>y</i> option is specified; if the <i>n</i> option is specified, the information is suppressed. -s Place all stabs in the .stabs section. By default, stabs are placed in stabs.excl sections, which are stripped out by the static linker, ld(1), during final execution. When the -s option is used, stabs remain in the final executable because .stab sections are not stripped by the static linker. -Uname When the -P option is in effect, this option is passed to the cpp(1) preprocessor without interpretation by the as command; otherwise, it is ignored. -V Write the version number of the assembler being run on the standard error output.

SPARC Options	<p>-b Enable Sun SourceBrowser.</p> <p>-K PIC Generate position-independent code.</p> <p>-L Save all symbols, including temporary labels that are normally discarded to save space, in the ELF symbol table.</p> <p>-q Perform a quick assembly. When the -q option is used, the assembler's internal node list is not built and the assembler simply emits instructions as they are read. Note: This option disables many error checks. It is recommended that you do not use this option to assemble handwritten assembly language.</p> <p>-S[a/C] Produce a disassembly of the emitted code to the standard output. <ul style="list-style-type: none"> • Adding the character <i>a</i> to the option appends a comment line to each assembly code which indicates its relative address in its own section. • Adding the character <i>C</i> to the option prevents comment lines from appearing in the output. </p> <p>-T This is a migration option for 4.1 assembly files to be assembled on 5.0 systems. With this option, the symbol names in 4.1 assembly files will be interpreted as 5.0 symbol names.</p> <p>-xF Generates additional information for performance analysis of the executable using SPARCworks analyzer. If the input file does not contain any stabs (debugging directives), then the assembler will generate some default stabs which are needed by the SPARCworks analyzer. Also see the manual pages analyzer(1), dbx(1), and collector(1).</p>
ENVIRONMENT	<p>TMPDIR as normally creates temporary files in the directory /tmp. You may specify another directory by setting the environment variable TMPDIR to your chosen directory. (If TMPDIR isn't a valid directory, then as will use /tmp).</p>
FILES	<p>By default, as creates its temporary files in /tmp.</p>
SEE ALSO	<p>cc(1B), cpp(1), ld(1), m4(1), nm(1), strip(1), tmpnam(3S), a.out(4)</p>
NOTES	<p>If the -m (invoke the m4(1) macro processor) option is used, keywords for m4(1) cannot be used as symbols (variables, functions, labels) in the input file since m4(1) cannot determine which keywords are assembler symbols and which keywords are real m4(1) macros.</p> <p>Whenever possible, you should access the assembler through a compilation system interface program such as cc(1B).</p> <p>All undefined symbols are treated as global.</p>

NAME	at, batch – execute commands at a later time
SYNOPSIS	at [-csm] [-f <i>script</i>] [-q <i>queue</i>] <i>time</i> [<i>date</i>] [+ <i>increment</i>] at -l [<i>job...</i>] at -r <i>job...</i> batch
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>at and batch read commands from standard input to be executed at a later time. at allows you to specify when the commands should be executed, while jobs queued with batch will execute when system load level permits.</p> <p>Standard output and standard error output are mailed to the user unless they are redirected elsewhere. No mail will be sent to the user if the job does not produce any standard output and/or standard error output (and -m option is not specified). The shell environment variables (except TERM and TERMCAP), current directory, umask, and ulimit are retained when the commands are executed. Open file descriptors, traps, and priority are lost. The SHELL environment variable determines which shell is used.</p> <p>at and batch write the job number and scheduled time to standard error.</p>
at Access Control	<p>Users: Access to at or batch is allowed:</p> <ul style="list-style-type: none"> • if the user's name appears in /etc/cron.d/at.allow. • if /etc/cron.d/at.allow does not exist and the user's name is not in /etc/cron.d/at.deny. <p>Users: Access to at or batch is denied:</p> <ul style="list-style-type: none"> • if /etc/cron.d/at.allow exists and the user's name is not in it. • if /etc/cron.d/at.allow does not exist and user's name is in /etc/cron.d/at.deny. • if neither file exists. <p>Note: The rules for allow and deny apply to root only if the allow/deny files exist. These files can only be modified by the privileged user.</p>
time, date, increment Specifications	<p>If the DATEMSK environment variable is set, it points to a template file that at will use to determine the valid <i>time</i> and <i>date</i> values instead of the values described below. For more information about using DATEMSK, see sub-section 'at and DATEMSK'.</p> <p><i>time</i> <i>time</i> may be specified as follows, where <i>h</i> is hours and <i>m</i> is minutes: <i>h</i>, <i>hh</i>, <i>hhmm</i>, <i>h:m</i>, <i>h:mm</i>, <i>hh:m</i>, <i>hh:mm</i>. A 24-hour clock is assumed, unless am or pm is appended to <i>time</i>. If zulu is appended to <i>time</i>, it means Greenwich Mean Time (GMT). <i>time</i> can also take on the values: noon, midnight, and now. at now responds with the error message too late; use now with the <i>increment</i> argument, such as: at now + 1 minute.</p> <p><i>date</i> An optional <i>date</i> may be specified as either a month name followed by a</p>

day number (and possibly a year number preceded by a comma) or a day of the week. (Both the month name and the day of the week may be spelled out or abbreviated to three characters.) Two special “days”, **today** and **tomorrow** are recognized. If no *date* is given, **today** is assumed if the given hour is greater than the current hour and **tomorrow** is assumed if it is less. If the given month is less than the current month (and no year is given), next year is assumed.

increment The optional *increment* is simply a number suffixed by one of the following: **minutes**, **hours**, **days**, **weeks**, **months**, or **years**. (The singular form is also accepted.) The modifier **next** may precede the *increment*; it means “+ 1.”

Thus valid commands include:

at 0815am Jan 24
at 8:15am Jan 24
at now + 1 day
at now next day
at 5 pm Friday

Removing at and batch jobs

at -r removes jobs previously scheduled by **at** or **batch**. The job number is the number returned to you previously by the **at** or **batch** command. You can also get job numbers by typing **at -l**. You can only remove your own jobs unless you are the privileged user.

at and DATEMSK

If the environment variable **DATEMSK** is set, **at** will use its value as the full path name of a template file containing format strings. The strings consist of field descriptors and text characters and are used to provide a richer set of allowable date formats in different languages by appropriate settings of the environment variable **LANG** or **LC_TIME** (see **environ(5)**). (See **getdate(3C)** for the allowable list of field descriptors; this list is a subset of the descriptors allowed by **calendar(1)** that are listed on the **date(1)** manual page.) The formats described above for the *time* and *date* arguments, the special names **noon**, **midnight**, **now**, **next**, **today**, **tomorrow**, and the *increment* argument are not recognized when **DATEMSK** is set.

OPTIONS

- c** C shell. **csh(1)** is used to execute *script*. By default, the **SHELL** environment variable determines which shell to use.
- s** Standard (Bourne) shell. **sh(1)** is used to execute the job.
- m** Sends mail to the user after the job has been completed, indicating that the job is finished, even if the job produces no output. Mail is sent only if the job has not already generated a mail message.
- f script** Reads commands to be executed from the named *script* file.

- qqueue** Submit the job in queue *queue* rather than the default queue **a**. The valid *queues* are **a** through **z**. **batch** submits jobs in queue **b**. Queue **c** is reserved for **cron(1M)** and jobs cannot be submitted to that queue.
- l [job]** Reports all jobs scheduled for the invoking user, or just the *jobs* specified. This option excludes the use of all other options.
- r job** Removes specified *jobs* previously scheduled using **at**. This option excludes the use of all other options.

EXAMPLES

The **at** and **batch** commands read from standard input the commands to be executed at a later time. **sh(1)** provides different ways of specifying standard input. Within your commands, it may be useful to redirect standard output.

This sequence can be used at a terminal:

```
example% batch
example% sort filename > outfile
<control-D> (hold down 'control' and depress 'd')
```

This sequence, which shows redirecting standard error to a pipe, is useful in a shell procedure (the sequence of output redirection specifications is significant):

```
example% batch <<!
example% sort filename 2>&1 > outfile | mail loginid
!
```

To have a job reschedule itself, invoke **at** from within the shell procedure, by including code similar to the following within the shell file:

```
example% echo "sh shellfile" | at 1900 thursday next week
```

The following example shows the possible contents of a template file **AT.TEMPL** in **/var/tmp**.

```
%I %p, the %est of %B of the year %Y run the following job
%I %p, the %end of %B of the year %Y run the following job
%I %p, the %erd of %B of the year %Y run the following job
%I %p, the %eth of %B of the year %Y run the following job
%d/%m/%y
%H:%M:%S
%I:%M%p
```

The following are examples of valid invocations if the environment variable **DATMSK** is set to **/var/tmp/AT.TEMPL**.

```
at 2 PM, the 3rd of July of the year 2000 run the following job
at 3/4/99
at 10:30:30
at 2:30PM
```

FILES

/etc/cron.d	main cron directory
/etc/cron.d/at.allow	list of allowed users
/etc/cron.d/at.deny	list of denied users

/etc/cron.d/queuedefs scheduling information
/var/spool/cron/atjobs spool area for **at**.

SEE ALSO **atq(1)**, **atrm(1)**, **calendar(1)**, **crontab(1)**, **date(1)**, **kill(1)**, **mail(1)**, **nice(1)**, **ps(1)**, **sh(1)**, **sort(1)**, **cron(1M)**, **getdate(3C)**, **environ(5)**

DIAGNOSTICS Complains about various syntax errors and times out of range.

NAME	atq – display the jobs queued to run at specified times
SYNOPSIS	atq [-c] [-n] [<i>username</i> . . .]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>atq displays the at jobs queued up for the current user. at(1) is a utility that allows users to execute commands at a later date. If invoked by the privileged user, atq will display all jobs in the queue.</p> <p>If no options are given, the jobs are displayed in chronological order of execution.</p> <p>When a privileged user invokes atq without specifying <i>username</i>, the entire queue is displayed; when a <i>username</i> is specified, only those jobs belonging to the named user are displayed.</p>
OPTIONS	<p>-c Display the queued jobs in the order they were created (that is, the time that the at command was given).</p> <p>-n Display only the total number of jobs currently in the queue.</p>
FILES	<p>/var/spool/cron/atjobs spool area for at jobs.</p>
SEE ALSO	at(1) , atrm(1) , cron(1M)

NAME	atrm – remove jobs spooled by at or batch	
SYNOPSIS	atrm [-afi] [[<i>job #</i>] [<i>user</i>] ...]	
AVAILABILITY	SUNWcsu	
DESCRIPTION	<p>atrm removes delayed-execution jobs that were created with the at(1) command, but have not yet executed. The list of these jobs and associated job numbers can be displayed by using atq(1).</p> <p>atrm removes each job-number you specify, and/or all jobs belonging to the user you specify, provided that you own the indicated jobs.</p> <p>You can only remove jobs belonging to other users if you have super-user privileges.</p>	
OPTIONS	-a	All. Remove all unexecuted jobs that were created by the current user. If invoked by the privileged user, the entire queue will be flushed.
	-f	Force. All information regarding the removal of the specified jobs is suppressed.
	-i	Interactive. atrm asks if a job should be removed. If you respond with a y , the job will be removed.
FILES	/var/spool/cron/atjobs	spool area for at jobs
SEE ALSO	at(1) , atq(1) , cron(1M)	

NAME	audioconvert – convert audio file formats
SYNOPSIS	audioconvert [-pF] [-f outfmt] [-o outfile] [[-i infmt] [<i>file ...</i>]] ...
DESCRIPTION	<p>audioconvert converts audio data between a set of supported audio encodings and file formats. It can be used to compress and decompress audio data, to add audio file headers to raw audio data files, and to convert between standard data encodings, such as μ-law and linear PCM.</p> <p>If no filenames are present, audioconvert reads the data from the standard input stream and writes an audio file to the standard output. Otherwise, input files are processed in order, concatenated, and written to the output file.</p> <p>Input files are expected to contain audio file headers that identify the audio data format. If the audio data does not contain a recognizable header, the format must be specified with the -i option, using the <i>rate</i>, <i>encoding</i>, and <i>channels</i> keywords to identify the input data format.</p> <p>The output file format is derived by updating the format of the first input file with the format options in the -f specification. If -p is not specified, all subsequent input files are converted to this resulting format and concatenated together. The output file will contain an audio file header, unless <i>format=raw</i> is specified in the output format options.</p> <p>Input files may be converted in place by using the -p option. When -p is in effect, the format of each input file is modified according to the -f option to determine the output format. The existing files are then overwritten with the converted data.</p> <p>The file(1) command decodes and prints the audio data format of Sun audio files.</p>
OPTIONS	<p>-p <i>In Place</i>: The input files are individually converted to the format specified by the -f option and rewritten. If a target file is a symbolic link, the underlying file will be rewritten. The -o option may not be specified with -p.</p> <p>-F <i>Force</i>: This option forces audioconvert to ignore any file header for input files whose format is specified by the -i option. If -F is not specified, audioconvert ignores the -i option for input files that contain valid audio file headers.</p> <p>-f outfmt <i>Output Format</i>: This option is used to specify the file format and data encoding of the output file. Defaults for unspecified fields are derived from the input file format. Valid keywords and values are listed in the next section.</p> <p>-o outfile <i>Output File</i>: All input files are concatenated, converted to the output format, and written to the named output file. If -o and -p are not specified, the concatenated output is written to the standard output. The -p option may not be specified with -o.</p>

- i infmt** *Input Format:* This option is used to specify the data encoding of raw input files. Ordinarily, the input data format is derived from the audio file header. This option is required when converting audio data that is not preceded by a valid audio file header. If **-i** is specified for an input file that contains an audio file header, the input format string will be ignored, unless **-F** is present. The format specification syntax is the same as the **-f** output file format.
- Multiple input formats may be specified. An input format describes all input files following that specification, until a new input format is specified.
- file** *File Specification:* The named audio files are concatenated, converted to the output format, and written out. If no filename is present, or if the special filename '-' is specified, audio data is read from the standard input.
- ?** *Help:* Print a command line usage message.

FORMAT SPECIFICATION

The syntax for the input and output format specification is:

keyword=value[,keyword=value ...]

with no intervening whitespace. Unambiguous values may be used without the preceding *keyword=*.

- rate** The audio sampling rate is specified in samples per second. If a number is followed by the letter **k**, it is multiplied by 1000 (for example, 44.1k = 44100). Standard of the commonly used sample rates are: 8k, 16k, 32k, 44.1k, and 48k.
- channels** The number of interleaved channels is specified as an integer. The words **mono** and **stereo** may also be used to specify one and two channel data, respectively.
- encoding** This option specifies the digital audio data representation. Encodings determine precision implicitly (**ulaw** implies 8-bit precision) or explicitly as part of the name (for example, **linear16**). Valid encoding values are:
- ulaw** CCITT G.711 μ -law encoding. This is an 8-bit format primarily used for telephone quality speech.
- alaw** CCITT G.711 A-law encoding. This is an 8-bit format primarily used for telephone quality speech in Europe.
- linear8, linear16, linear32** Linear Pulse Code Modulation (PCM) encoding. The name identifies the number of bits of precision. **linear16** is typically used for high quality audio data.
- pcm** Same as **linear16**.

- g721** CCITT G.721 compression format. This encoding uses Adaptive Delta Pulse Code Modulation (ADPCM) with 4-bit precision. It is primarily used for compressing μ -law voice data (achieving a 2:1 compression ratio).
- g723** CCITT G.723 compression format. This encoding uses Adaptive Delta Pulse Code Modulation (ADPCM) with 3-bit precision. It is primarily used for compressing μ -law voice data (achieving an 8:3 compression ratio). The audio quality is similar to G.721, but may result in lower quality when used for non-speech data.

The following encoding values are also accepted as shorthand to set the sample rate, channels, and encoding:

voice Equivalent to **encoding=ulaw,rate=8k,channels=mono**.

cd Equivalent to **encoding=linear16,rate=44.1k,channels=stereo**.

dat Equivalent to **encoding=linear16,rate=48k,channels=stereo**.

format This option specifies the audio file format. Valid formats are:

sun Sun compatible file format (the default).

raw Use this format when reading or writing raw audio data (with no audio header), or in conjunction with an **offset** to import a foreign audio file format.

offset (*-i only*) Specify a byte offset to locate the start of the audio data. This option may be used to import audio data that contains an unrecognized file header.

EXAMPLES

Record voice data and compress it before storing it to a file:

```
example% audiorecord | audioconvert -f g721 > mydata.au
```

Concatenate two Sun format audio files, regardless of their data format, and output an 8-bit μ -law, 16 kHz, mono file:

```
example% audioconvert -f ulaw,rate=16k,mono -o outfile.au infile1 infile2
```

Convert a directory containing raw voice data files, in place, to Sun format (adds a file header to each file):

```
example% audioconvert -p -i voice -f sun *.au
```

SEE ALSO

audioplay(1), audiorecord(1), file(1)

NOTES

The algorithm used for converting multi-channel data to mono is implemented by simply summing the channels together. If the input data is perfectly in phase (as would be the case if a mono file is converted to stereo and back to mono), the resulting data may contain some distortion.

NAME	audioplay – play audio files
SYNOPSIS	audioplay [-iV] [-v vol] [-b bal] [-p speaker headphone line] [-d dev] [<i>file ...</i>]
AVAILABILITY	SUNWaudio
DESCRIPTION	<p>audioplay copies the named audio files (or the standard input if no filenames are present) to the audio device. If no input file is specified and standard input is a tty, the port, volume, and balance settings specified on the command line will be applied and the program will exit.</p> <p>The input files must contain a valid audio file header. The encoding information in this header is matched against the capabilities of the audio device and, if the data formats are incompatible, an error message is printed and the file is skipped. Compressed ADPCM (G.721) monaural audio data is automatically uncompressed before playing.</p> <p>Minor deviations in sampling frequency (that is, less than 1%) are ordinarily ignored. This allows, for instance, data sampled at 8012 Hz to be played on an audio device that only supports 8000 Hz. If the -V option is present, such deviations are flagged with warning messages.</p>
OPTIONS	<p>-i <i>Immediate</i>: If the audio device is unavailable (that is, another process currently has write access), audioplay ordinarily waits until it can obtain access to the device. When the -i option is present, audioplay prints an error message and exits immediately if the device is busy.</p> <p>-V <i>Verbose</i>: Print messages on the standard error when waiting for access to the audio device or when sample rate deviations are detected.</p> <p>-v vol <i>Volume</i>: The output volume is set to the specified value before playing begins, and is reset to its previous level when audioplay exits. The <i>vol</i> argument is an integer value between 0 and 100, inclusive. If this argument is not specified, the output volume remains at the level most recently set by any process.</p> <p>-b bal <i>Balance</i>: The output balance is set to the specified value before playing begins, and is reset to its previous level when audioplay exits. The <i>bal</i> argument is an integer value between -100 and 100, inclusive. A value of -100 indicates left balance, 0 middle, and 100 right. If this argument is not specified, the output balance remains at the level most recently set by any process.</p> <p>-p speaker headphone line <i>Output Port</i>: Select the built-in speaker, (the default), headphone jack, or line out as the destination of the audio output signal. If this argument is not specified, the output port will remain unchanged. <i>Not all audio adapters support all of the output ports. If the named port does not exist, an appropriate substitute will be used.</i></p> <p>-d dev <i>Device</i>: The <i>dev</i> argument specifies an alternate audio device to which output should be directed. If the -d option is not specified, the AUDIODEV</p>

environment variable is consulted (see below). Otherwise, **/dev/audio** is used as the default audio device.

file **File Specification:** Audio files named on the command line are played sequentially. If no filenames are present, the standard input stream (if it is not a tty) is played (it, too, must contain an audio file header). The special filename '-' may be used to read the standard input stream instead of a file. If a relative path name is supplied, the **AUDIOPATH** environment variable is consulted (see below).

-\? **Help:** Print a command line usage message.

ENVIRONMENT

AUDIODEV

The full path name of the audio device to write to, if no **-d** argument is supplied. If the **AUDIODEV** variable is not set, **/dev/audio** is used.

AUDIOPATH

A colon-separated list of directories in which to search for audio files whose names are given by relative pathnames. The current directory (".") may be specified explicitly in the search path. If the **AUDIOPATH** variable is not set, only the current directory will be searched.

SEE ALSO

Sparc Only
x86 Only

audioconvert(1), **audiorecord(1)**
audio(7), **audioamd(7)**, **dbri(7)**
sbpro(7)

BUGS

audioplay currently supports a limited set of audio format conversions. If the audio file is not in a format supported by the audio device, it must first be converted. For example, to convert to voice format on the fly, use the command:

example% audioconvert -f voice myfile | audioplay

The format conversion will not always be able to keep up with the audio output. If this is the case, you should convert to a temporary file before playing the data.

NAME	audiorecord – record an audio file
SYNOPSIS	audiorecord [-af] [-v vol] [-b bal] [-m monvol] [-p mic line internal-cd] [-c channels] [-s rate] [-e encoding] [-t time] [-i info] [-d dev] [<i>file</i>]
AVAILABILITY	SUNWaudio
DESCRIPTION	<p>audiorecord copies audio data from the audio device to a named audio file (or the standard output if no filename is present). If no output file is specified and standard output is a tty, the volume, balance, monitor volume, port, and audio format settings specified on the command line will be applied and the program will exit.</p> <p>By default, monaural audio data is recorded at 8 kHz and encoded in μ-law format. If the audio device supports additional configurations, the -c, -s, and -e options may be used to specify the data format. The output file is prefixed by an audio file header that identifies the format of the data encoded in the file.</p> <p>Recording begins immediately and continues until a SIGINT signal (for example, CTRL-C) is received. If the -t option is specified, audiorecord stops when the specified quantity of data has been recorded.</p> <p>If the audio device is unavailable (that is, another process currently has read access), audiorecord prints an error message and exits immediately.</p>
OPTIONS	<p>-a <i>Append</i>: Append the data on the end of the named audio file. The audio device must support the audio data format of the existing file.</p> <p>-f <i>Force</i>: When the -a flag is specified, the sample rate of the audio device must match the sample rate at which the original file was recorded. If the -f flag is also specified, sample rate differences are ignored, with a warning message printed on the standard error.</p> <p>-v vol <i>Volume</i>: The recording gain is set to the specified value before recording begins, and is reset to its previous level when audiorecord exits. The <i>vol</i> argument is an integer value between 0 and 100, inclusive. If this argument is not specified, the input volume will remain at the level most recently set by any process.</p> <p>-b bal <i>Balance</i>: The recording balance is set to the specified value before recording begins, and is reset to its previous level when audiorecord exits. The <i>bal</i> argument is an integer value between -100 and 100, inclusive. A value of -100 indicates left balance, 0 middle, and 100 right. If this argument is not specified, the input balance will remain at the level most recently set by any process.</p>

- m monvol** *Monitor Volume*: The input monitor volume is set to the specified value before recording begins, and is reset to its previous level when **audiorecord** exits. The *monval* argument is an integer value between 0 and 100, inclusive. A non-zero value allows a directly connected input source to be heard on the output speaker while recording is in-progress. If this argument is not specified, the monitor volume will remain at the level most recently set by any process.
- p mic | line | internal-cd**
Input Port: Select the **mic**, **line**, or **internal-cd** input as the source of the audio output signal. If this argument is not specified, the input port will remain unchanged. *Some systems will not support all possible input ports. If the named port does not exist, this option is ignored.*
- c channels**
Channels: Specify the number of audio channels (1 or 2). The value may be specified as an integer or as the string **mono** or **stereo**. The default value is **mono**.
- s rate** *Sample Rate*: Specify the sample rate, in samples per second. If a number is followed by the letter **k**, it is multiplied by 1000 (for example, 44.1k = 44100). The default sample rate is 8 kHz.
- e encoding**
Encoding: Specify the audio data encoding. This value may be one of **ulaw**, **alaw**, or **linear**. The default encoding is **ulaw**.
- t time** *Time*: The *time* argument specifies the maximum length of time to record. Time can be specified as a floating-point value, indicating the number of seconds, or in the form: *hh:mm:ss.dd*, where the hour and minute specifications are optional.
- i info** *Information*: The 'information' field of the output file header is set to the string specified by the *info* argument. This option cannot be specified in conjunction with the **-a** argument.
- d dev** *Device*: The *dev* argument specifies an alternate audio device from which input should be taken. If the **-d** option is not specified, the **AUDIODEV** environment variable is consulted (see below). Otherwise, **/dev/audio** is used as the default audio device.
- file* *File Specification*: The named audio file is rewritten (or appended). If no filename is present (and standard output is not a tty), or if the special filename '-' is specified, output is directed to the the standard output.
- \?** *Help*: Print a command line usage message.

ENVIRONMENT

- AUDIODEV** The full path name of the audio device to record from, if no **-d** argument is supplied. If the **AUDIODEV** variable is not set, **/dev/audio** is used.

SEE ALSO
SPARC Only
x86 Only

audioconvert(1), audioplay(1)
audio(7), audioamd(7), dbri(7)
sbpro(7)

NAME	awk – pattern scanning and processing language
SYNOPSIS	awk [<i>-f program-file</i>] [<i>-Fc</i>] [<i>'prog'</i>] [<i>parameters</i>] [<i>filename...</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	awk scans each input <i>filename</i> for lines that match any of a set of patterns specified in <i>prog</i> . The <i>prog</i> string must be enclosed in single quotes (') to protect it from the shell. For each pattern in <i>prog</i> there may be an associated action performed when a line of a <i>filename</i> matches the pattern. The set of pattern-action statements may appear literally as <i>prog</i> or in a file specified with the <i>-f program-file</i> option. Input files are read in order; if there are no files, the standard input is read. The file name '-' means the standard input.
OPTIONS	<p><i>-f program-file</i> awk uses the set of patterns it reads from <i>program-file</i>.</p> <p><i>-Fc</i> Use the character <i>c</i> as the field separator (FS) character. See the discussion of FS below.</p>
USAGE	
Input Lines	<p>Each input line is matched against the pattern portion of every pattern-action statement; the associated action is performed for each matched pattern. Any <i>filename</i> of the form <i>var=value</i> is treated as an assignment, not a filename, and is executed at the time it would have been opened if it were a filename. <i>Variables</i> assigned in this manner are not available inside a BEGIN rule, and are assigned after previously specified files have been read.</p> <p>An input line is normally made up of fields separated by white spaces. (This default can be changed by using the FS built-in variable or the <i>-Fc</i> option.) The default is to ignore leading blanks and to separate fields by blanks and/or tab characters. However, if FS is assigned a value that does not include any of the white spaces, then leading blanks are not ignored. The fields are denoted \$1, \$2, ...; \$0 refers to the entire line.</p>
Pattern-action Statements	<p>A pattern-action statement has the form:</p> <p style="text-align: center;"><i>pattern { action }</i></p> <p>Either pattern or action may be omitted. If there is no action, the matching line is printed. If there is no pattern, the action is performed on every input line. Pattern-action statements are separated by newlines or semicolons.</p> <p>Patterns are arbitrary Boolean combinations (!, , &&, and parentheses) of relational expressions and regular expressions. A relational expression is one of the following:</p> <p style="text-align: center;"><i>expression relop expression</i> <i>expression matchop regular_expression</i></p>

where a *relop* is any of the six relational operators in C, and a *matchop* is either `~` (contains) or `!` (does not contain). An *expression* is an arithmetic expression, a relational expression, the special expression

`var in array`

or a Boolean combination of these.

Regular expressions are as in **egrep**(1). In patterns they must be surrounded by slashes. Isolated regular expressions in a pattern apply to the entire line. Regular expressions may also occur in relational expressions. A pattern may consist of two patterns separated by a comma; in this case, the action is performed for all lines between the occurrence of the first pattern to the occurrence of the second pattern.

The special patterns **BEGIN** and **END** may be used to capture control before the first input line has been read and after the last input line has been read respectively. These keywords do not combine with any other patterns.

Built-in Variables

Built-in variables include:

FILENAME	name of the current input file
FS	input field separator regular expression (default blank and tab)
NF	number of fields in the current record
NR	ordinal number of the current record
OFMT	output format for numbers (default %6g)
OFS	output field separator (default blank)
ORS	output record separator (default new-line)
RS	input record separator (default new-line)

An action is a sequence of statements. A statement may be one of the following:

```

if ( expression ) statement [ else statement ]
while ( expression ) statement
do statement while ( expression )
for ( expression ; expression ; expression ) statement
for ( var in array ) statement
break
continue
{ [ statement ] ... }
expression      # commonly variable = expression
print [ expression-list ] [ >expression ]
printf format [ , expression-list ] [ >expression ]
next           # skip remaining patterns on this input line
exit [expr]  # skip the rest of the input; exit status is expr

```

Statements are terminated by semicolons, newlines, or right braces. An empty *expression-list* stands for the whole input line. Expressions take on string or numeric values as appropriate, and are built using the operators `+`, `-`, `*`, `/`, `%`, `^` and concatenation (indicated by a blank). The operators `++` `--` `+=` `-=` `*=` `/=` `%=` `^=` `>` `>=` `<` `<=` `==` `!=` `?:` are also

available in expressions. Variables may be scalars, array elements (denoted `x[i]`), or fields. Variables are initialized to the null string or zero. Array subscripts may be any string, not necessarily numeric; this allows for a form of associative memory. String constants are quoted ("`\"`"), with the usual C escapes recognized within.

The **print** statement prints its arguments on the standard output, or on a file if `>expression` is present, or on a pipe if `'|cmd'` is present. The output resulted from the print statement is terminated by the output record separator with each argument separated by the current output field separator. The **printf** statement formats its expression list according to the format (see **printf(3S)**).

The mathematical functions: **exp**, **log**, **sqrt**, are built-in.

Other built-in functions include:

- index(s, t)** returns the position in string *s* where string *t* first occurs, or 0 if it does not occur at all.
- int(s)** truncates *s* to an integer value. If *s* is not specified, `$0` is used.
- length(s)** returns the length of its argument taken as a string, or of the whole line if there is no argument.
- match(s, re)** returns the position in string *s* where the regular expression *re* occurs, or 0 if it does not occur at all.
- split(s, a, fs)** splits the string *s* into array elements *a[1]*, *a[2]*, ... *a[n]*, and returns *n*. The separation is done with the regular expression *fs* or with the field separator **FS** if *fs* is not given.
- sprintf(fmt, expr, expr, ...)** formats the expressions according to the **printf(3S)** format given by *fmt* and returns the resulting string.
- substr(s, m, n)** returns the *n*-character substring of *s* that begins at position *m*.

The input/output built-in function is:

- getline** sets **\$0** to the next input record from the current input file. **getline** returns 1 for successful input, 0 for end of file, and -1 for an error.

EXAMPLES

Print lines longer than 72 characters:

```
length > 72
```

Print first two fields in opposite order:

```
{ print $2, $1 }
```

Same, with input fields separated by comma and/or blanks and tabs:

```
BEGIN { FS = ",[ \t]*|[ \t]+" }
       { print $2, $1 }
```

Add up first column, print sum and average:

```
    { s += $1 }
  END { print "sum is", s, " average is", s/NR }
```

Print fields in reverse order:

```
  { for (i = NF; i > 0; --i) print $i }
```

Print all lines between start/stop pairs:

```
  /start/, /stop/
```

Print all lines whose first field is different from previous one:

```
  $1 != prev { print; prev = $1 }
```

Print a file, filling in page numbers starting at 5:

```
  /Page/ { $2 = n++; }
        { print }
```

Assuming this program is in a file named **prog**, the following command line prints the file **input** numbering its pages starting at 5: **awk -f prog n=5 input**.

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S/English).

SEE ALSO

egrep(1), **grep(1)**, **nawk(1)**, **sed(1)**, **printf(3S)**, **environ(5)**

NOTES

Input white space is not preserved on output if fields are involved.

There are no explicit conversions between numbers and strings. To force an expression to be treated as a number add 0 to it; to force it to be treated as a string concatenate the null string ("") to it.

NAME	banner – make posters
SYNOPSIS	banner <i>strings</i>
AVAILABILITY	SUNWesu
DESCRIPTION	banner prints its arguments (each up to 10 characters long) in large letters on the standard output.
SEE ALSO	echo (1)

NAME	basename, dirname – deliver portions of pathnames
SYNOPSIS	basename <i>string</i> [<i>suffix</i>] dirname <i>string</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	basename deletes any prefix ending in / and the <i>suffix</i> (if present in <i>string</i>) from <i>string</i> , and prints the result on the standard output. It is normally used inside substitution marks (. .) within shell procedures. The <i>suffix</i> is a pattern as defined on the ed (1) manual page. dirname delivers all but the last level of the path name in <i>string</i> .
EXAMPLES	The following example, invoked with the argument /home/sms/personal/mail sets the environment variable NAME to the file named mail and the environment variable MYMAILPATH to the string /home/sms/personal : <pre>example% NAME=`basename \$HOME/personal/mail` example% MYMAILPATH=`dirname \$HOME/personal/mail`</pre> This shell procedure, invoked with the argument /usr/src/bin/cat.c , compiles the named file and moves the output to cat in the current directory: <pre>example% cc \$1 example% mv a.out `basename \$1 .c`</pre>
SEE ALSO	ed (1), sh (1)

NAME	basename – display portions of pathnames
SYNOPSIS	<code>/usr/ucb/basename <i>string</i> [<i>suffix</i>]</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	basename deletes any prefix ending in '/' and the <i>suffix</i> , if present in <i>string</i> . It directs the result to the standard output, and is normally used inside substitution marks (` `) within shell procedures. The <i>suffix</i> is a pattern as defined on ed (1).
EXAMPLES	This shell procedure invoked with the argument <code>/usr/src/bin/cat.c</code> compiles the named file and moves the output to cat in the current directory: <pre>example% cc \$1 example% mv a.out `basename \$1 .c`</pre>
SEE ALSO	ed (1), sh (1)

NAME	bc – arbitrary precision arithmetic language
SYNOPSIS	bc [-c] [-l] [<i>filename...</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	<p>bc is an interactive processor for a language that resembles C but provides unlimited precision arithmetic. It takes input from any files given, then reads the standard input. bc is actually a preprocessor for the desk calculator program dc, which it invokes automatically unless the -c option is present. In this case the dc input is sent to the standard output instead.</p> <p>The syntax for bc programs is as follows: <i>L</i> means letter a–z, <i>E</i> means expression, <i>S</i> means statement.</p>
USAGE	
Comments	Are enclosed in /* and */ .
Names	<p>Simple variables: <i>L</i>.</p> <p>Array elements: <i>L</i> [<i>E</i>].</p> <p>The words ibase, obase, and scale.</p>
Other Operands	<p>Arbitrarily long numbers with optional sign and decimal point.</p> <p>(<i>E</i>)</p> <p>sqrt (<i>E</i>)</p> <p>length (<i>E</i>) Number of significant decimal digits.</p> <p>scale (<i>E</i>) Number of digits right of decimal point.</p> <p><i>L</i> (<i>E</i> , ... , <i>E</i>)</p>
Operators	<p>+ - * / % ^ (%is remainder; ^is power)</p> <p>++ — (prefix and postfix; apply to names)</p> <p>== <= >= != < ></p> <p>= == == ==* != !=% ==^</p>
Statements	<p><i>E</i></p> <p>{ <i>S</i>; ... ; <i>S</i> }</p> <p>if (<i>E</i>) <i>S</i></p> <p>while (<i>E</i>) <i>S</i></p> <p>for (<i>E</i>; <i>E</i>; <i>E</i>) <i>S</i></p> <p>null statement</p> <p>break</p> <p>quit</p>

Function Definitions

```
define L ( L ,... L ) {
    auto L ,... L
    S ;... S
    return ( E )
}
```

Functions in -l Math Library

```
s(x)    sine
c(x)    cosine
e(x)    exponential
l(x)    log
a(x)    arctangent
j(n,x)  Bessel function
```

All function arguments are passed by value.

The value of a statement that is an expression is printed unless the main operator is an assignment. Either semicolons or new-lines may separate statements. Assignment to **scale** influences the number of digits to be retained on arithmetic operations in the manner of **dc**. Assignments to **ibase** or **obase** set the input and output number radix respectively.

The same letter may be used as an array, a function, and a simple variable simultaneously. All variables are global to the program. **auto** variables are pushed down during function calls. When using arrays as function arguments or defining them as automatic variables, empty square brackets must follow the array name.

OPTIONS

```
-c      Compile only. The output is sent to the standard output.
-l      Argument stands for the name of an arbitrary precision math library.
```

EXAMPLES

Defines a function to compute an approximate value of the exponential function:

```
scale = 20
define e(x){
    auto a, b, c, i, s
    a = 1
    b = 1
    s = 1
    for(i=1; 1==1; i++){
        a = a*x
        b = b*i
        c = a/b
        if(c == 0) return(s)
        s = s+c
    }
}
```

Prints approximate values of the exponential function of the first ten integers:

```
for(i=1; i<=10; i++) e(i)
```

FILES /usr/bin/dc desk calculator proper
 /usr/lib/lib.b mathematical library

SEE ALSO dc(1)

NOTES The **bc** command does not recognize the logical operators **&&** and **|** .
 The **for** statement must have all three expressions (*E*'s).
 The **quit** statement is interpreted when read, not when executed.

NAME	bdiff – big diff
SYNOPSIS	bdiff <i>filename1 filename2</i> [<i>n</i>] [-s]
AVAILABILITY	SUNWesu
DESCRIPTION	<p>bdiff is used in a manner analogous to diff to find which lines in <i>filename1</i> and <i>filename2</i> must be changed to bring the files into agreement. Its purpose is to allow processing of files too large for diff. If <i>filename1</i> (<i>filename2</i>) is -, the standard input is read.</p> <p>bdiff ignores lines common to the beginning of both files, splits the remainder of each file into <i>n</i>-line segments, and invokes diff on corresponding segments. If both optional arguments are specified, they must appear in the order indicated above.</p> <p>The output of bdiff is exactly that of diff, with line numbers adjusted to account for the segmenting of the files (that is, to make it look as if the files had been processed whole). Note: Because of the segmenting of the files, bdiff does not necessarily find a smallest sufficient set of file differences.</p>
OPTIONS	<p><i>n</i> The number of line segments. The value of <i>n</i> is 3500 by default. If the optional third argument is given and it is numeric, it is used as the value for <i>n</i>. This is useful in those cases in which 3500-line segments are too large for diff, causing it to fail.</p> <p>-s Specifies that no diagnostics are to be printed by bdiff (silent option). Note: However, this does not suppress possible diagnostic messages from diff, which bdiff calls.</p>
FILES	/tmp/bd?????
SEE ALSO	diff (1)
DIAGNOSTICS	Use help for explanations.

NAME	biff – give notice of incoming mail messages
SYNOPSIS	<code>/usr/ucb/biff [y n]</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>biff turns mail notification on or off for the terminal session. With no arguments, biff displays the current notification status for the terminal.</p> <p>If notification is allowed, the terminal rings the bell and displays the header and the first few lines of each arriving mail message. biff operates asynchronously. For synchronized notices, use the MAIL variable of sh(1) or the mail variable of csh(1).</p> <p>A 'biff y' command can be included in your <code>~/.login</code> or <code>~/.profile</code> file for execution when you log in.</p>
OPTIONS	<p>y Allow mail notification for the terminal.</p> <p>n Disable notification for the terminal.</p>
FILES	<code>~/.login</code> <code>~/.profile</code>
SEE ALSO	cs h(1), mail (1), sh (1)

NAME	break, continue – shell built-in functions to escape from or advance within a controlling while, for, foreach, or until loop
SYNOPSIS	
sh	break [<i>n</i>] continue [<i>n</i>]
csch	break continue
ksh	† break [<i>n</i>] † continue [<i>n</i>]
DESCRIPTION	
sh	break exits from the enclosing for or while loop, if any. If <i>n</i> is specified, break <i>n</i> levels. continue resumes the next iteration of the enclosing for or while loop. If <i>n</i> is specified, resume at the <i>n</i> -th enclosing loop.
csch	break resumes execution after the end of the nearest enclosing foreach or while loop. The remaining commands on the current line are executed. This allows multilevel breaks to be written as a list of break commands, all on one line. continue continues execution of the next iteration of the nearest enclosing while or foreach loop.
ksh	break exits from the enclosed for , while , until , or select loop, if any. If <i>n</i> is specified then break <i>n</i> levels. continue resumes the next iteration of the enclosed for , while , until , or select loop. If <i>n</i> is specified then resume at the <i>n</i> -th enclosed loop. On this man page, ksh (1) commands that are preceded by one or two † (daggers) are treated specially in the following ways: 1. Variable assignment lists preceding the command remain in effect when the command completes. 2. I/O redirections are processed after variable assignments. 3. Errors cause a script that contains them to abort. 4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.
SEE ALSO	csch (1), exit (1), for (1), foreach (1), ksh (1), select (1), sh (1), until (1), while (1)

NAME	cal – display a calendar
SYNOPSIS	cal [[<i>month</i>] <i>year</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	cal prints a calendar for the specified year. If a month is also specified, a calendar just for that month is printed. If neither is specified, a calendar for the present month is printed. The calendar produced is that for England and the United States.
OPTIONS	<i>month</i> The <i>month</i> is a number between 1 and 12. <i>year</i> The <i>year</i> can be between 1 and 9999.
SEE ALSO	calendar(1)
NOTES	An unusual calendar is printed for September 1752. That is the month 11 days were skipped to make up for lack of leap year adjustments. To see this calendar, type: cal 9 1752 The command cal 83 refers to the year 83, not 1983. The year is always considered to start in January even though this is historically naive.

NAME	calendar – reminder service
SYNOPSIS	calendar [-]
AVAILABILITY	SUNWesu
DESCRIPTION	<p>calendar consults the file calendar in the current directory and prints out lines that contain today's or tomorrow's date anywhere in the line. Most reasonable month-day dates such as Aug. 24, august 24, 8/24, and so forth, are recognized, but not 24 August or 24/8. On weekends "tomorrow" extends through Monday. calendar can be invoked regularly by using the crontab(1) or at(1) commands.</p> <p>When an argument is present, calendar does its job for every user who has a file calendar in his or her login directory and sends them any positive results by mail(1). Normally this is done daily by facilities in the UNIX operating system (see cron(1M)).</p> <p>If the environment variable DATMSK is set, calendar will use its value as the full path name of a template file containing format strings. The strings consist of field descriptors and text characters and are used to provide a richer set of allowable date formats in different languages by appropriate settings of the environment variable LANG or LC_TIME (see environ(5)). (See date(1) for the allowable list of field descriptors.)</p>
EXAMPLES	<p>The following example shows the possible contents of a template:</p> <p style="padding-left: 40px;">%B %eth of the year %Y</p> <p>%B represents the full month name, %e the day of month and %Y the year (4 digits). If DATMSK is set to this template, the following calendar file would be valid:</p> <p style="padding-left: 40px;">March 7th of the year 1989 < Reminder ></p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of calendar for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how calendar behaves.</p> <p>LC_CTYPE</p> <p>Determines how calendar handles characters. When LC_CTYPE is set to a valid value, calendar can display and handle text and filenames containing valid characters for that locale. calendar can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. calendar can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p> <p>LC_TIME</p> <p>Determines how calendar handles date and time formats. In the "C" locale, date and time handling follows the U.S. rules.</p>

FILES /etc/passwd

/tmp/cal*

/usr/lib/calprog program used to figure out today's and tomorrow's dates

SEE ALSO at(1), crontab(1), date(1), mail(1), cron(1M), environ(5)

NOTES Appropriate lines beginning with white space will not be printed.

Your calendar must be public information for you to get reminder service.

calendar's extended idea of "tomorrow" does not account for holidays.

NAME	case, switch, select – shell built-in functions to choose from among a list of actions
SYNOPSIS	
sh	case <i>word</i> in [(<i>pattern</i> [<i>pattern</i>] ...) <i>actions</i> ;;] ... esac
csh	switch (<i>expression</i>) case <i>comparison1</i> : ... breaksw case <i>comparison2</i> : ... breaksw ... default : endsw
ksh	case <i>word</i> in [[(<i>pattern</i> [<i>pattern</i>] ...) <i>actions</i> ;;] ... esac select <i>identifier</i> [in <i>word</i> ...]; do <i>list</i> ; done
DESCRIPTION	
sh	A case command executes the <i>actions</i> associated with the first <i>pattern</i> that matches <i>word</i> . The form of the patterns is the same as that used for file-name generation except that a slash, a leading dot, or a dot immediately following a slash need not be matched explicitly.
csh	The c-shell uses the switch statement, in which each <i>comparison</i> is successively matched, against the specified <i>expression</i> , which is first command and filename expanded. The file metacharacters *, ? and [...] may be used in the case comparison, which are variable expanded. If none of the comparisons match before a “default” comparison is found, execution begins after the default comparison. Each case statement and the default statement must appear at the beginning of a line. The command breaksw continues execution after the endsw . Otherwise control falls through subsequent case and default statements as with C. If no comparison matches and there is no default, execution continues after the endsw . case comparison : A compared-expression in a switch statement. default : If none of the preceding <i>comparisons</i> match <i>expression</i> , then this is the default case in a switch statement. The default should come after all case comparisons. Any remaining commands on the command line are first executed. breaksw exits from a switch , resuming after the endsw .
ksh	A case command executes the <i>actions</i> associated with the first <i>pattern</i> that matches <i>word</i> . The form of the patterns is the same as that used for file-name generation (see File Name Generation in ksh(1)).

A **select** command prints to standard error (file descriptor 2), the set of *words*, each preceded by a number. If **in** *word* . . . is omitted, then the positional parameters are used instead. The **PS3** prompt is printed and a line is read from the standard input. If this line consists of the number of one of the listed *words*, then the value of the variable *identifier* is set to the *word* corresponding to this number. If this line is empty the selection list is printed again. Otherwise the value of the variable *identifier* is set to **NULL**. The contents of the line read from standard input is saved in the shell variable **REPLY**. The *list* is executed for each selection until a **break** or *end-of-file* is encountered. If the **REPLY** variable is set to **NULL** by the execution of *list*, then the selection list is printed before displaying the **PS3** prompt for the next selection.

SEE ALSO **break(1)**, **cs(1)**, **ksh(1)**, **sh(1)**

NAME	cat – concatenate and display files
SYNOPSIS	cat [-bnsuvet] <i>filename</i> . . .
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>cat reads each <i>filename</i> in sequence and writes it on the standard output. Thus:</p> <p style="padding-left: 40px;">example% cat filename</p> <p>prints <i>filename</i> on your terminal, and:</p> <p style="padding-left: 40px;">example% cat filename1 filename2 >filename3</p> <p>concatenates <i>filename1</i> and <i>filename2</i>, and writes the results in <i>filename3</i>.</p> <p>If no input file is given, or if the argument - is encountered, cat reads from the standard input file.</p>
OPTIONS	<p>-b Number the lines, as -n, but omit the line numbers from blank lines.</p> <p>-n Precede each line output with its line number.</p> <p>-u The output is not buffered. (The default is buffered output.)</p> <p>-s cat is silent about non-existent files.</p> <p>-v Non-printing characters (with the exception of tabs, new-lines and form-feeds) are printed visibly. ASCII control characters (octal 000 – 037) are printed as ^n, where <i>n</i> is the corresponding ASCII character in the range octal 100 – 137 (@, A, B, C, . . . , X, Y, Z, [, \,], ^, and _); the DEL character (octal 0177) is printed ?^. Other non-printable characters are printed as M-x, where <i>x</i> is the ASCII character specified by the low-order seven bits.</p> <p>When used with the -v option, the following options may be used:</p> <p>-t Tabs will be printed as ^T's and formfeeds to be printed as ^L's.</p> <p>-e A \$ character will be printed at the end of each line (prior to the new-line).</p> <p>The -t and -e options are ignored if the -v option is not specified.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of cat for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how cat behaves.</p>

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

cp(1), **pg(1)**, **pr(1)**, **environ(5)**

NOTES

Redirecting the output of **cat** onto one of the files being read will cause the loss of the data originally in the file being read. For example,

example% cat filename1 filename2 >filename1

causes the original data in **filename1** to be lost.

NAME	cc – C compiler																		
SYNOPSIS	/usr/ucb/cc [<i>options</i>]																		
AVAILABILITY	SUNWscpu																		
DESCRIPTION	<p>/usr/ucb/cc is the interface to the BSD Compatibility Package C compiler. It is a script that looks for the link /usr/ccs/bin/ucbcc to the C compiler. /usr/ccs/bin/ucbcc is available only with the SPROcc package, whose default location is /opt/SUNWspro.</p> <p>/usr/ucb/cc is identical to /usr/ccs/bin/ucbcc, except that BSD headers are used and BSD libraries are linked <i>before</i> base libraries. The /opt/SUNWspro/man/man1/acc.1 man page is available only with the SPROcc package.</p>																		
OPTIONS	<p>/usr/ucb/cc accepts the same options as /usr/ccs/bin/ucbcc, with the following exceptions:</p> <p>-I<i>dir</i> Search <i>dir</i> for included files whose names do not begin with a slash (/) prior to searching the usual directories. The directories for multiple -I options are searched in the order specified. The preprocessor first searches for #include files in the directory containing <i>sourcefile</i>, and then in directories named with -I options (if any), then /usr/ucbinclude, and finally, in /usr/include.</p> <p>-L<i>dir</i> Add <i>dir</i> to the list of directories searched for libraries by /usr/ccs/bin/ucbcc. This option is passed to /usr/ccs/bin/ld and /usr/ccs/lib. Directories specified with this option are searched before /usr/ucb/lib and /usr/lib.</p> <p>-Y P, <i>dir</i> Change the default directory used for finding libraries.</p>																		
FILES	<table border="0"> <tr> <td>/usr/ccs/bin/ld</td> <td>link editor</td> </tr> <tr> <td>/usr/lib/libc</td> <td>C library</td> </tr> <tr> <td>/usr/ucbinclude</td> <td>BSD Compatibility directory for header files</td> </tr> <tr> <td>/usr/ucb/lib</td> <td>BSD Compatibility directory for libraries</td> </tr> <tr> <td>/usr/ucb/lib/libucb</td> <td>BSD Compatibility C library</td> </tr> <tr> <td>/usr/lib/libsocket</td> <td>library containing socket routines</td> </tr> <tr> <td>/usr/lib/libnsl</td> <td>library containing network functions</td> </tr> <tr> <td>/usr/lib/libelf</td> <td>library containing routines to process ELF object files</td> </tr> <tr> <td>/usr/lib/libaio</td> <td>library containing asynchronous I/O routines</td> </tr> </table>	/usr/ccs/bin/ld	link editor	/usr/lib/libc	C library	/usr/ucbinclude	BSD Compatibility directory for header files	/usr/ucb/lib	BSD Compatibility directory for libraries	/usr/ucb/lib/libucb	BSD Compatibility C library	/usr/lib/libsocket	library containing socket routines	/usr/lib/libnsl	library containing network functions	/usr/lib/libelf	library containing routines to process ELF object files	/usr/lib/libaio	library containing asynchronous I/O routines
/usr/ccs/bin/ld	link editor																		
/usr/lib/libc	C library																		
/usr/ucbinclude	BSD Compatibility directory for header files																		
/usr/ucb/lib	BSD Compatibility directory for libraries																		
/usr/ucb/lib/libucb	BSD Compatibility C library																		
/usr/lib/libsocket	library containing socket routines																		
/usr/lib/libnsl	library containing network functions																		
/usr/lib/libelf	library containing routines to process ELF object files																		
/usr/lib/libaio	library containing asynchronous I/O routines																		
SEE ALSO	ld(1) , a.out(4)																		
NOTES	The -Y P, <i>dir</i> option may have unexpected results, and should not be used.																		

NAME	cd, chdir, pushd, popd, dirs – shell built-in functions to change the current working directory
SYNOPSIS	
sh	cd [<i>argument</i>]
csh	cd [<i>dir</i>] chdir [<i>dir</i>] pushd [+ <i>n</i> <i>dir</i>] popd [+ <i>n</i>] dirs [- <i>l</i>]
ksh	cd [<i>arg</i>] cd <i>old new</i>
DESCRIPTION	
sh	<p>cd changes the current directory to <i>argument</i>. The shell parameter HOME is the default <i>argument</i>. The shell parameter CDPATH defines the search path for the directory containing <i>argument</i>. Alternative directory names are separated by a colon (:). The default path is <null> (specifying the current directory). Note: The current directory is specified by a null path name, which can appear immediately after the equal sign or between the colon delimiters anywhere else in the path list. If <i>argument</i> begins with '/', '.', or '..', the search path is not used. Otherwise, each directory in the path is searched for <i>argument</i>. cd must have execute (search) permission in <i>argument</i>.</p> <p>Because a new process is created to execute each command, cd would be ineffective if it were written as a normal command; therefore, it is recognized by and is internal to the shell. (See pwd(1), sh(1), and chdir(2)).</p>
csh	<p>If <i>dir</i> is not specified, cd uses the value of shell parameter HOME as the new working directory. If <i>dir</i> specifies a complete path starting with '/', '.', or '..', <i>dir</i> becomes the new working directory. If neither case applies, cd tries to find the designated directory relative to one of the paths specified by the CDPATH shell variable. CDPATH has the same syntax as, and similar semantics to, the PATH shell variable. cd must have execute (search) permission in <i>dir</i>. Because a new process is created to execute each command, cd would be ineffective if it were written as a normal command; therefore, it is recognized by and is internal to the C-shell. (See pwd(1), sh(1), and chdir(2)).</p> <p>chdir changes the shell's working directory to directory <i>dir</i>. If no argument is given, change to the home directory of the user. If no argument is given, change to the home directory of the user. If <i>dir</i> is a relative pathname not found in the current directory, check for it in those directories listed in the cdpath variable. If <i>dir</i> is the name of a shell variable whose value starts with a /, change to the directory named by that value.</p> <p>pushd will push a directory onto the directory stack. With no arguments, exchange the top two elements.</p> <p>+<i>n</i> Rotate the <i>n</i>'th entry to the top of the stack and cd to it.</p>

dir Push the current working directory onto the stack and change to *dir*.

popd pops the directory stack and **cd** to the new top directory. The elements of the directory stack are numbered from 0 starting at the top.

+n Discard the *n*'th entry in the stack.

dirs will print the directory stack, most recent to the left; the first directory shown is the current directory. With the **-l** argument, produce an unabbreviated printout; use of the **~** notation is suppressed.

ksh The **cd** command can be in either of two forms. In the first form it changes the current directory to *arg*. If *arg* is **-** the directory is changed to the previous directory. The shell variable **HOME** is the default *arg*. The variable **PWD** is set to the current directory. The shell variable **CDPATH** defines the search path for the directory containing *arg*. Alternative directory names are separated by a colon (:). The default path is **<null>** (specifying the current directory). Note that the current directory is specified by a **NULL** path name, which can appear immediately after the equal sign or between the colon delimiters anywhere else in the path list. If *arg* begins with a **'/'**, **'.'**, or **'..'**, then the search path is not used. Otherwise, each directory in the path is searched for *arg*.

The second form of **cd** substitutes the string *new* for the string *old* in the current directory name, **PWD** and tries to change to this new directory.

The **cd** command may not be executed by **rksh**. Because a new process is created to execute each command, **cd** would be ineffective if it were written as a normal command; therefore, it is recognized by and is internal to the Korn shell. (See **pwd(1)**, **sh(1)**, and **chdir(2)**).

SEE ALSO **cs(1)**, **ksh(1)**, **pwd(1)**, **sh(1)**, **chdir(2)**

NAME	checknr – check nroff and troff input files; report possible errors
SYNOPSIS	checknr [-fs] [-a .x1 .y1 .x2 .y2xn .yn] [-c .x1 .x2 .x3xn] [<i>filename ...</i>]
AVAILABILITY	SUNWdoc
DESCRIPTION	<p>checknr checks a list of nroff(1) or troff(1) input files for certain kinds of errors involving mismatched opening and closing delimiters and unknown commands. If no files are specified, checknr checks the standard input. Delimiters checked are:</p> <ul style="list-style-type: none"> • Font changes using \fx ... \fP. • Size changes using \sx ... \s0. • Macros that come in open ... close forms, for example, the .TS and .TE macros which must always come in pairs. <p>checknr knows about the ms(5) and me(5) macro packages.</p> <p>checknr is intended to be used on documents that are prepared with checknr in mind. It expects a certain document writing style for \f and \s commands, in that each \fx must be terminated with \fP and each \sx must be terminated with \s0. While it will work to directly go into the next font or explicitly specify the original font or point size, and many existing documents actually do this, such a practice will produce complaints from checknr. Since it is probably better to use the \fP and \s0 forms anyway, you should think of this as a contribution to your document preparation style.</p>
OPTIONS	<p>-f Ignore \f font changes.</p> <p>-s Ignore \s size changes.</p> <p>-a .x1 .y1 ... Add pairs of macros to the list. The pairs of macros are assumed to be those (such as .DS and .DE) that should be checked for balance. The -a option must be followed by groups of six characters, each group defining a pair of macros. The six characters are a period, the first macro name, another period, and the second macro name. For example, to define a pair .BS and .ES, use '-a.BS.ES'</p> <p>-c .x1 ... Define commands which checknr would otherwise complain about as undefined.</p>
SEE ALSO	eqn (1), nroff (1), troff (1), me (5), ms (5)
BUGS	There is no way to define a one-character macro name using the -a option.

NAME	chgrp – change the group ownership of a file
SYNOPSIS	chgrp [-f] [-h] [-R] <i>gid filename</i> ...
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>chgrp changes the group ID (GID) of the <i>files</i> given as arguments to <i>group</i>. The group may be either a decimal group ID or a group name found in the group ID file, <i>/etc/group</i>.</p> <p>You must be the owner of the file, or be the super-user to use this command. If you are not the super-user, and successfully change the group, the set-group-id bit is cleared.</p> <p>The operating system has a configuration option { _POSIX_CHOWN_RESTRICTED }, to restrict ownership changes. When this option is in effect, the owner of the file may change the group of the file only to a group to which the owner belongs. Only the super-user can arbitrarily change owner IDs whether this option is in effect or not.</p>
OPTIONS	<p>-f Force. Do not report errors.</p> <p>-h If the file is a symbolic link, change the group of the symbolic link. Without this option, the group of the file referenced by the symbolic link is changed.</p> <p>-R Recursive. chgrp descends through the directory, and any subdirectories, setting the specified group ID as it proceeds. When a symbolic link is encountered, the group of the target file is changed (unless the -h option is specified), but no recursion takes place.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of chgrp for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how chgrp behaves.</p> <p>LC_CTYPE Determines how chgrp handles characters. When LC_CTYPE is set to a valid value, chgrp can display and handle text and filenames containing valid characters for that locale. chgrp can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. chgrp can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p> <p>LC_MESSAGES Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).</p>

FILES**/etc/group****SEE ALSO****chmod(1), chown(1), id(1M), chown(2), group(4), passwd(4), environ(5)**

NAME	chkey – change user’s secure RPC key pair
SYNOPSIS	chkey [-p] [-s nisplus nis files]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>chkey is used to change a user’s secure RPC public key and secret key pair. chkey prompts for the old secure-rpc password and verifies that it is correct by decrypting the secret key. If the user has not already keylogged in, chkey registers the secret key with the local keyserv(1M) daemon. If the secure-rpc password does not match the login password, chkey prompts for the login password. chkey uses the login password to encrypt the user’s secret Diffie-Hellman (192 bit) cryptographic key.</p> <p>chkey ensures that the login password and the secure-rpc password are kept the same, thus enabling password shadowing, (see shadow(4)).</p> <p>The key pair can be stored in the /etc/publickey file, (see publickey(4)), NIS publickey map or NIS+ cred.org_dir table. If a new secret key is generated, it will be registered with the local keyserv(1M) daemon.</p> <p>If the source of the publickey is not specified with the -s option, chkey consults the publickey entry in the name service switch configuration file (see nsswitch.conf(4)). If the publickey entry specifies one and only one source, then chkey will change the key in the specified name service. However, if multiple name services are listed, chkey can not decide which source to update and will display an error message. The user should specify the source explicitly with the -s option.</p> <p>Non root users are not allowed to change their key pair in the files database.</p>
OPTIONS	<p>-p Re-encrypt the existing secret key with the user’s login password.</p> <p>-s nisplus Update the NIS+ database.</p> <p>-s nis Update the NIS database.</p> <p>-s files Update the files database.</p>
FILES	<p>/etc/nsswitch.conf</p> <p>/etc/publickey</p>
SEE ALSO	keylogin (1), keylogout (1), keyserv (1M), newkey (1M), nisaddcred (1M), nsswitch.conf (4), publickey (4), shadow (4)

NAME	chmod – change the permissions mode of a file																																		
SYNOPSIS	chmod [-fR] <i>mode filename...</i> chmod [ugoa] { + - = } [rxwlsStTugo] <i>filename...</i>																																		
AVAILABILITY	SUNWcsu																																		
DESCRIPTION	<p>chmod changes or assigns the mode of a file. The mode of a file specifies its permissions and other attributes. The mode may be absolute or symbolic.</p> <p>An absolute <i>mode</i> is specified using octal numbers:</p> <p style="padding-left: 40px;">chmod <i>nnnn filename...</i></p> <p>where <i>n</i> is a number from 0 to 7. An absolute mode is constructed from the OR of any of the following modes:</p> <table border="0" style="margin-left: 40px;"> <tr> <td style="padding-right: 20px;">4000</td> <td>Set user ID on execution.</td> </tr> <tr> <td>20#0</td> <td>Set group ID on execution if # is 7, 5, 3, or 1.</td> </tr> <tr> <td></td> <td>Enable mandatory locking if # is 6, 4, 2, or 0.</td> </tr> <tr> <td></td> <td>For directories, files are created with BSD semantics for propagation of the group ID. With this option, files inherit the group ID of the directory. It may be cleared only by using symbolic mode.</td> </tr> <tr> <td>1000</td> <td>Turn on sticky bit. See chmod(2).</td> </tr> <tr> <td>0400</td> <td>Allow read by owner.</td> </tr> <tr> <td>0200</td> <td>Allow write by owner.</td> </tr> <tr> <td>0100</td> <td>Allow execute (search in directory) by owner.</td> </tr> <tr> <td>0070</td> <td>Allow read, write, and execute (search) by group.</td> </tr> <tr> <td>0007</td> <td>Allow read, write, and execute (search) by others.</td> </tr> </table> <p>A symbolic <i>mode</i> is specified in the following format:</p> <p style="padding-left: 40px;">chmod [<i>who</i>] <i>operator</i> [<i>permission(s)</i>] <i>filename...</i></p> <p><i>who</i> is zero or more of the characters u, g, o, and a specifying whose permissions are to be changed or assigned:</p> <table border="0" style="margin-left: 40px;"> <tr> <td style="padding-right: 20px;">u</td> <td>user's permissions</td> </tr> <tr> <td>g</td> <td>group's permissions</td> </tr> <tr> <td>o</td> <td>others' permissions</td> </tr> <tr> <td>a</td> <td>all permissions (user, group, and other)</td> </tr> </table> <p>If <i>who</i> is omitted, it defaults to a, but the setting of the file creation mask (see umask in sh(1) or csh(1) for more information) is taken into account. When <i>who</i> is omitted, chmod will not override the restrictions of your user mask.</p> <p><i>operator</i> is one of +, -, or =, signifying how permissions are to be changed:</p> <table border="0" style="margin-left: 40px;"> <tr> <td style="padding-right: 20px;">+</td> <td>Add permissions.</td> </tr> <tr> <td>-</td> <td>Take away permissions.</td> </tr> <tr> <td>=</td> <td>Assign permissions absolutely.</td> </tr> </table>	4000	Set user ID on execution.	20#0	Set group ID on execution if # is 7, 5, 3, or 1.		Enable mandatory locking if # is 6, 4, 2, or 0.		For directories, files are created with BSD semantics for propagation of the group ID. With this option, files inherit the group ID of the directory. It may be cleared only by using symbolic mode.	1000	Turn on sticky bit. See chmod(2) .	0400	Allow read by owner.	0200	Allow write by owner.	0100	Allow execute (search in directory) by owner.	0070	Allow read, write, and execute (search) by group.	0007	Allow read, write, and execute (search) by others.	u	user's permissions	g	group's permissions	o	others' permissions	a	all permissions (user, group, and other)	+	Add permissions.	-	Take away permissions.	=	Assign permissions absolutely.
4000	Set user ID on execution.																																		
20#0	Set group ID on execution if # is 7, 5, 3, or 1.																																		
	Enable mandatory locking if # is 6, 4, 2, or 0.																																		
	For directories, files are created with BSD semantics for propagation of the group ID. With this option, files inherit the group ID of the directory. It may be cleared only by using symbolic mode.																																		
1000	Turn on sticky bit. See chmod(2) .																																		
0400	Allow read by owner.																																		
0200	Allow write by owner.																																		
0100	Allow execute (search in directory) by owner.																																		
0070	Allow read, write, and execute (search) by group.																																		
0007	Allow read, write, and execute (search) by others.																																		
u	user's permissions																																		
g	group's permissions																																		
o	others' permissions																																		
a	all permissions (user, group, and other)																																		
+	Add permissions.																																		
-	Take away permissions.																																		
=	Assign permissions absolutely.																																		

Unlike other symbolic operations, = has an absolute effect in that it resets all other bits. Omitting *permission(s)* is useful only with = to take away all permissions.

permission(s) is any compatible combination of the following letters:

r	read permission
w	write permission
x	execute permission
l	mandatory locking
s	user or group set-ID
S	undefined bit-state (the set-user-ID bit is on and the user execution bit is off)
t	sticky bit
T	the 1000 bit is turned on, and execution is off (undefined bit-state)
u,g,o	indicate that <i>permission</i> is to be taken from the current user, group or other mode respectively.

Permissions to a file may vary depending on your user identification number (UID) or group identification number (GID). Permissions are described in three sequences each having three characters:

User	Group	Other
rwX	rwX	rwX

This example (user, group, and others all have permission to read, write, and execute a given file) demonstrates two categories for granting permissions: the access class and the permissions themselves.

Multiple symbolic modes separated by commas may be given, though no spaces may intervene between these modes. Operations are performed in the order given. Multiple symbolic letters following a single operator cause the corresponding operations to be performed simultaneously.

The letter **s** is only meaningful with **u** or **g**, and **t** only works with **u**.

Mandatory file and record locking (**l**) refers to a file's ability to have its reading or writing permissions locked while a program is accessing that file. It is not possible to permit group execution and enable a file to be locked on execution at the same time. In addition, it is not possible to turn on the set-group-ID bit and enable a file to be locked on execution at the same time. The following examples, therefore, are invalid and elicit error messages:

```
chmod g+x,+l filename
chmod g+s,+l filename
```

Only the owner of a file or directory (or the super-user) may change that file's or directory's mode. Only the super-user may set the sticky bit on a non-directory file. If you are not super-user, **chmod** will mask the sticky-bit but will not return an error. In order to turn on a file's set-group-ID bit, your own group ID must correspond to the file's and group execution must be set.

OPTIONS

- f** Force. **chmod** will not complain if it fails to change the mode of a file.
- R** Recursively descend through directory arguments, setting the mode for each file as described above. When symbolic links are encountered, the mode of the target file is changed, but no recursion takes place.

EXAMPLES

Deny execute permission to everyone:

```
example% chmod a-x filename
```

Allow read permission to everyone:

```
example% chmod 444 filename
```

Make a file readable and writable by the group and others:

```
example% chmod go+rw filename
```

```
example% chmod 066 filename
```

Cause a file to be locked during access:

```
example% chmod +l filename
```

Allow everyone to read, write, and execute the file and turn on the set group-ID.

```
example% chmod =rwx,g+s filename
```

```
example% chmod 2777 filename
```

Absolute changes don't work for the set-group-ID bit of a directory. You must use **g+s** or **g-s**.

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO | **ls(1), chmod(2), environ(5)**

NOTES | **chmod** permits you to produce useless modes so long as they are not illegal (for instance, making a text file executable). **chmod** does not check the file type to see if mandatory locking is meaningful.

If the filesystem is mounted with the *nosuid* option, *setuid* execution is not allowed.

NAME	chown – change owner of file
SYNOPSIS	chown [-fhR] <i>owner filename</i> ...
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>chown changes the owner of the <i>files</i> to <i>owner</i>. The <i>owner</i> may be either a decimal user ID or a login name found in <i>/etc/passwd</i> file.</p> <p>If chown is invoked by other than the super-user, the set-user-ID bit of the file mode, 04000, is cleared.</p> <p>Only the owner of a file (or the super-user) may change the owner of that file.</p> <p>The operating system has a configuration option {_POSIX_CHOWN_RESTRICTED}, to restrict ownership changes. When this option is in effect the owner of the file is prevented from changing the owner ID of the file. Only the super-user can arbitrarily change owner IDs whether this option is in effect or not.</p>
OPTIONS	<p>-f Do not report errors.</p> <p>-h If the file is a symbolic link, change the owner of the symbolic link. Without this option, the owner of the file referenced by the symbolic link is changed.</p> <p>-R Recursive. chown descends through the directory, and any subdirectories, setting the ownership ID as it proceeds. When a symbolic link is encountered, the owner of the target file is changed (unless the -h option is specified), but no recursion takes place.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of chown for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how chown behaves.</p> <p>LC_CTYPE</p> <p>Determines how chown handles characters. When LC_CTYPE is set to a valid value, chown can display and handle text and filenames containing valid characters for that locale. chown can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. chown can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p>

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

FILES

/etc/passwd

SEE ALSO

chgrp(1), chmod(1), chown(2), passwd(4), environ(5)

NAME	chown – change owner
SYNOPSIS	<code>/usr/ucb/chown [-fR] owner[.group] filename ...</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>chown changes the owner of the <i>filenames</i> to <i>owner</i>. The owner may be either a decimal user ID (UID) or a login name found in the password file. An optional <i>group</i> may also be specified. The group may be either a decimal group ID (GID) or a group name found in the GID file.</p> <p>Only the super-user can change owner, in order to simplify accounting procedures.</p>
OPTIONS	<p>-f Do not report errors.</p> <p>-R Recursively descend into directories setting the ownership of all files in each directory encountered. When symbolic links are encountered, their ownership is changed, but they are not traversed.</p>
FILES	<code>/etc/passwd</code> password file
SEE ALSO	<code>chgrp(1)</code> , <code>chown(2)</code> , <code>group(4)</code> , <code>passwd(4)</code>

NAME	ckdate, errdate, helpdate, valdate – prompts for and validates a date
SYNOPSIS	<pre>ckdate [-Q] [-W width] [-f format] [-d default] [-h help] [-e error] [-p prompt] [-k pid [-s signal]] /usr/sadm/bin/errdate [-W width] [-e error] [-f format] /usr/sadm/bin/helpdate [-Wwidth] [-h help] [-f format] /usr/sadm/bin/valdate [-f format] [input]</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>ckdate prompts a user and validates the response. It defines, among other things, a prompt message whose response should be a date, text for help and error messages, and a default value (which will be returned if the user responds with a RETURN). The user response must match the defined format for a date.</p> <p>All messages are limited in length to 70 characters and are formatted automatically. Any white space used in the definition (including newline) is stripped. The -W option cancels the automatic formatting. When a tilde is placed at the beginning or end of a message definition, the default text will be inserted at that point, allowing both custom text and the default text to be displayed.</p> <p>If the prompt, help or error message is not defined, the default message (as defined under NOTES) will be displayed.</p> <p>Three visual tool modules are linked to the ckdate command. They are errdate (which formats and displays an error message), helpdate (which formats and displays a help message), and valdate (which validates a response). These modules should be used in conjunction with FML objects. In this instance, the FML object defines the prompt. When <i>format</i> is defined in the errdate and helpdate modules, the messages will describe the expected format.</p>
OPTIONS	<pre>-Q Specifies that quit will not be allowed as a valid response. -W width Specifies that prompt, help and error messages will be formatted to a line length of width. -f format Specifies the format against which the input will be verified. Possible formats and their definitions are: %b = abbreviated month name (jan, feb, mar) %B = full month name %d = day of month (01 - 31) %D = date as %m/%d/%y (the default format) %e = day of month (1 - 31; single digits are preceded by a blank) %h = abbreviated month name, identical to %b% %m = month number (01 - 12) %y = year within century (for instance, 89) %Y = year as CCYY (for instance, 1989)</pre>

	-d <i>default</i>	Defines the default value as <i>default</i> . The default does not have to meet the format criteria.
	-h <i>help</i>	Defines the help messages as <i>help</i> .
	-e <i>error</i>	Defines the error message as <i>error</i> .
	-p <i>prompt</i>	Defines the prompt message as <i>prompt</i> .
	-k <i>pid</i>	Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to abort.
	-s <i>signal</i>	Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal signal when quit is chosen. If no signal is specified, SIGTERM is used.
	<i>input</i>	Input to be verified against format criteria.
EXIT CODES	0	Successful execution
	1	EOF on input or negative width on -W option, or usage error
	3	User termination (quit)
	4	Garbled format argument
NOTES	The default prompt for ckdate is: Enter the date [?,q]:	
	The default error message is: ERROR - Please enter a date. Format is <format>.	
	The default help message is: Please enter a date. Format is <format>.	
	When the quit option is chosen (and allowed), q is returned along with the return code 3 . The valdate module will not produce any output. It returns zero for success and non-zero for failure.	

NAME	ckgid, errgid, helpgid, valgid – prompts for and validates a group id																				
SYNOPSIS	<pre>ckgid [-Q] [-W width] [-m] [-d default] [-h help] [-e error] [-p prompt] [-k pid [-s signal]] /usr/sadm/bin/errgid [-W width] [-e error] /usr/sadm/bin/helpgid [-W width] [-m] [-h help] /usr/sadm/bin/valgid input</pre>																				
AVAILABILITY	SUNWcsu																				
DESCRIPTION	<p>ckgid prompts a user and validates the response. It defines, among other things, a prompt message whose response should be an existing group ID, text for help and error messages, and a default value (which will be returned if the user responds with a carriage return).</p> <p>All messages are limited in length to 70 characters and are formatted automatically. Any white space used in the definition (including newline) is stripped. The -W option cancels the automatic formatting. When a tilde is placed at the beginning or end of a message definition, the default text will be inserted at that point, allowing both custom text and the default text to be displayed.</p> <p>If the prompt, help or error message is not defined, the default message (as defined under NOTES) will be displayed.</p> <p>Three visual tool modules are linked to the ckgid command. They are errgid (which formats and displays an error message), helpgid (which formats and displays a help message), and valgid (which validates a response). These modules should be used in conjunction with FML objects. In this instance, the FML object defines the prompt.</p>																				
OPTIONS	<table border="0"> <tr> <td style="padding-right: 10px;">-Q</td> <td>Specifies that quit will not be allowed as a valid response.</td> </tr> <tr> <td>-W width</td> <td>Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i>.</td> </tr> <tr> <td>-m</td> <td>Displays a list of all groups when help is requested or when the user makes an error.</td> </tr> <tr> <td>-d default</td> <td>Defines the default value as <i>default</i>. The default is not validated and so does not have to meet any criteria.</td> </tr> <tr> <td>-h help</td> <td>Defines the help messages as <i>help</i>.</td> </tr> <tr> <td>-e error</td> <td>Defines the error message as <i>error</i>.</td> </tr> <tr> <td>-p prompt</td> <td>Defines the prompt message as <i>prompt</i>.</td> </tr> <tr> <td>-k pid</td> <td>Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to abort.</td> </tr> <tr> <td>-s signal</td> <td>Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal signal when quit is chosen. If no signal is specified, SIGTERM is used.</td> </tr> <tr> <td><i>input</i></td> <td>Input to be verified against /etc/group.</td> </tr> </table>	-Q	Specifies that quit will not be allowed as a valid response.	-W width	Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i> .	-m	Displays a list of all groups when help is requested or when the user makes an error.	-d default	Defines the default value as <i>default</i> . The default is not validated and so does not have to meet any criteria.	-h help	Defines the help messages as <i>help</i> .	-e error	Defines the error message as <i>error</i> .	-p prompt	Defines the prompt message as <i>prompt</i> .	-k pid	Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to abort.	-s signal	Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal signal when quit is chosen. If no signal is specified, SIGTERM is used.	<i>input</i>	Input to be verified against /etc/group .
-Q	Specifies that quit will not be allowed as a valid response.																				
-W width	Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i> .																				
-m	Displays a list of all groups when help is requested or when the user makes an error.																				
-d default	Defines the default value as <i>default</i> . The default is not validated and so does not have to meet any criteria.																				
-h help	Defines the help messages as <i>help</i> .																				
-e error	Defines the error message as <i>error</i> .																				
-p prompt	Defines the prompt message as <i>prompt</i> .																				
-k pid	Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to abort.																				
-s signal	Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal signal when quit is chosen. If no signal is specified, SIGTERM is used.																				
<i>input</i>	Input to be verified against /etc/group .																				

EXIT CODES	0	Successful execution
	1	EOF on input or negative width on -W option, or usage error
	3	User termination (quit)

NOTES

The default prompt for **ckgid** is:
Enter the name of an existing group [?,q]:

The default error message is:
ERROR: Please enter one of the following group names: [List]

If the **-m** option of **ckgid** is used, a list of valid groups is displayed here.

The default help message is:
ERROR: Please enter one of the following group names: [List]

If the **-m** option of **ckgid** is used, a list of valid groups is displayed here.

When the quit option is chosen (and allowed), **q** is returned along with the return code **3**. The **valgid** module will not produce any output. It returns zero for success and non-zero for failure.

NAME	ckint, errint, helpint, valint – display a prompt; verify and return an integer value
SYNOPSIS	<p>ckint [-Q] [-W <i>width</i>] [-b <i>base</i>] [-d <i>default</i>] [-h <i>help</i>] [-e <i>error</i>] [-p <i>prompt</i>] [-k <i>pid</i> [-s <i>signal</i>]]</p> <p>/usr/sadm/bin/errint [-W <i>width</i>] [-b <i>base</i>] [-e <i>error</i>] /usr/sadm/bin/helpint [-W <i>width</i>] [-b <i>base</i>] [-h <i>help</i>] /usr/sadm/bin/valint [-b <i>base</i>] <i>input</i></p>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>ckint prompts a user, then validates the response. It defines, among other things, a prompt message whose response should be an integer, text for help and error messages, and a default value (which will be returned if the user responds with a carriage return).</p> <p>All messages are limited in length to 70 characters and are formatted automatically. Any white space used in the definition (including newline) is stripped. The -W option cancels the automatic formatting. When a tilde is placed at the beginning or end of a message definition, the default text will be inserted at that point, allowing both custom text and the default text to be displayed.</p> <p>If the prompt, help or error message is not defined, the default message (as defined under NOTES) will be displayed.</p> <p>Three visual tool modules are linked to the ckint command. They are errint (which formats and displays an error message), helpint (which formats and displays a help message), and valint (which validates a response). These modules should be used in conjunction with FML objects. In this instance, the FML object defines the prompt. When <i>base</i> is defined in the errint and helpint modules, the messages will include the expected base of the input.</p>
OPTIONS	<p>-Q Specifies that quit will not be allowed as a valid response.</p> <p>-W <i>width</i> Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i>.</p> <p>-b <i>base</i> Defines the base for input. Must be 2 to 36, default is 10.</p> <p>-d <i>default</i> Defines the default value as <i>default</i>. The default is not validated and so does not have to meet any criteria.</p> <p>-h <i>help</i> Defines the help messages as <i>help</i>.</p> <p>-e <i>error</i> Defines the error message as <i>error</i>.</p> <p>-p <i>prompt</i> Defines the prompt message as <i>prompt</i>.</p> <p>-k <i>pid</i> Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to abort.</p>

	<i>-s signal</i>	Specifies that the process ID <i>pid</i> defined with the <i>-k</i> option is to be sent signal signal when quit is chosen. If no signal is specified, SIGTERM is used.
	<i>input</i>	Input to be verified against <i>base</i> criterion.
EXIT CODES	0	Successful execution
	1	EOF on input, or negative width on <i>-W</i> option, or usage error
	3	User termination (quit)
NOTES	<p>The default base 10 prompt for ckint is:</p> <p style="padding-left: 40px;">Enter an integer [?,q]:</p> <p>The default base 10 error message is:</p> <p style="padding-left: 40px;">ERROR - Please enter an integer.</p> <p>The default base 10 help message is:</p> <p style="padding-left: 40px;">Please enter an integer.</p> <p>The messages are changed from “integer” to “base <i>base</i> integer” if the base is set to a number other than 10.</p> <p>When the quit option is chosen (and allowed), q is returned along with the return code 3. The valint module will not produce any output. It returns 0 for success and non-zero for failure.</p>	

NAME	ckitem, erritem, helpitem – build a menu; prompt for and return a menu item
SYNOPSIS	<pre>ckitem [-Q] [-W width] [-uno] [-f filename] [-l label] [[-i invis] [, ...]] [-m max] [-d default] [-h help] [-e error] [-p prompt] [-k pid [-s signal]] [choice [...]] /usr/sadm/bin/erritem [-W width] [-e error] [choice [...]] /usr/sadm/bin/helpitem [-W width] [-h help] [choice [...]]</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>ckitem builds a menu and prompts the user to choose one item from a menu of items. It then verifies the response. Options for this command define, among other things, a prompt message whose response will be a menu item, text for help and error messages, and a default value (which will be returned if the user responds with a carriage return).</p> <p>By default, the menu is formatted so that each item is prepended by a number and is printed in columns across the terminal. Column length is determined by the longest choice. Items are alphabetized.</p> <p>All messages are limited in length to 70 characters and are formatted automatically. Any white space used in the definition (including newline) is stripped. The -W option cancels the automatic formatting. When a tilde is placed at the beginning or end of a message definition, the default text will be inserted at that point, allowing both custom text and the default text to be displayed.</p> <p>If the prompt, help or error message is not defined, the default message (as defined under NOTES) will be displayed.</p> <p>Two visual tool modules are linked to the ckitem command. They are erritem (which formats and displays an error message) and helpitem (which formats and displays a help message). These modules should be used in conjunction with FML objects. In this instance, the FML object defines the prompt. When <i>choice</i> is defined in these modules, the messages will describe the available menu choice (or choices).</p>
OPTIONS	<p>-Q Specify that quit will not be allowed as a valid response.</p> <p>-W width Specify that prompt, help and error messages will be formatted to a line length of <i>width</i>.</p> <p>-u Specify that menu items should be displayed as an unnumbered list.</p> <p>-n Specify that menu items should not be displayed in alphabetical order.</p> <p>-o Specify that only one menu token will be returned.</p> <p>-f filename Define a file, <i>filename</i>, which contains a list of menu items to be displayed. (The format of this file is: token<tab>description. Lines beginning with a pound sign (#) are designated as comments and ignored.)</p> <p>-l label Define a label, <i>label</i>, to print above the menu.</p>

- i *invis*** Define invisible menu choices (those which will not be printed in the menu). (For example, "all" used as an invisible choice would mean it is a legal option but does not appear in the menu. Any number of invisible choices may be defined.) Invisible choices should be made known to a user either in the prompt or in a help message.
- m *max*** Define the maximum number of menu choices that the user can choose. The default is 1.
- d *default*** Define the default value as *default*. The default is not validated and so does not have to meet any criteria.
- h *help*** Define the help messages as *help*.
- e *error*** Define the error message as *error*.
- p *prompt*** Define the prompt message as *prompt*.
- k *pid*** Specify that the process ID *pid* is to be sent a signal if the user chooses to abort.
- s *signal*** Specify that process ID *pid* defined with the **-k** option is to be sent signal **signal** when quit is chosen. If no signal is specified, SIGTERM is used.
- choice*** Define menu items. Items should be separated by white space or newline.

EXIT CODES

- 0** Successful execution
- 1** EOF on input, or negative width on **-W** option, or inability to open file on **-f** option, or usage error
- 3** User termination (quit)
- 4** No choices from which to choose

NOTES

The user may input the number of the menu item if choices are numbered or as much of the string required for a unique identification of the item. Long menus are paged with 10 items per page.

When menu entries are defined both in a file (by using the **-f** option) and also on the command line, they are usually combined alphabetically. However, if the **-n** option is used to suppress alphabetical ordering, then the entries defined in the file are shown first, followed by the options defined on the command line.

The default prompt for **ckitem** is:

Enter selection [?,??,q]:

One question mark will give a help message and then redisplay the prompt. Two question marks will give a help message and then redisplay the menu label, the menu and the prompt.

The default error message if you typed a number is:

ERROR: Bad numeric choice specification

The default error message if you typed a string is:

ERROR: Entry does not match available menu selection. Enter the number of the menu item you wish to select, the token which is associated with the menu item, or a partial string which uniquely identifies the token for the menu item. Enter ?? to reprint the menu.

The default help message is:

Enter the number of the menu item you wish to select, the token which is associated with the menu item, or a partial string which uniquely identifies the token for the menu item. Enter ?? to reprint the menu.

When the quit option is chosen (and allowed), **q** is returned along with the return code **3**.

NAME	ckkeywd – prompts for and validates a keyword
SYNOPSIS	ckkeywd [-Q] [-W <i>width</i>] [-d <i>default</i>] [-h <i>help</i>] [-e <i>error</i>] [-p <i>prompt</i>] [-k <i>pid</i> [-s <i>signal</i>]] <i>keyword</i> [...]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>ckkeywd prompts a user and validates the response. It defines, among other things, a prompt message whose response should be one of a list of keywords, text for help and error messages, and a default value (which will be returned if the user responds with a carriage return). The answer returned from this command must match one of the defined list of keywords.</p> <p>All messages are limited in length to 70 characters and are formatted automatically. Any white space used in the definition (including newline) is stripped. The -W option cancels the automatic formatting. When a tilde is placed at the beginning or end of a message definition, the default text will be inserted at that point, allowing both custom text and the default text to be displayed.</p> <p>If the prompt, help or error message is not defined, the default message (as defined under NOTES) will be displayed.</p>
OPTIONS	<p>-Q Specifies that quit will not be allowed as a valid response.</p> <p>-W <i>width</i> Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i>.</p> <p>-d <i>default</i> Defines the default value as <i>default</i>. The default is not validated and so does not have to meet any criteria.</p> <p>-h <i>help</i> Defines the help messages as <i>help</i>.</p> <p>-e <i>error</i> Defines the error message as <i>error</i>.</p> <p>-p <i>prompt</i> Defines the prompt message as <i>prompt</i>.</p> <p>-k <i>pid</i> Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to abort.</p> <p>-s <i>signal</i> Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal signal when quit is chosen. If no signal is specified, SIGTERM is used.</p> <p><i>keyword</i> Defines the keyword, or list of keywords, against which the answer will be verified.</p>
EXIT CODES	<p>0 Successful execution</p> <p>1 EOF on input, or negative width on -W option, or no keywords from which to choose, or usage error</p>

3 User termination (quit)**NOTES**

The default prompt for **ckkeywd** is:

Enter appropriate value [*keyword*,[...],?,**q**]:

The default error message is:

ERROR: Please enter one of the following keywords:

keyword,[...],**q**

The default help message is:

keyword,[...],**q**

When the quit option is chosen (and allowed), **q** is returned along with the return code **3**.

NAME	ckpath, errpath, helppath, valpath – display a prompt; verify and return a pathname
SYNOPSIS	<pre>ckpath [-Q] [-W width] [-a l] [-b c f y] [-n [o z]] [-rtwx] [-d default] [-h help] [-e error] [-p prompt] [-k pid [-s signal]] /usr/sadm/bin/errpath [-W width] [-a l] [-b c f y] [-n [o z]] [-rtwx] [-e error] /usr/sadm/bin/helppath [-W width] [-a l] [-b c f y] [-n [o z]] [-rtwx] [-h help] /usr/sadm/bin/valpath [-a l] [-b c f y] [-n [o z]] [-rtwx] input</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>ckpath prompts a user and validates the response. It defines, among other things, a prompt message whose response should be a pathname, text for help and error messages, and a default value (which is returned if the user responds with a RETURN).</p> <p>The pathname must obey the criteria specified by the first group of options. If no criteria is defined, the pathname must be for a normal file that does not yet exist. If neither -a (absolute) or -l (relative) is given, then either is assumed to be valid.</p> <p>All messages are limited in length to 79 characters and are formatted automatically. Tabs and newlines are removed after a single white space character in a message definition, but spaces are not removed. When a tilde is placed at the beginning or end of a message definition, the default text is inserted at that point, allowing both custom text and the default text to be displayed.</p> <p>If the prompt, help or error message is not defined, the default message (as defined under EXAMPLES) is displayed.</p> <p>Three visual tool modules are linked to the ckpath command. They are errpath (which formats and displays an error message on the standard output), helppath (which formats and displays a help message on the standard output), and valpath (which validates a response). These modules should be used in conjunction with Framed Access Command Environment (FACE) objects. In this instance, the FACE object defines the prompt.</p>
OPTIONS	<p>-Q Specify that quit is not allowed as a valid response.</p> <p>-W width Specify that prompt, help and error messages be formatted to a line length of <i>width</i>.</p> <p>-a Pathname must be an absolute path.</p> <p>-l Pathname must be a relative path.</p> <p>-b Pathname must be a block special file.</p> <p>-c Pathname must be a character special file.</p> <p>-f Pathname must be a regular file.</p> <p>-y Pathname must be a directory.</p>

-n	Pathname must not exist (must be new).
-o	Pathname must exist (must be old).
-z	Pathname must have a file having a size greater than 0 bytes.
-r	Pathname must be readable.
-t	Pathname must be creatable (touchable). Pathname will be created if it does not already exist.
-w	Pathname must be writable.
-x	Pathname must be executable.
-d default	Defines the default value as <i>default</i> . The default is not validated and so does not have to meet any criteria.
-h help	Defines the help message as <i>help</i> .
-e error	Defines the error message as <i>error</i> .
-p prompt	Defines the prompt message as <i>prompt</i> .
-k pid	Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to quit.
-s signal	Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal <i>signal</i> when quit is chosen. If no signal is specified, SIGTERM is used.
<i>input</i>	Input to be verified against validation options.

EXIT CODES

0	Successful execution
1	EOF on input, or negative width on -W option, or usage error
2	Mutually exclusive options
3	User termination (quit)
4	Mutually exclusive options

EXAMPLES

The text of the default messages for **ckpath** depends upon the criteria options that have been used. An example default prompt for **ckpath** (using the **-a** option) is:

```
example% ckpath -a
Enter an absolute pathname [?,q]
```

An example default error message (using the **-a** option) is:

```
example% /usr/sadm/bin/errpath -a
ERROR: A pathname is a filename, optionally preceded by parent directories.
The pathname you enter: - must begin with a slash (/)
```

An example default help message (using the **-a** option) is:

```
example% /usr/sadm/bin/helppath -a
A pathname is a filename, optionally preceded by parent directories.
The pathname you enter: - must begin with a slash (/)
```

When the quit option is chosen (and allowed), **q** is returned along with the return code **3**. Quit input gets a trailing newline.

The **valpath** module will produce a usage message on stderr. It returns **0** for success and non-zero for failure.

example% /usr/sadm/bin/valpath

usage: valpath [-[a | l][b | c | f | y][n | [o | z]]rtwx] input

.
.
.

SEE ALSO **face(1), signal(5)**

NAME	ckrange, errrange, helprange, valrange – prompts for and validates an integer
SYNOPSIS	<pre>ckrange [-Q] [-W width] [-l lower] [-u upper] [-b base] [-d default] [-h help] [-e error] [-p prompt] [-k pid [-s signal]] /usr/sadm/bin/errrange [-W width] [-e error] [-l lower] [-u upper] [-b base] /usr/sadm/bin/helprange [-W width] [-h help] [-l lower] [-u upper] [-b base] /usr/sadm/bin/valrange [-l lower] [-u upper] [-b base] input</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>ckrange prompts a user for an integer between a specified range and determines whether this response is valid. It defines, among other things, a prompt message whose response should be an integer in the range specified, text for help and error messages, and a default value (which is returned if the user responds with a RETURN).</p> <p>This command also defines a range for valid input. If either the lower or upper limit is left undefined, then the range is bounded on only one end.</p> <p>All messages are limited in length to 79 characters and are formatted automatically. Tabs and newlines are removed after a single whitespace character in a message definition, but spaces are not removed. When a tilde is placed at the beginning or end of a message definition, the default text will be inserted at that point, allowing both custom text and the default text to be displayed.</p> <p>If the prompt, help or error message is not defined, the default message (as defined under EXAMPLES) is displayed.</p> <p>Three visual tool modules are linked to the ckrange command. They are errrange (which formats and displays an error message on the standard output), helprange (which formats and displays a help message on the standard output), and valrange (which validates a response). These modules should be used in conjunction with Framed Access Command Environment (FACE) objects. In this instance, the FACE object defines the prompt.</p> <p>Note: Negative "input" arguments confuse getopt in valrange. By inserting a "--" before the argument, getopt processing will stop. See getopt(1) and intro(1) about getopt parameter handling. getopt is used to parse positional parameters and to check for legal options.</p>
OPTIONS	<p>-Q Specifies that quit will not be allowed as a valid response.</p> <p>-W width Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i>.</p> <p>-l lower Defines the lower limit of the range as <i>lower</i>. Default is the machine's largest negative long.</p>

	-u <i>upper</i>	Defines the upper limit of the range as <i>upper</i> . Default is the machine's largest positive long.
	-b <i>base</i>	Defines the base for input. Must be 2 to 36, default is 10. Base conversion uses strtol(3C) . Output is always base 10.
	-d <i>default</i>	Defines the default value as <i>default</i> . <i>default</i> is converted using strtol(3C) in the desired base. Any characters invalid in the specified base will terminate the strtol conversion without error.
	-h <i>help</i>	Defines the help message as <i>help</i> .
	-e <i>error</i>	Defines the error message as <i>error</i> .
	-p <i>prompt</i>	Defines the prompt message as <i>prompt</i> .
	-k <i>pid</i>	Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to quit.
	-s <i>signal</i>	Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal <i>signal</i> when quit is chosen. If no signal is specified, SIGTERM is used.
	<i>input</i>	Input to be verified against upper and lower limits and base.
EXIT CODES	0	Successful execution
	1	EOF on input, or negative width on -W option, or usage error
	2	Usage error
	3	User termination (quit)
EXAMPLES	The default base 10 prompt for ckrange is: <pre>example% ckrange Enter an integer between lower_bound and upper_bound [lower_bound-upper_bound,?,q]:</pre>	
	The default base 10 error message is: <pre>example% /usr/sadm/bin/errrange ERROR: Please enter an integer between lower_bound and upper_bound.</pre>	
	The default base 10 help message is: <pre>example% /usr/sadm/bin/helprange Please enter an integer between lower_bound and upper_bound.</pre>	
	The messages are changed from "integer" to "base <i>base</i> integer" if the base is set to a number other than 10, for example, example% /usr/sadm/bin/helprange -b 36 .	
	When the quit option is chosen (and allowed), q is returned along with the return code 3 . Quit input gets a trailing newline.	

The **valrange** module will produce a usage message on stderr. It returns **0** for success and non-zero for failure.

example% /usr/sadm/bin/valrange

usage: valrange [-l lower] [-u upper] [-b base] input

SEE ALSO

intro(1), face(1), getopt(1), strtol(3C), signal(5)

NAME	ckstr, errstr, helpstr, valstr – display a prompt; verify and return a string answer
SYNOPSIS	<pre>ckstr [-Q] [-W width] [[-r regexp] [...]] [-l length] [-d default] [-h help] [-e error] [-p prompt] [-k pid [-s signal]] /usr/sadm/bin/errstr [-W width] [-e error] [-l length] [[-r regexp] [...]] /usr/sadm/bin/helpstr [-W width] [-h help] [-l length] [[-r regexp] [...]] /usr/sadm/bin/valstr [-l length] [[-r regexp] [...]] input</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>ckstr prompts a user and validates the response. It defines, among other things, a prompt message whose response should be a string, text for help and error messages, and a default value (which are returned if the user responds with a RETURN).</p> <p>The answer returned from this command must match the defined regular expression and be no longer than the length specified. If no regular expression is given, valid input must be a string with a length less than or equal to the length defined with no internal, leading or trailing white space. If no length is defined, the length is not checked.</p> <p>All messages are limited in length to 79 characters and are formatted automatically. Tabs and newlines are removed after a single white space character in a message definition, but spaces are not removed. When a tilde is placed at the beginning or end of a message definition, the default text will be inserted at that point, allowing both custom text and the default text to be displayed.</p> <p>If the prompt, help or error message is not defined, the default message (as defined under EXAMPLES) is displayed.</p> <p>Three visual tool modules are linked to the ckstr command. They are errstr (which formats and displays an error message on the standard output), helpstr (which formats and displays a help message on the standard output), and valstr (which validates a response). These modules should be used in conjunction with Framed Access Command Environment (FACE) objects. In this instance, the FACE object defines the prompt.</p>
OPTIONS	<p>-Q Specifies that quit will not be allowed as a valid response.</p> <p>-W width Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i>.</p> <p>-r regexp Specifies a regular expression, regexp, against which the input should be validated. May include white space. If multiple expressions are defined, the answer need match only one of them.</p> <p>-l length Specifies the maximum length of the input.</p> <p>-d default Defines the default value as <i>default</i>. The default is not validated and so does not have to meet any criteria.</p> <p>-h help Defines the help message as <i>help</i>.</p> <p>-e error Defines the error message as <i>error</i>.</p>

-p <i>prompt</i>	Defines the prompt message as <i>prompt</i> .
-k <i>pid</i>	Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to quit.
-s <i>signal</i>	Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal <i>signal</i> when quit is chosen. If no signal is specified, SIGTERM is used.
<i>input</i>	Input to be verified against format length and/or regular expression criteria.
EXIT CODES	
0	Successful execution
1	EOF on input, or negative width on -W option, or usage error
2	Invalid regular expression
3	User termination (quit)

EXAMPLES

The default prompt for **ckstr** is:

```
example% ckstr
Enter an appropriate value [?,q]:
```

The default error message is dependent upon the type of validation involved. The user will be told either that the length or the pattern matching failed. The default error message is:

```
example% /usr/sadm/bin/errstr
ERROR: Please enter a string which contains no embedded,
leading or trailing spaces or tabs.
```

The default help message is also dependent upon the type of validation involved. If a regular expression has been defined, the message is:

```
example% /usr/sadm/bin/helpstr -r regexp
Please enter a string which matches the following pattern:
regexp
```

Other messages define the length requirement and the definition of a string.

When the quit option is chosen (and allowed), **q** is returned along with the return code **3**. Quit input gets a trailing newline.

The **valstr** module will produce a usage message on stderr. It returns **0** for success and non-zero for failure.

```
example% /usr/sadm/bin/valstr
usage: valstr [-l length] [--r regexp] [...] input
```

SEE ALSO

face(1), **signal(5)**

NAME	cktime, errtime, helptime, valtime – display a prompt; verify and return a time of day
SYNOPSIS	<pre>cktime [-Q] [-W width] [-f format] [-d default] [-h help] [-e error] [-p prompt] [-k pid [-s signal]] /usr/sadm/bin/errtime [-W width] [-e error] [-f format] /usr/sadm/bin/helptime [-W width] [-h help] [-f format] /usr/sadm/bin/valtime [-f format] input</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>cktime prompts a user and validates the response. It defines, among other things, a prompt message whose response should be a time, text for help and error messages, and a default value (which is returned if the user responds with a RETURN). The user response must match the defined format for the time of day.</p> <p>All messages are limited in length to 70 characters and are formatted automatically. Any white space used in the definition (including NEWLINE) is stripped. The -W option cancels the automatic formatting. When a tilde is placed at the beginning or end of a message definition, the default text is inserted at that point, allowing both custom text and the default text to be displayed.</p> <p>If the prompt, help or error message is not defined, the default message (as defined under NOTES) is displayed.</p> <p>Three visual tool modules are linked to the cktime command. They are errtime (which formats and displays an error message), helptime (which formats and displays a help message), and valtime (which validates a response). These modules should be used in conjunction with FML objects. In this instance, the FML object defines the prompt. When <i>format</i> is defined in the errtime and helptime modules, the messages will describe the expected format.</p>
OPTIONS	<p>-Q Specifies that quit will not be allowed as a valid response.</p> <p>-W width Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i>.</p> <p>-f format Specifies the format against which the input will be verified. Possible formats and their definitions are:</p> <ul style="list-style-type: none"> %H = hour (00 - 23) %I = hour (00 - 12) %M = minute (00 - 59) %p = ante meridian or post meridian %r = time as %I:%M:%S %p %R = time as %H:%M (the default format) %S = seconds (00 - 59) %T = time as %H:%M:%S

	-d <i>default</i>	Defines the default value as <i>default</i> . The default is not validated and so does not have to meet any criteria.
	-h <i>help</i>	Defines the help messages as <i>help</i> .
	-e <i>error</i>	Defines the error message as <i>error</i> .
	-p <i>prompt</i>	Defines the prompt message as <i>prompt</i> .
	-k <i>pid</i>	Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to abort.
	-s <i>signal</i>	Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal signal when quit is chosen. If no signal is specified, SIGTERM is used.
	<i>input</i>	Input to be verified against format criteria.
EXIT CODES	0	Successful execution
	1	EOF on input, or negative width on -W option, or usage error
	3	User termination (quit)
	4	Garbled format argument
NOTES	The default prompt for cktime is: <p style="text-align: center;">Enter a time of day [?,q]:</p> The default error message is: <p style="text-align: center;">ERROR: Please enter the time of day. Format is <format>.</p> The default help message is: <p style="text-align: center;">Please enter the time of day. Format is <format>.</p> When the quit option is chosen (and allowed), q is returned along with the return code 3 . The valtime module will not produce any output. It returns 0 for success and non-zero for failure.	

NAME	ckuid, erruid, helpuid, valuid – prompts for and validates a user ID																				
SYNOPSIS	<pre>ckuid [-Q] [-W width] [-m] [-d default] [-h help] [-e error] [-p prompt] [-k pid [-s signal]] /usr/sadm/bin/erruid [-W width] [-e error] /usr/sadm/bin/helpuid [-W width] [-m] [-h help] /usr/sadm/bin/valuid input</pre>																				
AVAILABILITY	SUNWcsu																				
DESCRIPTION	<p>ckuid prompts a user and validates the response. It defines, among other things, a prompt message whose response should be an existing user ID, text for help and error messages, and a default value (which are returned if the user responds with a RETURN).</p> <p>All messages are limited in length to 70 characters and are formatted automatically. Any white space used in the definition (including NEWLINE) is stripped. The -W option cancels the automatic formatting. When a tilde is placed at the beginning or end of a message definition, the default text is inserted at that point, allowing both custom text and the default text to be displayed.</p> <p>If the prompt, help or error message is not defined, the default message (as defined under NOTES) is displayed.</p> <p>Three visual tool modules are linked to the ckuid command. They are erruid (which formats and displays an error message), helpuid (which formats and displays a help message), and valuid (which validates a response). These modules should be used in conjunction with FML objects. In this instance, the FML object defines the prompt.</p>																				
OPTIONS	<table border="0"> <tr> <td style="vertical-align: top;">-Q</td> <td>Specifies that quit will not be allowed as a valid response.</td> </tr> <tr> <td style="vertical-align: top;">-W width</td> <td>Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i>.</td> </tr> <tr> <td style="vertical-align: top;">-m</td> <td>Displays a list of all logins when help is requested or when the user makes an error.</td> </tr> <tr> <td style="vertical-align: top;">-d default</td> <td>Defines the default value as <i>default</i>. The default is not validated and so does not have to meet any criteria.</td> </tr> <tr> <td style="vertical-align: top;">-h help</td> <td>Defines the help messages as <i>help</i>.</td> </tr> <tr> <td style="vertical-align: top;">-e error</td> <td>Defines the error message as <i>error</i>.</td> </tr> <tr> <td style="vertical-align: top;">-p prompt</td> <td>Defines the prompt message as <i>prompt</i>.</td> </tr> <tr> <td style="vertical-align: top;">-k pid</td> <td>Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to abort.</td> </tr> <tr> <td style="vertical-align: top;">-s signal</td> <td>Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal signal when quit is chosen. If no signal is specified, SIGTERM is used.</td> </tr> <tr> <td style="vertical-align: top;"><i>input</i></td> <td>Input to be verified against /etc/passwd.</td> </tr> </table>	-Q	Specifies that quit will not be allowed as a valid response.	-W width	Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i> .	-m	Displays a list of all logins when help is requested or when the user makes an error.	-d default	Defines the default value as <i>default</i> . The default is not validated and so does not have to meet any criteria.	-h help	Defines the help messages as <i>help</i> .	-e error	Defines the error message as <i>error</i> .	-p prompt	Defines the prompt message as <i>prompt</i> .	-k pid	Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to abort.	-s signal	Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal signal when quit is chosen. If no signal is specified, SIGTERM is used.	<i>input</i>	Input to be verified against /etc/passwd .
-Q	Specifies that quit will not be allowed as a valid response.																				
-W width	Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i> .																				
-m	Displays a list of all logins when help is requested or when the user makes an error.																				
-d default	Defines the default value as <i>default</i> . The default is not validated and so does not have to meet any criteria.																				
-h help	Defines the help messages as <i>help</i> .																				
-e error	Defines the error message as <i>error</i> .																				
-p prompt	Defines the prompt message as <i>prompt</i> .																				
-k pid	Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to abort.																				
-s signal	Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal signal when quit is chosen. If no signal is specified, SIGTERM is used.																				
<i>input</i>	Input to be verified against /etc/passwd .																				

EXIT CODES	0	Successful execution
	1	EOF on input, or negative width on -W option, or usage error
	2	Usage error
	3	User termination (quit)

NOTES

The default prompt for **ckuid** is:
Enter the login name of an existing user [?,q]:

The default error message is:
ERROR - Please enter the login name of an existing user.

If the **-m** option is used, the default error message is:
ERROR: Please enter one of the following login names: <List>

The default help message is:
Please enter the login name of an existing user.

If the **-m** option is used, the default help message is:
Please enter one of the following login names: <List>

When the quit option is chosen (and allowed), **q** is returned along with the return code **3**. The **valuid** module will not produce any output. It returns **0** for success and non-zero for failure.

NAME	ckyornd, erryornd, helpyornd, valyornd – prompts for and validates yes/no																		
SYNOPSIS	<pre>ckyornd [-Q] [-W width] [-d default] [-h help] [-e error] [-p prompt] [-k pid [-s signal]] /usr/sadm/bin/erryornd [-W width] [-e error] /usr/sadm/bin/helpyornd [-W width] [-h help] /usr/sadm/bin/valyornd input</pre>																		
AVAILABILITY	SUNWcsu																		
DESCRIPTION	<p>ckyornd prompts a user and validates the response. It defines, among other things, a prompt message for a yes or no answer, text for help and error messages, and a default value (which is returned if the user responds with a RETURN).</p> <p>All messages are limited in length to 70 characters and are formatted automatically. Any white space used in the definition (including newline) is stripped. The -W option cancels the automatic formatting. When a tilde is placed at the beginning or end of a message definition, the default text is inserted at that point, allowing both custom text and the default text to be displayed.</p> <p>If the prompt, help or error message is not defined, the default message (as defined under NOTES) is displayed.</p> <p>Three visual tool modules are linked to the ckyornd command. They are erryornd (which formats and displays an error message), helpyornd (which formats and displays a help message), and valyornd (which validates a response). These modules should be used in conjunction with FACE objects. In this instance, the FACE object defines the prompt.</p>																		
OPTIONS	<table border="0"> <tr> <td style="vertical-align: top;">-Q</td> <td>Specifies that quit will not be allowed as a valid response.</td> </tr> <tr> <td style="vertical-align: top;">-W width</td> <td>Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i>.</td> </tr> <tr> <td style="vertical-align: top;">-d default</td> <td>Defines the default value as <i>default</i>. The default is not validated and so does not have to meet any criteria.</td> </tr> <tr> <td style="vertical-align: top;">-h help</td> <td>Defines the help messages as <i>help</i>.</td> </tr> <tr> <td style="vertical-align: top;">-e error</td> <td>Defines the error message as <i>error</i>.</td> </tr> <tr> <td style="vertical-align: top;">-p prompt</td> <td>Defines the prompt message as <i>prompt</i>.</td> </tr> <tr> <td style="vertical-align: top;">-k pid</td> <td>Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to abort.</td> </tr> <tr> <td style="vertical-align: top;">-s signal</td> <td>Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal signal when quit is chosen. If no signal is specified, SIGTERM is used.</td> </tr> <tr> <td style="vertical-align: top;"><i>input</i></td> <td>Input to be verified as y, yes, or n, no (in any combination of upper- and lower-case letters).</td> </tr> </table>	-Q	Specifies that quit will not be allowed as a valid response.	-W width	Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i> .	-d default	Defines the default value as <i>default</i> . The default is not validated and so does not have to meet any criteria.	-h help	Defines the help messages as <i>help</i> .	-e error	Defines the error message as <i>error</i> .	-p prompt	Defines the prompt message as <i>prompt</i> .	-k pid	Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to abort.	-s signal	Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal signal when quit is chosen. If no signal is specified, SIGTERM is used.	<i>input</i>	Input to be verified as y , yes , or n , no (in any combination of upper- and lower-case letters).
-Q	Specifies that quit will not be allowed as a valid response.																		
-W width	Specifies that prompt, help and error messages will be formatted to a line length of <i>width</i> .																		
-d default	Defines the default value as <i>default</i> . The default is not validated and so does not have to meet any criteria.																		
-h help	Defines the help messages as <i>help</i> .																		
-e error	Defines the error message as <i>error</i> .																		
-p prompt	Defines the prompt message as <i>prompt</i> .																		
-k pid	Specifies that process ID <i>pid</i> is to be sent a signal if the user chooses to abort.																		
-s signal	Specifies that the process ID <i>pid</i> defined with the -k option is to be sent signal signal when quit is chosen. If no signal is specified, SIGTERM is used.																		
<i>input</i>	Input to be verified as y , yes , or n , no (in any combination of upper- and lower-case letters).																		

EXIT CODES	0	Successful execution
	1	EOF on input, or negative width on -W option, or usage error
	2	Usage error
	3	User termination (quit)

NOTES

The default prompt for **ckyornd** is:
Yes or No [y,n,?,q]:

The default error message is:
ERROR - Please enter yes or no.

The default help message is:
To respond in the affirmative, enter y, yes, Y, or YES.
To respond in the negative, enter n, no, N, or NO.

When the quit option is chosen (and allowed), **q** is returned along with the return code **3**. The **valyorn** module will not produce any output. It returns **0** for success and non-zero for failure.

NAME	clear – clear the terminal screen
SYNOPSIS	clear
AVAILABILITY	SUNWcsu
DESCRIPTION	clear clears your screen if this is possible. It looks in the environment for the terminal type and then in the terminfo database to figure out how to clear the screen.

NAME	cmp – compare two files
SYNOPSIS	cmp [-l] [-s] <i>filename1 filename2</i> [<i>skip1</i>] [<i>skip2</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	The two files are compared. (If <i>filename1</i> is - , the standard input is used.) Under default options, cmp makes no comment if the files are the same; if they differ, it announces the byte and line numbers at which the first difference occurred. If one file is an initial subsequence of the other, that fact is noted. <i>skip1</i> and <i>skip2</i> are initial byte offsets into <i>filename1</i> and <i>filename2</i> respectively, and may be either octal or decimal; a leading 0 denotes octal.
OPTIONS	<p>-l Print the byte number (decimal) and the differing bytes (octal) for each difference.</p> <p>-s Print nothing for differing files; return exit codes only.</p>
EXAMPLE	<p>The following example:</p> <p style="padding-left: 40px;">example% cmp file1 file2 0 1024</p> <p>does a byte for byte comparison of <i>file1</i> and <i>file2</i>. It skips the first 1024 bytes in <i>file2</i> before starting the comparison.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of cmp for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how cmp behaves.</p> <p>LC_CTYPE</p> <p>Determines how cmp handles characters. When LC_CTYPE is set to a valid value, cmp can display and handle text and filenames containing valid characters for that locale. cmp can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. cmp can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p> <p>LC_MESSAGES</p> <p>Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).</p>
SEE ALSO	comm(1) , diff(1) , environ(5)

DIAGNOSTICS

Exit code **0** is returned for identical files, **1** for different files, and **2** for an inaccessible or missing argument.

NAME	cof2elf – COFF to ELF object file translation
SYNOPSIS	cof2elf [-iqV] [-Qy n] [-s <i>directory</i>] <i>files</i>
AVAILABILITY	x86 SUNWbtool
DESCRIPTION	<p>cof2elf converts one or more COFF object <i>files</i> to ELF. This translation occurs in place, which means that the original file contents are modified. If an input file is an archive, each member will be translated as necessary, and the archive will be rebuilt with its members in the original order. cof2elf does not change input files that are not COFF.</p> <p>Options have the following meanings:</p> <ul style="list-style-type: none"> -i Normally, the files are modified only when full translation occurs. Unrecognized data, such as unknown relocation types, are treated as errors and prevent translation. Using the -i flag causes the program to ignore these partial translation conditions and modify the file anyway. -q Normally, cof2elf prints a message for each file it examines, reporting whether the file was translated, ignored, etc. Using the -q (quiet) flag suppresses these messages. -Qarg If <i>arg</i> is <i>y</i>, identification information about cof2elf will be added to the output files. This can be useful for software administration. Giving <i>n</i> for <i>arg</i> explicitly asks for no such information, which is the default behavior. -s directory As mentioned above, cof2elf modifies the input files. This option saves a copy of the original files in the specified <i>directory</i>, which must already exist. cof2elf does not save files it does not modify. -V This flag tells cof2elf to print a version message on standard error.
SEE ALSO	ld(1) , elf(3E) , a.out(4) , ar(4)
NOTES	<p>Some debugging information is discarded. Although this does not affect the behavior of a running program, it may affect the information available for symbolic debugging.</p> <p>cof2elf translates only COFF relocatable files. It does not translate executable or static shared library files for two main reasons. First, the operating system supports executable files and static shared libraries, making translation unnecessary. Second, those files have specific address and alignment constraints determined by the file format. Matching the constraints with a different object file format is problematic.</p> <p>When possible, programmers should recompile their source code to build new object files. cof2elf is provided for those times when source code is unavailable.</p>

NAME	<code>col</code> – reverse line-feeds filter
SYNOPSIS	<code>col [-b] [-f] [-x] [-p]</code>
AVAILABILITY	SUNWesu
DESCRIPTION	<p><code>col</code> reads from the standard input and writes onto the standard output. It performs the line overlays implied by reverse line feeds (ASCII code ESC-7), and by forward and reverse half-line-feeds (ESC-9 and ESC-8). <code>col</code> is particularly useful for filtering multicolumn output made with the <code>.rt</code> command of <code>nroff</code> and output resulting from use of the <code>tbl(1)</code> preprocessor.</p> <p>The ASCII control characters SO (\017) and SI (\016) are assumed by <code>col</code> to start and end text in an alternate character set. The character set to which each input character belongs is remembered, and on output SI and SO characters are generated as appropriate to ensure that each character is printed in the correct character set.</p> <p>On input, the only control characters accepted are space, backspace, tab, return, new-line, SI, SO, VT (\013), and ESC followed by 7, 8, or 9. The VT character is an alternate form of full reverse line-feed, included for compatibility with some earlier programs of this type. All other non-printing characters are ignored.</p>
OPTIONS	<p>-b <code>col</code> assumes that the output device in use is not capable of backspacing. In this case, if two or more characters are to appear in the same place, only the last one read will be output.</p> <p>-f Although <code>col</code> accepts half-line motions in its input, it normally does not emit them on output. Instead, text that would appear between lines is moved to the next lower full-line boundary. This treatment can be suppressed by the <code>-f</code> (fine) option; in this case, the output from <code>col</code> may contain forward half-line-feeds (ESC-9), but will still never contain either kind of reverse line motion.</p> <p>-x Unless the <code>-x</code> option is given, <code>col</code> will convert white space to tabs on output wherever possible to shorten printing time.</p> <p>-p Normally, <code>col</code> will ignore any escape sequences unknown to it that are found in its input; the <code>-p</code> option may be used to cause <code>col</code> to output these sequences as regular characters, subject to overprinting from reverse line motions. The use of this option is highly discouraged unless the user is fully aware of the textual position of the escape sequences.</p>
ENVIRONMENT	<p>If any of the <code>LC_*</code> variables (<code>LC_CTYPE</code>, <code>LC_MESSAGES</code>, <code>LC_TIME</code>, <code>LC_COLLATE</code>, <code>LC_NUMERIC</code>, and <code>LC_MONETARY</code>) (see <code>environ(5)</code>) are not set in the environment, the operational behavior of <code>col</code> for each corresponding locale category is determined by the value of the <code>LANG</code> environment variable. If <code>LC_ALL</code> is set, its contents are used to override both the <code>LANG</code> and the other <code>LC_*</code> variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how <code>col</code> behaves.</p>

LC_CTYPE

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

SEE ALSO

nroff(1), **tbl(1)**, **ascii(5)**, **environ(5)**

NOTES

The input format accepted by **col** matches the output produced by **nroff** with either the **-T37** or **-Tlp** options. Use **-T37** (and the **-f** option of **col**) if the ultimate disposition of the output of **col** will be a device that can interpret half-line motions, and **-Tlp** otherwise.

col cannot back up more than 128 lines or handle more than 800 characters per line.

Local vertical motions that would result in backing up over the first line of the document are ignored. As a result, the first line must not have any superscripts.

NAME	comm – select or reject lines common to two sorted files
SYNOPSIS	comm [-1 -2 -3 -12 -13 -23] <i>filename1 filename2</i>
AVAILABILITY	SUNWesu
DESCRIPTION	comm reads <i>filename1</i> and <i>filename2</i> , which should be ordered in ASCII collating sequence (see sort (1)), and produces a three-column output: lines only in <i>filename1</i> ; lines only in <i>filename2</i> ; and lines in both files. The file name – means the standard input.
OPTIONS	The following options can be used to suppress the indicated columns from the display. <ul style="list-style-type: none"> -1 Suppress column 1; omit lines only in <i>filename1</i>. -2 Suppress column 2; omit lines only in <i>filename2</i>. -3 Suppress column 3; omit lines common to both files. -12 Suppress columns 1 and 2; only show lines common to both files. -13 Suppress columns 1 and 3; only show lines in <i>filename2</i>. -23 Suppress columns 2 and 3; only show lines in <i>filename1</i>.
ENVIRONMENT	If any of the LC_* variables (LC_CTYPE , LC_MESSAGES , LC_TIME , LC_COLLATE , LC_NUMERIC , and LC_MONETARY) (see environ (5)) are not set in the environment, the operational behavior of comm for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how comm behaves. <p>LC_CTYPE Determines how comm handles characters. When LC_CTYPE is set to a valid value, comm can display and handle text and filenames containing valid characters for that locale. comm can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. comm can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p> <p>LC_MESSAGES Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).</p>
SEE ALSO	cmp (1), diff (1), sort (1), uniq (1), environ (5)
NOTES	You can specify -123 , but doing so suppresses all output.

NAME	compress, uncompress, zcat – compress, uncompress files or display expanded files
SYNOPSIS	compress [-cfv] [-b bits] [<i>filename...</i>] uncompress [-cfv] [<i>filename...</i>] zcat [<i>filename...</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	compress reduces the size of the named files using adaptive Lempel-Ziv coding. Whenever possible, each file is replaced by one with a .Z , extension. The ownership modes (if super-user), access time and modification time will stay the same. If no files are specified, the standard input is compressed to the standard output. The amount of compression obtained depends on the size of the input, the number of <i>bits</i> per code, and the distribution of common substrings. Typically, text such as source code or English is reduced by 50–60%. Compression is generally much better than that achieved by Huffman coding (as used in pack(1)), and takes less time to compute. The <i>bits</i> parameter specified during compression is encoded within the compressed file, along with a magic number to ensure that neither decompression of random data nor recompression of compressed data is subsequently allowed. Compressed files can be restored to their original form using uncompress . zcat produces uncompressed output on the standard output, but leaves the compressed .Z file intact.
OPTIONS	-c Write to the standard output; no files are changed. The nondestructive behavior of zcat is identical to that of ' uncompress -c '. -f When compressing, force compression, even if the file does not actually shrink, or the corresponding .Z file already exists. Except when running in the background (under /usr/bin/sh), if -f is not given, prompt to verify whether an existing .Z file should be overwritten. When uncompressing, forces uncompression, even if filename exists. -v Verbose. Display the percentage reduction for each file compressed. -b bits Set the upper limit (in bits) for common substring codes. <i>bits</i> must be between 9 and 16 (16 is the default). Lowering the number of bits will result in larger, less compressed files.
FILES	/usr/bin/sh
SEE ALSO	pack(1)
DIAGNOSTICS	Exit status is normally 0 . If the last file was not compressed because it became larger, the status is 2 . If an error occurs, exit status is 1 . Usage: compress [-fvc] [-b maxbits] [<i>filename ...</i>] Invalid options were specified on the command line.

Missing maxbits

Maxbits must follow **-b**.

filename: not in compressed format

The file specified to **uncompress** has not been compressed.

filename: compressed with xxbits, can only handle yybits

filename was compressed by a program that could deal with more *bits* than the compress code on this machine. Recompress the file with smaller *bits*.

filename: already has .Z suffix -- no change

The file is assumed to be already compressed. Rename the file and try again.

filename: already exists; do you wish to overwrite (y or n)?

Respond **y** if you want the output file to be replaced; **n** if not.

uncompress: corrupt input

A SIGSEGV violation was detected, which usually means that the input file is corrupted.

Compression: xx.xx%

Percentage of the input saved by compression. (Relevant only for **-v**.)

-- not a regular file: unchanged

When the input file is not a regular file, (such as a directory), it is left unaltered.

-- has xx other links: unchanged

The input file has links; it is left unchanged. See **ln(1)** for more information.

-- file unchanged

No savings are achieved by compression. The input remains uncompressed.

NOTES

Although compressed files are compatible between machines with large memory, **-b 12** should be used for file transfer to architectures with a small process data space (64KB or less).

compress should be more flexible about the existence of the **.Z** suffix.

NAME	coproc, cocreate, cosend, cocheck, coreceive, codestroy – communicate with a process
SYNOPSIS	<pre> cocreate [-r <i>rpath</i>] [-w <i>wpath</i>] [-i <i>id</i>] [-R <i>refname</i>] [-s <i>send_string</i>] [-e <i>expect_string</i>] <i>command</i> cosend [-n] <i>proc_id</i> <i>string</i> cocheck <i>proc_id</i> coreceive <i>proc_id</i> codestroy [-R <i>refname</i>] <i>proc_id</i> [<i>string</i>] </pre>
DESCRIPTION	<p>These co-processing functions provide a flexible means of interaction between FMLI and an independent process; especially, they enable FMLI to be responsive to asynchronous activity.</p> <p>The cocreate function starts <i>command</i> as a co-process and initializes communications by setting up pipes between FMLI and the standard input and standard output of <i>command</i>. The argument <i>command</i> must be an executable and its arguments (if any). This means that <i>command</i> expects strings on its input (supplied by cosend) and sends information on its output that can be handled in various ways by FMLI.</p> <p>The cosend function sends <i>string</i> to the co-process identified by <i>proc_id</i> via the pipe set up by cocreate (optionally <i>wpath</i>), where <i>proc_id</i> can be either the <i>command</i> or <i>id</i> specified in cocreate. By default, cosend blocks, waiting for a response from the co-process. Also by default, FMLI does not send a <i>send_string</i> and does not expect an <i>expect_string</i> (except a newline). That is, it reads only one line of output from the co-process. If -e <i>expect_string</i> was not defined when the pipe was created, then the output of the co-process is any single string followed by a newline; any other lines of output remain on the pipe. If the -e option was specified when the pipe was created, cosend reads lines from the pipe until it reads a line starting with <i>expect_string</i>. All lines except the line starting with <i>expect_string</i> become the output of cosend.</p> <p>The cocheck function determines if input is available from the process identified by <i>proc_id</i>, where <i>proc_id</i> can be either the <i>command</i> or <i>id</i> specified in cocreate. It returns a Boolean value, which makes cocheck useful in if statements and in other backquoted expressions in Boolean descriptors. cocheck receives no input from the co-process; it simply indicates if input is available from the co-process. You must use coreceive to actually accept the input. The cocheck function can be called from a reread descriptor to force a frame to update when new data is available. This is useful when the default value of a field in a form includes coreceive.</p> <p>The coreceive function is used to read input from the co-process identified by <i>proc_id</i>, where <i>proc_id</i> can be either the <i>command</i> or <i>id</i> specified in cocreate. It should only be used when it has been determined, using cocheck, that input is actually available. If the -e option was used when the co-process was created, coreceive will continue to return lines of input until <i>expect_string</i> is read. At this point, coreceive will terminate. The output of coreceive is all the lines that were read excluding the line starting with <i>expect_string</i>. If the -e option was not used in the cocreate, each invocation of coreceive</p>

will return exactly one line from the co-process. If no input is available when **coreceive** is invoked, it will simply terminate without producing output.

The **codestroy** function terminates the read/write pipes to *proc-id*, where *proc-id* can be either the *command* or *id* specified in **cocreate**. It generates a **SIGPIPE** signal to the (child) co-process. This kills the co-process, unless the co-process ignores the **SIGPIPE** signal. If the co-process ignores the **SIGPIPE**, it will not die, even after the FMLI process terminates (the parent process id of the co-process will be 1).

The optional argument *string* is sent to the co-process before the co-process dies. If *string* is not supplied, a NULL string is passed, followed by the normal *send_string* (newline by default). That is, **codestroy** will call **cosend proc-id string**: this implies that **codestroy** will write any output generated by the co-process to *stdout*. For example, if an interactive co-process is written to expect a "quit" string when the communication is over, the **close** descriptor could be defined;

```
close=`codestroy ID 'quit' | message`
```

and any output generated by the co-process when the string **quit** is sent to it via **codes-troy** (using **cosend**) would be redirected to the message line.

The **codestroy** function should usually be given the **-R** option, since you may have more than one process with the same name, and you do not want to kill the wrong one. **codes-troy** keeps track of the number of *refnames* you have assigned to a process with **cocreate**, and when the last instance is killed, it kills the process (*id*) for you. **codestroy** is typically called as part of a **close** descriptor because **close** is evaluated when a frame is closed. This is important because the co-process will continue to run if **codestroy** is not issued.

When writing programs to use as co-processes, the following tips may be useful. If the co-process program is written in C language, be sure to flush output after writing to the pipe. (Currently, **awk**(1) and **sed**(1) cannot be used in a co-process program because they do not flush after lines of output.) Shell scripts are well-mannered, but slow. C language is recommended. If possible, use the default *send_string*, *rpath* and *wpath*. In most cases, *expect_string* will have to be specified. This, of course, depends on the co-process.

In the case where asynchronous communication from a co-process is desired, a co-process program should use **vsig** to force strings into the pipe and then signal FMLI that output from the co-process is available. This causes the **reread** descriptor of all frames to be evaluated immediately.

OPTIONS

cocreate options are:

- r *rpath*** If **-r** is specified, *rpath* is the pathname from which FMLI reads information. This option is usually used to set up communication with processes that naturally write to a certain path. If **-r** is not specified, **cocreate** will choose a unique path in **/var/tmp**.
- w *wpath*** If **-w** is specified, *wpath* is the pathname to which **cosend** writes information. This option is usually used so that one process can talk to many different FMLI processes through the same pipe. If **-w** is not specified, **cocreate** will choose a unique path in **/var/tmp**.
- i *id*** If **-i** is specified, *id* is an alternative name for the co-process initialized by this **cocreate**. If **-i** is not specified, *id* defaults to *command*. The argument *id* can later be used with the other co-processing functions rather than *command*. This option is typically used, since it facilitates the creation of two or more co-processes generated from the same *command*. (For example, **cocreate -i ID1 program args** and **cocreate -i ID2 program different_args**).
- R *refname*** If **-R** is specified, *refname* is a local name for the co-process. Since the **cocreate** function can be issued more than once, a *refname* is useful when the same co-process is referenced a second or subsequent time. With the **-R** option, if the co-process already exists a new one will not be created: the same pipes will be shared. Then, *refname* can be used as an argument to the **-R** option to **codestroy** when you want to end a particular connection to a co-process and leave other connections undisturbed. (The co-process is only killed after **codestroy -R** has been called as many times as **cocreate -R** was called.)
- s *send_string*** The **-s** option specifies *send_string* as a string that will be appended to all output sent to the co-process using **cosend**. This option allows a co-process to know when input from FMLI has completed. The default *send_string* is a newline if **-s** is not specified.
- e *expect_string*** The **-e** option specifies *expect_string* as a string that identifies the end of all output returned by the co-process. (Note: *expect_string* need only be the initial part of a line, and there must be a newline at the end of the co-process output.) This option allows FMLI to know when output from the co-process has completed. The default *expect_string* is a newline if **-e** is not specified.

cosend options are:

- n** If the **-n** option is specified, **cosend** will not wait for a response from the co-process. It simply returns, providing no output. If the **-n** option is not used, a co-process that does not answer will cause FMLI to permanently hang, waiting for input from the co-process.

EXAMPLES

```
.  
. .  
init=`cocreate -i BIGPROCESS initialize`  
close=`codestroy BIGPROCESS`  
. .  
reread=`cocheck BIGPROCESS`  
  
name=`cosend -n BIGPROCESS field1`  
. .  
name="Receive field"  
inactive=TRUE  
value=`coreceive BIGPROCESS`
```

SEE ALSO [awk\(1\)](#), [cat\(1\)](#), [sed\(1\)](#), [vsig\(1F\)](#)

NOTES

If **cosend** is used without the **-n** option, a co-process that does not answer will cause FMLI to permanently hang.

The use of non-alphabetic characters in input and output strings to a co-process should be avoided because they may not get transferred correctly.

NAME	cp – copy files
SYNOPSIS	cp [-i] [-p] [-r] [<i>filename . . .</i>] <i>target</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The cp command copies <i>filename</i> to <i>target</i>. <i>filename</i> and <i>target</i> cannot have the same name. Be careful when using sh(1) metacharacters. If <i>target</i> is not a directory, only one file may be specified before it; if it is a directory, more than one file may be specified. If <i>target</i> does not exist, cp creates a file named <i>target</i>. If <i>target</i> exists and is not a directory, its contents are overwritten. If <i>target</i> is a directory, the file(s) are copied to that directory. <i>target</i> and <i>filename</i> do not have to share the same parent directory.</p> <p>If <i>filename</i> is a file and <i>target</i> is a link to another file with links, the other links remain and <i>target</i> becomes a new file.</p> <p>If <i>target</i> does not exist, cp creates a new file named <i>target</i> which has the same mode as <i>filename</i> except that the sticky bit is not set unless the user is a privileged user; the owner and group of <i>target</i> are those of the user.</p> <p>If <i>target</i> is a file, its contents are overwritten, but the mode, owner, and group associated with it are not changed. The last modification time of <i>target</i> and the last access time of <i>filename</i> are set to the time the copy was made.</p> <p>If <i>target</i> is a directory, then for each file named, a new file with the same mode is created in the target directory; the owner and the group are those of the user making the copy.</p>
OPTIONS	<p>-i Interactive. cp will prompt for confirmation whenever the copy would overwrite an existing <i>target</i>. A y answer means that the copy should proceed. Any other answer prevents cp from overwriting <i>target</i>.</p> <p>-p Preserve. cp will duplicate not only the contents of <i>filename</i>, but also preserves the modification time and permission modes.</p> <p>-r If <i>filename</i> is a directory, cp will copy the directory and all its files, including any subdirectories and their files; <i>target</i> must be a directory.</p>
EXAMPLES	<p>To copy a file:</p> <pre>example% cp goodies goodies.old example% ls goodies* goodies goodies.old</pre>

To copy a directory, first to a new, and then to an existing destination directory:

```
example% ls ~/bkup  
/usr/example/fred/bkup not found  
example% cp -r ~/src ~/bkup  
example% ls -R ~/bkup  
x.c y.c z.sh  
example% cp -r ~/src ~/bkup  
example% ls -R ~/bkup  
src x.c y.c z.sh
```

```
src:  
x.c y.c z.sh
```

To copy a list of files to a destination directory:

```
example% cp ~/src/* /tmp
```

SEE ALSO

chmod(1), cpio(1), rm(1)

NOTES

The permission modes of the source file are preserved in the copy.

A `--` permits the user to mark the end of any command line options explicitly, thus allowing `cp` to recognize filename arguments that begin with a `-`. If a `--` and a `-` both appear on the same command line, the second will be interpreted as a filename.

NAME	cpio – copy file archives in and out
SYNOPSIS	<pre> cpio -i [bBcdfkmrsStuvV6] [-C bufsize] [-E filename] [-H header] [-I filename [-M message]] [-R id] [<i>pattern ...</i>] cpio -o [aABcLvV] [-C bufsize] [-H header] [-O filename [-M message]] cpio -p [adlLmuvV] [-R id] <i>directory</i> </pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>cpio copies files in to and out from a cpio archive. The cpio archive may span multiple volumes. The -i, -o, and -p options select the action to be performed. The following list describes each of the actions (which are mutually exclusive).</p> <p>cpio -i (copy in) extracts files from the standard input, which is assumed to be the product of a previous cpio -o. Only files with names that match <i>patterns</i> are selected. <i>patterns</i> are regular expressions given in the filename-generating notation of sh(1). In <i>patterns</i>, meta-characters ?, *, and [...] match the slash (/) character, and backslash (\) is an escape character. A ! meta-character means <i>not</i>. (For example, the !abc* pattern would exclude all files that begin with abc.) Multiple <i>patterns</i> may be specified and if no <i>patterns</i> are specified, the default for <i>patterns</i> is * (that is, select all files). Each pattern must be enclosed in double quotes; otherwise, the name of a file in the current directory might be used. Extracted files are conditionally created and copied into the current directory tree based on the options described below. The permissions of the files will be those of the previous cpio -o. Owner and group will be the same as the current user unless the current user is super-user. If this is true, owner and group will be the same as those resulting from the previous cpio -o. Note: If cpio -i tries to create a file that already exists and the existing file is the same age or younger (newer), cpio will output a warning message and not replace the file. (The -u option can be used to overwrite, unconditionally, the existing file.)</p> <p>cpio -o (copy out) reads the standard input to obtain a list of path names and copies those files onto the standard output together with path name and status information. Output is padded to a 512-byte boundary by default or to the user specified block size (with the -B or -C options) or to some device-dependent block size where necessary (as with the CTC tape).</p> <p>cpio -p (pass) reads the standard input to obtain a list of path names of files that are conditionally created and copied into the destination <i>directory</i> tree based on the options described below.</p> <p>Note: cpio assumes four-byte words.</p> <p>If, when writing to a character device (-o) or reading from a character device (-i), cpio reaches the end of a medium (such as the end of a diskette), and the -O and -I options are not used, cpio prints the following message:</p> <p style="text-align: center;">To continue, type device/file name when ready.</p>

To continue, you must replace the medium and type the character special device name (*/dev/rdiskette* for example) and press RETURN. You may want to continue by directing **cpio** to use a different device. For example, if you have two floppy drives you may want to switch between them so **cpio** can proceed while you are changing the floppies. (Simply pressing RETURN causes the **cpio** process to exit.)

OPTIONS

- i** **cpio -i** (copy in) extracts files from the standard input.
- o** **cpio -o** (copy out) reads the standard input to obtain a list of path names and copies those files onto the standard output.
- p** **cpio -p** (pass) reads the standard input to obtain a list of path names of files.
- a** Reset access times of input files after they have been copied. Access times are not reset for linked files when **cpio -pla** is specified (mutually exclusive with **-m**).
- A** Append files to an archive. The **-A** option requires the **-O** option. Valid only with archives that are files, or that are on floppy diskettes or hard disk partitions.
- b** Reverse the order of the bytes within each word. (Use only with the **-i** option.)
- B** Block input/output 5120 bytes to the record. The default buffer size is 512 bytes when this and the **-C** options are not used. (**-B** does not apply to the *pass* option; **-B** is meaningful only with data directed to or from a character special device, for example, */dev/rmt/0m*.)
- c** Read or write header information in ASCII character form for portability. Always use this option (or the **-H** option) when the origin and the destination machines are different types (mutually exclusive with **-H** and **-6**). The **-c** option implies expanded device numbers. When transferring files between Solaris 1.x and Solaris 2.x use regular **cpio** format. Don't use the **-c** or **-H** options.
- C bufsize** Block input/output *bufsize* bytes to the record, where *bufsize* is replaced by a positive integer. The default buffer size is 512 bytes when this and **-B** options are not used. (**-C** does not apply to the *pass* option; **-C** is meaningful only with data directed to or from a character special device, for example, */dev/rmt/0m*.)
- d** Create directories as needed.
- E filename** Specify an input file (*file*) that contains a list of filenames to be extracted from the archive (one filename per line).
- f** Copy in all files except those in *patterns*. (See the paragraph on **cpio -i** for a description of *patterns*.)
- H header** Read or write header information in *header* format. Always use this option or the **-c** option when the origin and the destination machines are different types (mutually exclusive with **-c** and **-6**). When

transferring files between Solaris 1.x and Solaris 2.x use regular **cpio** format. Don't use the **-c** or **-H** options.

Valid values for *header* are:

bar	bar head and format. Used only with the -i option (read only)
crc CRC	ASCII header with expanded device numbers and an additional per-file checksum
odc	ASCII header with small device numbers
tar TAR	tar header and format
ustar USTAR	IEEE/P1003 Data Interchange Standard header and format

- I filename** Read the contents of *filename* as an input archive. If *filename* is a character special device, and the current medium has been completely read, replace the medium and press RETURN to continue to the next medium. This option is used only with the **-i** option.
- k** Attempt to skip corrupted file headers and I/O errors that may be encountered. If you want to copy files from a medium that is corrupted or out of sequence, this option lets you read only those files with good headers. (For **cpio** archives that contain other **cpio** archives, if an error is encountered **cpio** may terminate prematurely. **cpio** will find the next good header, which may be one for a smaller archive, and terminate when the smaller archive's trailer is encountered.) Used only with the **-i** option.
- l** Whenever possible, link files rather than copying them. (Usable only with the **-p** option.)
- L** Follow symbolic links. The default is not to follow symbolic links.
- m** Retain previous file modification time. This option is ineffective on directories that are being copied (mutually exclusive with **-a**).
- M message** Define a *message* to use when switching media. When you use the **-O** or **-I** options and specify a character special device, you can use this option to define the message that is printed when you reach the end of the medium. One **%d** can be placed in *message* to print the sequence number of the next medium needed to continue.
- O filename** Direct the output of **cpio** to *filename*. If *filename* is a character special device and the current medium is full, replace the medium and type a carriage return to continue to the next medium. Use only with the **-o** option.
- r** Interactively rename files. If the user types a carriage return alone, the file is skipped. If the user types a "." the original pathname will be retained. (Not available with **cpio -p**.)
- R id** Reassign ownership and group information for each file to *user ID* (*ID*

- must be a valid login ID from `/etc/passwd`). This option is valid only for the super-user.
- `-s` Swap bytes within each half word.
 - `-S` Swap halfwords within each word.
 - `-t` Print a table of contents of the input. No files are created (mutually exclusive with `-V`).
 - `-u` Copy unconditionally (normally, an older file will not replace a newer file with the same name).
 - `-v` Verbose. Print a list of file names. When used with the `-t` option, the table of contents looks like the output of an `ls -l` command (see `ls(1)`).
 - `-V` Special verbose. Print a dot for each file read or written. Useful to assure the user that **cpio** is working without printing out all file names.
 - `-6` Process a UNIX System Sixth Edition archive format file. Use only with the `-i` option (mutually exclusive with `-c` and `-H`).

EXAMPLES

The following examples show three uses of **cpio**.

When standard input is directed through a pipe to **cpio -o**, it groups the files so they can be directed (`>`) to a single file (`./newfile`). The `-c` option insures that the file will be portable to other machines (as would the `-H` option). Instead of `ls(1)`, you could use `find(1)`, `echo(1)`, `cat(1)`, and so on, to pipe a list of names to **cpio**. You could direct the output to a device instead of a file.

```
example% ls | cpio -oc > ./newfile
```

cpio -i uses the output file of **cpio -o** (directed through a pipe with `cat` in the example below), extracts those files that match the patterns (`memo/a1`, `memo/b*`), creates directories below the current directory as needed (`-d` option), and places the files in the appropriate directories. The `-c` option is used if the input file was created with a portable header. If no patterns were given, all files from `newfile` would be placed in the directory.

```
example% cat newfile | cpio -icd "memo/a1" "memo/b*"
```

cpio -p takes the file names piped to it and copies or links (`-l` option) those files to another directory (`newdir` in the example below). The `-d` option says to create directories as needed. The `-m` option says retain the modification time. (It is important to use the `-depth` option of `find(1)` to generate path names for **cpio**. This eliminates problems **cpio** could have trying to create files under read-only directories.) The destination directory, `newdir`, must exist.

```
example% find . -depth -print | cpio -pdlmv newdir
```

Note: When you use **cpio** in conjunction with `find`, if you use the `-L` option with **cpio** then you must use the `-follow` option with `find` and vice versa. Otherwise there will be undesirable results.

Note that for multi-reel archives, dismount the old volume, mount the new one, and continue to the next tape by typing the name of the next device (probably the same as the first reel). To stop, type a RETURN and **cpio** will end.

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

LC_TIME

Determines how **cpio** handles date and time formats. In the "C" locale, date and time handling follows the U.S. rules.

SEE ALSO

ar(1), **cat(1)**, **echo(1)**, **find(1)**, **ls(1)**, **sh(1)**, **tar(1)**, **archives(4)**, **environ(5)**

NOTES

Path names are restricted to 256 characters for the binary (the default) and **-H odc** header formats. Otherwise, path names are restricted to 1024 characters.

Only the super-user can copy special files.

Blocks are reported in 512-byte quantities.

If a file has **000** permissions, contains more than 0 characters of data, and the user is not root, the file will not be saved or restored.

The inode number stored in the header, (**/usr/include/archives.h**) is an unsigned short which is 2 bytes. This limits the range of inode numbers from 0 to 65535. Files which are hard linked must fall in this inode range. This could be a problem when moving **cpio** archives between different vendors' machines.

NAME	cpp – the C language preprocessor
SYNOPSIS	<code>/usr/ccs/lib/cpp [-BCHMpPRT] [-undef] [-Dname] [-Dname=def] [-Idirectory] [-Uname] [-Ydirectory] [input-file [output-file]]</code>
DESCRIPTION	<p>cpp is the C language preprocessor. It is invoked as the first pass of any C compilation started with the <code>cc(1B)</code> command; however, cpp can also be used as a first-pass preprocessor for other Sun compilers.</p> <p>Although cpp can be used as a macro processor, this is not normally recommended, as its output is geared toward that which would be acceptable as input to a compiler's second pass. Thus, the preferred way to invoke cpp is through the <code>cc(1B)</code> command, or some other compilation command. For general-purpose macro-processing, see <code>m4(1)</code>, and the chapter on <code>m4</code> in <i>Programming Utilities Guide</i>.</p> <p>cpp optionally accepts two filenames as arguments. <i>input-file</i> and <i>output-file</i> are, respectively, the input and output files for the preprocessor. They default to the standard input and the standard output.</p>
OPTIONS	<p>-B Support the C++ comment indicator <code>/// on the line after the <code>//</code> is treated as a comment.</code></p> <p>-C Pass all comments (except those that appear on cpp directive lines) through the preprocessor. By default, cpp strips out C-style comments.</p> <p>-H Print the pathnames of included files, one per line on the standard error.</p> <p>-M Generate a list of makefile dependencies and write them to the standard output. This list indicates that the object file which would be generated from the input file depends on the input file as well as the include files referenced.</p> <p>-p Use only the first eight characters to distinguish preprocessor symbols, and issue a warning if extra tokens appear at the end of a line containing a directive.</p> <p>-P Preprocess the input without producing the line control information used by the next pass of the C compiler.</p> <p>-R Allow recursive macros.</p> <p>-T Use only the first eight characters for distinguishing different preprocessor names. This option is included for backward compatibility with systems which always use only the first eight characters.</p> <p>-undef Remove initial definitions for all predefined symbols.</p> <p>-Dname Define <i>name</i> as 1 (one). This is the same as if a -Dname=1 option appeared on the cpp command line, or as if a <code>#define name 1</code> line appeared in the source file that cpp is processing.</p>

- Dname=def** Define *name* as if by a **#define** directive. This is the same as if a **#define name def** line appeared in the source file that **cpp** is processing. The **-D** option has lower precedence than the **-U** option. That is, if the same name is used in both a **-U** option and a **-D** option, the name will be undefined regardless of the order of the options.
- Idirectory** Insert *directory* into the search path for **#include** files with names not beginning with '/'. *directory* is inserted ahead of the standard list of "include" directories. Thus, **#include** files with names enclosed in double-quotes (") are searched for first in the directory of the file with the **#include** line, then in directories named with **-I** options, and lastly, in directories from the standard list. For **#include** files with names enclosed in angle-brackets (<>), the directory of the file with the **#include** line is not searched. See **Details** below for exact details of this search order.
- Uname** Remove any initial definition of *name*, where *name* is a symbol that is predefined by a particular preprocessor. Here is a partial list of symbols that may be predefined, depending upon the architecture of the system:
- | | |
|--------------------------|--|
| Operating System: | ibm, gcos, os, tss and unix |
| Hardware: | interdata, pdp11, u370, u3b, u3b2, u3b5, u3b15, u3b20d, vax, ns32000, iAPX286, i386, sparc, and sun |
| UNIX system variant: | RES, and RT |
| The lint command: | lint |
- The symbols **sun**, **sparc** and **unix** are defined for all Sun systems.
- Ydirectory** Use directory *directory* in place of the standard list of directories when searching for **#include** files.

USAGE Directives

All **cpp** directives start with a hash symbol (#) as the first character on a line. White space (SPACE or TAB characters) can appear after the initial # for proper indentation.

#define name token-string

Replace subsequent instances of *name* with *token-string*.

#define name(argument [, argument] ...) token-string

There can be no space between *name* and the '(' . Replace subsequent instances of *name*, followed by a parenthesized list of arguments, with *token-string*, where each occurrence of an *argument* in the *token-string* is replaced by the corresponding token in the comma-separated list. When a macro with arguments is expanded, the arguments are placed into the expanded *token-string* unchanged. After the entire *token-string* has been expanded, **cpp** re-starts its scan for names to expand at the beginning of the newly created *token-string*.

#undef name

Remove any definition for the symbol *name*. No additional tokens are permitted

on the directive line after *name*.

#include "*filename*"

#include <*filename*>

Read in the contents of *filename* at this location. This data is processed by **cpp** as if it were part of the current file. When the <*filename*> notation is used, *filename* is only searched for in the standard "include" directories. See the **-I** and **-Y** options above for more detail. No additional tokens are permitted on the directive line after the final '"' or '>'.

#line *integer-constant* "*filename*"

Generate line control information for the next pass of the C compiler. *integer-constant* is interpreted as the line number of the next line and *filename* is interpreted as the file from where it comes. If "*filename*" is not given, the current filename is unchanged. No additional tokens are permitted on the directive line after the optional *filename*.

#if *constant-expression*

Subsequent lines up to the matching **#else**, **#elif**, or **#endif** directive, appear in the output only if *constant-expression* yields a nonzero value. All binary non-assignment C operators, including '&&', '|', and ',', are legal in *constant-expression*. The '?' operator, and the unary '-', '!', and '~' operators, are also legal in *constant-expression*.

The precedence of these operators is the same as that for C. In addition, the unary operator **defined**, can be used in *constant-expression* in these two forms: '**defined** (*name*)' or '**defined** *name*'. This allows the effect of **#ifdef** and **#ifndef** directives (described below) in the **#if** directive. Only these operators, integer constants, and names that are known by **cpp** should be used within *constant-expression*. In particular, the **sizeof** operator is not available.

#ifdef *name*

Subsequent lines up to the matching **#else**, **#elif**, or **#endif** appear in the output only if *name* has been defined, either with a **#define** directive or a **-D** option, and in the absence of an intervening **#undef** directive. Additional tokens after *name* on the directive line will be silently ignored.

#ifndef *name*

Subsequent lines up to the matching **#else**, **#elif**, or **#endif** appear in the output only if *name* has *not* been defined, or if its definition has been removed with an **#undef** directive. No additional tokens are permitted on the directive line after *name*.

#elif *constant-expression*

Any number of **#elif** directives may appear between an **#if**, **#ifdef**, or **#ifndef** directive and a matching **#else** or **#endif** directive. The lines following the **#elif** directive appear in the output only if all of the following conditions hold:

- The *constant-expression* in the preceding **#if** directive evaluated to zero, the *name* in the preceding **#ifdef** is not defined, or the *name* in the preceding **#ifndef** directive was defined.

- The *constant-expression* in all intervening **#elif** directives evaluated to zero.
- The current *constant-expression* evaluates to non-zero.

If the *constant-expression* evaluates to non-zero, subsequent **#elif** and **#else** directives are ignored up to the matching **#endif**. Any *constant-expression* allowed in an **#if** directive is allowed in an **#elif** directive.

#else This inverts the sense of the conditional directive otherwise in effect. If the preceding conditional would indicate that lines are to be included, then lines between the **#else** and the matching **#endif** are ignored. If the preceding conditional indicates that lines would be ignored, subsequent lines are included in the output. Conditional directives and corresponding **#else** directives can be nested.

#endif End a section of lines begun by one of the conditional directives **#if**, **#ifdef**, or **#ifndef**. Each such directive must have a matching **#endif**.

Macros

Formal parameters for macros are recognized in **#define** directive bodies, even when they occur inside character constants and quoted strings. For instance, the output from:

```
#define abc(a) | `| a |
abc(xyz)
```

is the seven characters “ | `xyz | ” (SPACE, vertical-bar, backquote, x, y, z, vertical-bar). Macro names are not recognized within character constants or quoted strings during the regular scan. Thus:

```
#define abc xyz
printf("abc");
```

does not expand **abc** in the second line, since it is inside a quoted string that is not part of a **#define** macro definition.

Macros are not expanded while processing a **#define** or **#undef**. Thus:

```
#define abc zingo
#define xyz abc
#undef abc
xyz
```

produces **abc**. The token appearing immediately after an **#ifdef** or **#ifndef** is not expanded.

Macros are not expanded during the scan which determines the actual parameters to another macro call. Thus:

```
#define reverse(first,second)second first
#define greeting hello
reverse(greeting,
#define greeting goodbye
)
```

	produces “ #define hello goodbye hello ”.
Output	Output consists of a copy of the input file, with modifications, plus lines of the form: <pre>#lineno " filename" "level"</pre> indicating the original source line number and filename of the following output line and whether this is the first such line after an include file has been entered (<i>level=1</i>), the first such line after an include file has been exited (<i>level=2</i>), or any other such line (<i>level</i> is empty).
Details	
<i>Directory Search Order</i>	#include files are searched for in the following order: <ol style="list-style-type: none"> 1. The directory of the file that contains the #include request (that is, #include is relative to the file being scanned when the request is made). 2. The directories specified by -I options, in left-to-right order. 3. The standard directory(s) (/usr/include on UNIX systems).
<i>Special Names</i>	Two special names are understood by cpp . The name <code>_ _LINE_ _</code> is defined as the current line number (a decimal integer) as known by cpp , and <code>_ _FILE_ _</code> is defined as the current filename (a C string) as known by cpp . They can be used anywhere (including in macros) just as any other defined name.
<i>Newline Characters</i>	A NEWLINE character terminates a character constant or quoted string. An escaped NEWLINE (that is, a backslash immediately followed by a NEWLINE) may be used in the body of a #define statement to continue the definition onto the next line. The escaped NEWLINE is not included in the macro value.
<i>Comments</i>	Comments are removed (unless the -C option is used on the command line). Comments are also ignored, except that a comment terminates a token.
SEE ALSO	cc(1B) , m4(1) <i>Programming Utilities Guide</i>
DIAGNOSTICS	The error messages produced by cpp are intended to be self-explanatory. The line number and filename where the error occurred are printed along with the diagnostic.
NOTES	When NEWLINE characters were found in argument lists for macros to be expanded, some previous versions of cpp put out the NEWLINE characters as they were found and expanded. The current version of cpp replaces them with SPACE characters. Because the standard directory for included files may be different in different environments, this form of #include directive: <pre>#include <file.h></pre> should be used, rather than one with an absolute path, like: <pre>#include "/usr/include/file.h"</pre>

cpp warns about the use of the absolute pathname.

While the compiler allows 8-bit strings and comments, 8-bits are not allowed anywhere else.

NAME	crontab – user crontab file
SYNOPSIS	crontab [<i>filename</i>] crontab -e [<i>username</i>] crontab -r [<i>username</i>] crontab -l [<i>username</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	crontab manages a user's access with cron by copying, creating, listing, and removing crontab files. If invoked without options, crontab copies the specified file, or the standard input if no file is specified, into a directory that holds all users' crontabs.
crontab Access Control	<p>Users: Access to crontab is allowed:</p> <ul style="list-style-type: none"> • if the user's name appears in /etc/cron.d/cron.allow. • if /etc/cron.d/cron.allow does not exist and the user's name is not in /etc/cron.d/cron.deny. <p>Users: Access to crontab is denied:</p> <ul style="list-style-type: none"> • if /etc/cron.d/cron.allow exists and the user's name is not in it. • if /etc/cron.d/cron.allow does not exist and user's name is in /etc/cron.d/cron.deny. • if neither file exists. <p>Note: The rules for allow and deny apply to root only if the allow/deny files exist. The allow/deny files consist of one user name per line.</p>
crontab Entry Format	<p>A crontab file consists of lines of six fields each. The fields are separated by spaces or tabs. The first five are integer patterns that specify the following:</p> <p style="margin-left: 2em;">minute (0–59), hour (0–23), day of the month (1–31), month of the year (1–12), day of the week (0–6 with 0=Sunday).</p> <p>Each of these patterns may be either an asterisk (meaning all legal values) or a list of elements separated by commas. An element is either a number or two numbers separated by a minus sign (meaning an inclusive range). Note that the specification of days may be made by two fields (day of the month and day of the week). If both are specified as a list of elements, both are adhered to. For example, 0 0 1,15 * 1 would run a command on the first and fifteenth of each month, as well as on every Monday. To specify days by only one field, the other field should be set to * (for example, 0 0 * * 1 would run a command only on Mondays).</p>

The sixth field of a line in a **crontab** file is a string that is executed by the shell at the specified times. A percent character in this field (unless escaped by `\`) is translated to a new-line character. Only the first line (up to a `%` or end of line) of the command field is executed by the shell. Other lines are made available to the command as standard input. Any line beginning with a `#` is a comment and will be ignored.

The shell is invoked from your **\$HOME** directory with an **arg0** of **sh**. Users who desire to have their **.profile** executed must explicitly do so in the **crontab** file. **crontab** supplies a default environment for every shell, defining **HOME**, **LOGNAME**, **SHELL(=/bin/sh)**, **TZ**, and **PATH**. The default **PATH** for **user** cron jobs is **/usr/bin**; while **root** cron jobs default to **/usr/sbin:/usr/bin**. The default **PATH** can be set in **/etc/default/cron**. (See **crontab(1M)**).

If you do not redirect the standard output and standard error of your commands, any generated output or errors will be mailed to you.

OPTIONS

-e edits a copy of the current user's **crontab** file, or creates an empty file to edit if **crontab** does not exist. When editing is complete, the file is installed as the user's **crontab** file. If a *username* is given, the specified user's **crontab** file is edited, rather than the current user's **crontab** file; this may only be done by a privileged user. The environment variable **EDITOR** determines which editor is invoked with the **-e** option. The default editor is **ed(1)**.

Note: All crontab jobs should be submitted using **crontab**; you should not add jobs by just editing **crontab** file because **crontab** will not be aware of changes made this way.

-l lists the **crontab** file for the invoking user. Only a privileged user can specify a *username* following the **-r** or **-l** options to remove or list the **crontab** file of the specified user.

-r removes a user's crontab from the crontab directory.

FILES

/etc/cron.d	main cron directory
/etc/cron.d/cron.allow	list of allowed users
/etc/default/cron	contains cron default settings
/etc/cron.d/cron.deny	list of denied users
/var/cron/log	accounting information
/var/spool/cron/crontabs	spool area for crontab .

SEE ALSO

atq(1), **atrm(1)**, **ed(1)**, **sh(1)**, **crontab(1M)**, **su(1M)**

NOTES

If you inadvertently enter the **crontab** command with no argument(s), do not attempt to get out with a CTRL-D. This removes all entries in your **crontab** file. Instead, exit with a CTRL-C.

If a privileged user modifies another user's **crontab** file, resulting behavior may be unpredictable. Instead, the privileged user should first **su(1M)** to the other user's login before making any changes to the **crontab** file.

NAME	crypt – encode or decode a file
SYNOPSIS	crypt [<i>password</i>]
AVAILABILITY	SUNWcry
DESCRIPTION	<p>crypt encrypts and decrypts the contents of a file. crypt reads from the standard input and writes on the standard output. The <i>password</i> is a key that selects a particular transformation. If no <i>password</i> is given, crypt demands a key from the terminal and turns off printing while the key is being typed in. crypt encrypts and decrypts with the same key:</p> <p style="padding-left: 40px;">example% crypt key <clear.file>encrypted.file example% crypt key <encrypted.file> pr</p> <p>will print the contents of <i>clear. file</i>.</p> <p>Files encrypted by crypt are compatible with those treated by the editors ed(1), ex(1) and vi(1) in encryption mode.</p> <p>The security of encrypted files depends on three factors: the fundamental method must be hard to solve; direct search of the key space must be infeasible; “sneak paths” by which keys or cleartext can become visible must be minimized.</p> <p>crypt implements a one-rotor machine designed along the lines of the German Enigma, but with a 256-element rotor. Methods of attack on such machines are widely known, thus crypt provides minimal security.</p> <p>The transformation of a key into the internal settings of the machine is deliberately designed to be expensive, that is, to take a substantial fraction of a second to compute. However, if keys are restricted to (say) three lower-case letters, then encrypted files can be read by expending only a substantial fraction of five minutes of machine time.</p> <p>Since the key is an argument to the crypt command, it is potentially visible to users executing ps(1) or a derivative command. To minimize this possibility, crypt takes care to destroy any record of the key immediately upon entry. No doubt the choice of keys and key security are the most vulnerable aspect of crypt.</p>
FILES	/dev/tty for typed key
SEE ALSO	des (1), ed (1), ex (1), ps (1), vi (1), makekey (1)
RESTRICTIONS	This program is not available on software shipped outside the U.S.

NAME	cs – shell command interpreter with a C-like syntax
SYNOPSIS	cs [-bcefinstvVxX] [<i>argument</i> . . .]
AVAILABILITY	SUNWcsu
DESCRIPTION	cs , the C shell, is a command interpreter with a syntax reminiscent of the C language. It provides a number of convenient features for interactive use that are not available with the Bourne shell, including filename completion, command aliasing, history substitution, job control, and a number of built-in commands. As with the Bourne shell, the C shell provides variable, command and filename substitution.
Initialization and Termination	<p>When first started, the C shell normally performs commands from the .cshrc file in your home directory, provided that it is readable and you either own it or your real group ID matches its group ID. If the shell is invoked with a name that starts with ‘-’, as when started by login(1), the shell runs as a login shell. In this case, after executing commands from the .cshrc file, the shell executes commands from the .login file in your home directory; the same permission checks as those for .cshrc are applied to this file. Typically, the .login file contains commands to specify the terminal type and environment.</p> <p>As a login shell terminates, it performs commands from the .logout file in your home directory; the same permission checks as those for .cshrc are applied to this file.</p>
Interactive Operation	After startup processing is complete, an interactive C shell begins reading commands from the terminal, prompting with <i>hostname%</i> (or <i>hostname#</i> for the privileged user). The shell then repeatedly performs the following actions: a line of command input is read and broken into <i>words</i> . This sequence of words is placed on the history list and then parsed, as described under USAGE, below. Finally, the shell executes each command in the current line.
Noninteractive Operation	When running noninteractively, the shell does not prompt for input from the terminal. A noninteractive C shell can execute a command supplied as an <i>argument</i> on its command line, or interpret commands from a file, also known as a script.
OPTIONS	<p>-b Force a “break” from option processing. Subsequent command line arguments are not interpreted as C shell options. This allows the passing of options to a script without confusion. The shell does not run set-user-ID or set-group-ID scripts unless this option is present.</p> <p>-c Read commands from the first filename <i>argument</i> (which must be present). Remaining arguments are placed in argv, the argument-list variable, and passed directly to cs.</p> <p>-e Exit if a command terminates abnormally or yields a nonzero exit status.</p> <p>-f Fast start. Read neither the .cshrc file, nor the .login file (if a login shell) upon startup.</p>

- i** Forced interactive. Prompt for command line input, even if the standard input does not appear to be a terminal (character-special device).
- n** Parse (interpret), but do not execute commands. This option can be used to check C shell scripts for syntax errors.
- s** Take commands from the standard input.
- t** Read and execute a single command line. A `'\'` (backslash) can be used to escape each newline for continuation of the command line onto subsequent input lines.
- v** Verbose. Set the **verbose** predefined variable; command input is echoed after history substitution (but before other substitutions) and before execution.
- V** Set **verbose** before reading **.cshrc**.
- x** Echo. Set the **echo** variable; echo commands after all substitutions and just before execution.
- X** Set **echo** before reading **.cshrc**.

Except with the options `-c`, `-i`, `-s`, or `-t`, the first nonoption *argument* is taken to be the name of a command or script. It is passed as argument zero, and subsequent arguments are added to the argument list for that command or script.

USAGE

Filename Completion

When enabled by setting the variable **filec**, an interactive C shell can complete a partially typed filename or user name. When an unambiguous partial filename is followed by an ESC character on the terminal input line, the shell fills in the remaining characters of a matching filename from the working directory.

If a partial filename is followed by the EOF character (usually typed as CTRL-d), the shell lists all filenames that match. It then prompts once again, supplying the incomplete command line typed in so far.

When the last (partial) word begins with a tilde (`~`), the shell attempts completion with a user name, rather than a file in the working directory.

The terminal bell signals errors or multiple matches; this can be inhibited by setting the variable **nobeep**. You can exclude files with certain suffixes by listing those suffixes in the variable **ignore**. If, however, the only possible completion includes a suffix in the list, it is not ignored. **ignore** does not affect the listing of filenames by the EOF character.

Lexical Structure

The shell splits input lines into words at space and tab characters, except as noted below. The characters `&`, `|`, `;`, `<`, `>`, `(`, and `)` form separate words; if paired, the pairs form single words. These shell metacharacters can be made part of other words, and their special meaning can be suppressed by preceding them with a `'\'` (backslash). A newline preceded by a `\` is equivalent to a space character.

In addition, a string enclosed in matched pairs of single-quotes ('), double-quotes ("), or backquotes (`), forms a partial word; metacharacters in such a string, including any space or tab characters, do not form separate words. Within pairs of backquote (`) or double-quote (") characters, a newline preceded by a ' \ ' (backslash) gives a true newline character. Additional functions of each type of quote are described, below, under **Variable Substitution**, **Command Substitution**, and **Filename Substitution**.

When the shell's input is not a terminal, the character # introduces a comment that continues to the end of the input line. Its special meaning is suppressed when preceded by a \ or enclosed in matching quotes.

Command Line Parsing

A *simple command* is composed of a sequence of words. The first word (that is not part of an I/O redirection) specifies the command to be executed. A simple command, or a set of simple commands separated by | or | & characters, forms a *pipeline*. With |, the standard output of the preceding command is redirected to the standard input of the command that follows. With | &, both the standard error and the standard output are redirected through the pipeline.

Pipelines can be separated by semicolons (;), in which case they are executed sequentially. Pipelines that are separated by && or | | form conditional sequences in which the execution of pipelines on the right depends upon the success or failure, respectively, of the pipeline on the left.

A pipeline or sequence can be enclosed within parentheses (') to form a simple command that can be a component in a pipeline or sequence.

A sequence of pipelines can be executed asynchronously or "in the background" by appending an '&'; rather than waiting for the sequence to finish before issuing a prompt, the shell displays the job number (see **Job Control**, below) and associated process IDs and prompts immediately.

History Substitution

History substitution allows you to use words from previous command lines in the command line you are typing. This simplifies spelling corrections and the repetition of complicated commands or arguments. Command lines are saved in the history list, the size of which is controlled by the **history** variable. The most recent command is retained in any case. A history substitution begins with a ! (although you can change this with the **histchars** variable) and may occur anywhere on the command line; history substitutions do not nest. The ! can be escaped with \ to suppress its special meaning.

Input lines containing history substitutions are echoed on the terminal after being expanded, but before any other substitutions take place or the command gets executed.

Event Designators

An event designator is a reference to a command line entry in the history list.

- ! Start a history substitution, except when followed by a space character, tab, newline, = or (.
- !! Refer to the previous command. By itself, this substitution repeats the previous command.
- !n Refer to command line *n*.
- !-n Refer to the current command line minus *n*.

Word Designators

- !str** Refer to the most recent command starting with *str*.
- !?str[?]** Refer to the most recent command containing *str*.
- !{...}** Insulate a history reference from adjacent characters (if necessary).

A ':' (colon) separates the event specification from the word designator. It can be omitted if the word designator begins with a ^, \$, *, - or %. If the word is to be selected from the previous command, the second ! character can be omitted from the event specification. For instance, **!!:1** and **!:1** both refer to the first word of the previous command, while **!!\$** and **!\$** both refer to the last word in the previous command. Word designators include:

- #** The entire command line typed so far.
- 0** The first input word (command).
- n** The *n*'th argument.
- ^** The first argument, that is, **1**.
- \$** The last argument.
- %** The word matched by (the most recent) *?s* search.
- x-y** A range of words; *-y* abbreviates **0-y**.
- *** All the arguments, or a null value if there is just one word in the event.
- x*** Abbreviates **x-\$**.
- x-** Like **x*** but omitting word **\$**.

Modifiers

After the optional word designator, you can add a sequence of one or more of the following modifiers, each preceded by a :

- h** Remove a trailing pathname component, leaving the head.
- r** Remove a trailing suffix of the form '.xxx', leaving the basename.
- e** Remove all but the suffix.
- s//r[/]** Substitute *r* for *l*.
- t** Remove all leading pathname components, leaving the tail.
- &** Repeat the previous substitution.
- g** Apply the change to the first occurrence of a match in each word, by prefixing the above (for example, **g&**).
- p** Print the new command but do not execute it.
- q** Quote the substituted words, escaping further substitutions.
- x** Like **q**, but break into words at each space character, tab or newline.

Unless preceded by a **g**, the modification is applied only to the first string that matches *l*; an error results if no string matches.

The left-hand side of substitutions are not regular expressions, but character strings. Any character can be used as the delimiter in place of /. A backslash quotes the delimiter character. The character **&**, in the right hand side, is replaced by the text from the left-hand-side. The **&** can be quoted with a backslash. A null *l* uses the previous string either from a *l* or from a contextual scan string *s* from *!s*. You can omit the rightmost delimiter if a newline immediately follows *r*; the rightmost ? in a context scan can similarly be omitted.

Without an event specification, a history reference refers either to the previous command, or to a previous history reference on the command line (if any).

Quick Substitution

`^r[r]` This is equivalent to the history substitution: `!s//r/[]`.

Aliases

The C shell maintains a list of aliases that you can create, display, and modify using the **alias** and **unalias** commands. The shell checks the first word in each command to see if it matches the name of an existing alias. If it does, the command is reprocessed with the alias definition replacing its name; the history substitution mechanism is made available as though that command were the previous input line. This allows history substitutions, escaped with a backslash in the definition, to be replaced with actual command line arguments when the alias is used. If no history substitution is called for, the arguments remain unchanged.

Aliases can be nested. That is, an alias definition can contain the name of another alias. Nested aliases are expanded before any history substitutions is applied. This is useful in pipelines such as

```
alias lm 'ls -l \!* | more'
```

which when called, pipes the output of `ls(1)` through `more(1)`.

Except for the first word, the name of the alias may not appear in its definition, nor in any alias referred to by its definition. Such loops are detected, and cause an error message.

I/O Redirection

The following metacharacters indicate that the subsequent word is the name of a file to which the command's standard input, standard output, or standard error is redirected; this word is variable, command, and filename expanded separately from the rest of the command.

- < Redirect the standard input.
- << *word* Read the standard input, up to a line that is identical with *word*, and place the resulting lines in a temporary file. Unless *word* is escaped or quoted, variable and command substitutions are performed on these lines. Then, the pipeline is invoked with the temporary file as its standard input. *word* is not subjected to variable, filename, or command substitution, and each line is compared to it before any substitutions are performed by the shell.
- > >! >& >&! Redirect the standard output to a file. If the file does not exist, it is created. If it does exist, it is overwritten; its previous contents are lost. When set, the variable **noclobber** prevents destruction of existing files. It also prevents redirection to terminals and `/dev/null`, unless one of the **!** forms is used. The **&** forms redirect both standard output and the standard error (diagnostic output) to the file.

>> >>& >>! >>&!

Append the standard output. Like >, but places output at the end of the file rather than overwriting it. If **noclobber** is set, it is an error for the file not to exist, unless one of the ! forms is used. The & forms append both the standard error and standard output to the file.

Variable Substitution

The C shell maintains a set of *variables*, each of which is composed of a *name* and a *value*. A variable name consists of up to 20 letters and digits, and starts with a letter (the underscore is considered a letter). A variable's value is a space-separated list of zero or more words.

To refer to a variable's value, precede its name with a '\$'. Certain references (described below) can be used to select specific words from the value, or to display other information about the variable. Braces can be used to insulate the reference from other characters in an input-line word.

Variable substitution takes place after the input line is analyzed, aliases are resolved, and I/O redirections are applied. Exceptions to this are variable references in I/O redirections (substituted at the time the redirection is made), and backquoted strings (see Command Substitution).

Variable substitution can be suppressed by preceding the \$ with a \, except within double-quotes where it always occurs. Variable substitution is suppressed inside of single-quotes. A \$ is escaped if followed by a space character, tab or newline.

Variables can be created, displayed, or destroyed using the **set** and **unset** commands. Some variables are maintained or used by the shell. For instance, the **argv** variable contains an image of the shell's argument list. Of the variables used by the shell, a number are toggles; the shell does not care what their value is, only whether they are set or not.

Numerical values can be operated on as numbers (as with the @ built-in command). With numeric operations, an empty value is considered to be zero; the second and subsequent words of multiword values are ignored. For instance, when the **verbose** variable is set to any value (including an empty value), command input is echoed on the terminal.

Command and filename substitution is subsequently applied to the words that result from the variable substitution, except when suppressed by double-quotes, when **noglob** is set (suppressing filename substitution), or when the reference is quoted with the **:q** modifier. Within double-quotes, a reference is expanded to form (a portion of) a quoted string; multiword values are expanded to a string with embedded space characters. When the **:q** modifier is applied to the reference, it is expanded to a list of space-separated words, each of which is quoted to prevent subsequent command or filename substitutions.

Except as noted below, it is an error to refer to a variable that is not set.

\$var

\${var}

These are replaced by words from the value of *var*, each separated by a space character. If *var* is an environment variable, its value is returned (but ':' modifiers and the other forms given below are not available).

$\$var[index]$
 $\${var[index]}$ These select only the indicated words from the value of *var*. Variable substitution is applied to *index*, which may consist of (or result in) a either single number, two numbers separated by a ‘-’, or an asterisk. Words are indexed starting from 1; a ‘*’ selects all words. If the first number of a range is omitted (as with **$\$argv[-2]$**), it defaults to 1. If the last number of a range is omitted (as with **$\$argv[1-]$**), it defaults to **$\$#var$** (the word count). It is not an error for a range to be empty if the second argument is omitted (or within range).

$\$#name$
 $\${#name}$ These give the number of words in the variable.
 $\$0$ This substitutes the name of the file from which command input is being read. An error occurs if the name is not known.

$\$n$
 $\${n}$ Equivalent to **$\$argv[n]$** .
 $\$*$ Equivalent to **$\$argv[*]$** .

The modifiers **:e**, **:h**, **:q**, **:r**, **:t**, and **:x** can be applied (see **History Substitution**), as can **:gh**, **:gt**, and **:gr**. If {} (braces) are used, then the modifiers must appear within the braces. The current implementation allows only one such modifier per expansion.

The following references may not be modified with **:** modifiers.

$\$?var$
 $\${?var}$ Substitutes the string 1 if *var* is set or 0 if it is not set.
 $\$?0$ Substitutes 1 if the current input filename is known or 0 if it is not.
 $\$$$ Substitute the process number of the (parent) shell.
 $\$<$ Substitutes a line from the standard input, with no further interpretation thereafter. It can be used to read from the keyboard in a C shell script.

**Command and
 Filename
 Substitutions**

Command and filename substitutions are applied selectively to the arguments of built-in commands. Portions of expressions that are not evaluated are not expanded. For non-built-in commands, filename expansion of the command name is done separately from that of the argument list; expansion occurs in a subshell, after I/O redirection is performed.

**Command
 Substitution**

A command enclosed by backquotes (`...`) is performed by a subshell. Its standard output is broken into separate words at each space character, tab and newline; null words are discarded. This text replaces the backquoted string on the current command line. Within double-quotes, only newline characters force new words; space and tab characters are preserved. However, a final newline is ignored. It is therefore possible for a command substitution to yield a partial word.

**Filename
Substitution**

Unquoted words containing any of the characters *, ?, [or {, or that begin with ~, are expanded (also known as *globbing*) to an alphabetically sorted list of filenames, as follows:

*	Match any (zero or more) characters.
?	Match any single character.
[...]	Match any single character in the enclosed list(s) or range(s). A list is a string of characters. A range is two characters separated by a dash (-), and includes all the characters in between in the ASCII collating sequence (see ascii(5)).
{ <i>str</i> , <i>str</i> , ... }	Expand to each string (or filename-matching pattern) in the comma-separated list. Unlike the pattern-matching expressions above, the expansion of this construct is not sorted. For instance, { b,a } expands to 'b' 'a', (not 'a' 'b'). As special cases, the characters { and }, along with the string {}, are passed undisturbed.
~[<i>user</i>]	Your home directory, as indicated by the value of the variable home , or that of <i>user</i> , as indicated by the password entry for <i>user</i> .

Only the patterns *, ? and [...] imply pattern matching; an error results if no filename matches a pattern that contains them. The '.' (dot character), when it is the first character in a filename or pathname component, must be matched explicitly. The / (slash) must also be matched explicitly.

**Expressions and
Operators**

A number of C shell built-in commands accept expressions, in which the operators are similar to those of C and have the same precedence. These expressions typically appear in the @, **exit**, **if**, **set** and **while** commands, and are often used to regulate the flow of control for executing commands. Components of an expression are separated by white space.

Null or missing values are considered 0. The result of all expressions is a string, which may represent decimal numbers.

The following C shell operators are grouped in order of precedence:

(...)	grouping
~	one's complement
!	logical negation
* / %	multiplication, division, remainder (These are right associative, which can lead to unexpected results. Group combinations explicitly with parentheses.)
+ -	addition, subtraction (also right associative)
<< >>	bitwise shift left, bitwise shift right
< > <= >=	less than, greater than, less than or equal to, greater than or equal to
= = ! = = ~ ! ~	equal to, not equal to, filename-substitution pattern match (described below), filename-substitution pattern mismatch

&	bitwise AND
^	bitwise XOR (exclusive or)
	bitwise inclusive OR
&&	logical AND
	logical OR

The operators: ==, !=, =~, and !~ compare their arguments as strings; other operators use numbers. The operators =~ and !~ each check whether or not a string to the left matches a filename substitution pattern on the right. This reduces the need for **switch** statements when pattern-matching between strings is all that is required.

Also available are file inquiries:

-r <i>filename</i>	Return true, or 1 if the user has read access. Otherwise it returns false, or 0.
-w <i>filename</i>	True if the user has write access.
-x <i>filename</i>	True if the user has execute permission (or search permission on a directory).
-e <i>filename</i>	True if <i>filename</i> exists.
-o <i>filename</i>	True if the user owns <i>filename</i> .
-z <i>filename</i>	True if <i>filename</i> is of zero length (empty).
-f <i>filename</i>	True if <i>filename</i> is a plain file.
-d <i>filename</i>	True if <i>filename</i> is a directory.

If *filename* does not exist or is inaccessible, then all inquiries return false.

An inquiry as to the success of a command is also available:

{ <i>command</i> }	If <i>command</i> runs successfully, the expression evaluates to true, 1. Otherwise, it evaluates to false, 0. (Note: Conversely, <i>command</i> itself typically returns 0 when it runs successfully, or some other value if it encounters a problem. If you want to get at the status directly, use the value of the status variable rather than this expression).
--------------------	---

Control Flow

The shell contains a number of commands to regulate the flow of control in scripts and within limits, from the terminal. These commands operate by forcing the shell either to reread input (to *loop*), or to skip input under certain conditions (to *branch*).

Each occurrence of a **foreach**, **switch**, **while**, **if...then** and **else** built-in command must appear as the first word on its own input line.

If the shell's input is not seekable and a loop is being read, that input is buffered. The shell performs seeks within the internal buffer to accomplish the rereading implied by the loop. (To the extent that this allows, backward **goto** commands will succeed on nonseekable inputs.)

Command Execution

If the command is a C shell built-in command, the shell executes it directly. Otherwise, the shell searches for a file by that name with execute access. If the command name contains a /, the shell takes it as a pathname, and searches for it. If the command name does not contain a /, the shell attempts to resolve it to a pathname, searching each directory in

the **path** variable for the command. To speed the search, the shell uses its hash table (see the **rehash** built-in command) to eliminate directories that have no applicable files. This hashing can be disabled with the **-c** or **-t**, options, or the **unhash** built-in command.

As a special case, if there is no / in the name of the script and there is an alias for the word **shell**, the expansion of the **shell** alias is prepended (without modification) to the command line. The system attempts to execute the first word of this special (late-occurring) alias, which should be a full pathname. Remaining words of the alias's definition, along with the text of the input line, are treated as arguments.

When a pathname is found that has proper execute permissions, the shell forks a new process and passes it, along with its arguments, to the kernel using the **execve()** system call (see **exec(2)**). The kernel then attempts to overlay the new process with the desired program. If the file is an executable binary (in **a.out(4)** format) the kernel succeeds and begins executing the new process. If the file is a text file and the first line begins with #!, the next word is taken to be the pathname of a shell (or command) to interpret that script. Subsequent words on the first line are taken as options for that shell. The kernel invokes (overlays) the indicated shell, using the name of the script as an argument.

If neither of the above conditions holds, the kernel cannot overlay the file and the **execve()** call fails (see **exec(2)**); the C shell then attempts to execute the file by spawning a new shell, as follows:

- If the first character of the file is a #, a C shell is invoked.
- Otherwise, a Bourne shell is invoked.

Signal Handling

The shell normally ignores QUIT signals. Background jobs are immune to signals generated from the keyboard, including hangups (HUP). Other signals have the values that the C shell inherited from its environment. The shell's handling of interrupt and terminate signals within scripts can be controlled by the **onintr** built-in command. Login shells catch the TERM signal; otherwise, this signal is passed on to child processes. In no case are interrupts allowed when a login shell is reading the **.logout** file.

Job Control

The shell associates a numbered *job* with each command sequence to keep track of those commands that are running in the background or have been stopped with TSTP signals (typically CTRL-z). When a command or command sequence (semicolon separated list) is started in the background using the **&** metacharacter, the shell displays a line with the job number in brackets and a list of associated process numbers:

```
[1] 1234
```

To see the current list of jobs, use the **jobs** built-in command. The job most recently stopped (or put into the background if none are stopped) is referred to as the *current* job and is indicated with a '+'. The previous job is indicated with a '-'; when the current job is terminated or moved to the foreground, this job takes its place (becomes the new current job).

To manipulate jobs, refer to the **bg**, **fg**, **kill**, **stop**, and **%** built-in commands.

A reference to a job begins with a '%'. By itself, the percent-sign refers to the current job.

% %+ %% The current job.

%- The previous job.

%*j* Refer to job *j* as in: '**kill -9 %*j***'. *j* can be a job number, or a string that uniquely specifies the command line by which it was started; '**fg %*vi***' might bring a stopped **vi** job to the foreground, for instance.

/?*string* Specify the job for which the command line uniquely contains *string*.

A job running in the background stops when it attempts to read from the terminal. Background jobs can normally produce output, but this can be suppressed using the '**stty tostop**' command.

Status Reporting

While running interactively, the shell tracks the status of each job and reports whenever the job finishes or becomes blocked. It normally displays a message to this effect as it issues a prompt, in order to avoid disturbing the appearance of your input. When set, the **notify** variable indicates that the shell is to report status changes immediately. By default, the **notify** command marks the current process; after starting a background job, type **notify** to mark it.

Built-In Commands

Built-in commands are executed within the C shell. If a built-in command occurs as any component of a pipeline except the last, it is executed in a subshell.

: Null command. This command is interpreted, but performs no action.

alias [*name* [*def*]]

Assign *def* to the alias *name*. *def* is a list of words that may contain escaped history-substitution metasyntax. *name* is not allowed to be **alias** or **unalias**. If *def* is omitted, the alias *name* is displayed along with its current definition. If both *name* and *def* are omitted, all aliases are displayed.

bg [%*job*...]

Run the current or specified jobs in the background.

break Resume execution after the **end** of the nearest enclosing **foreach** or **while** loop. The remaining commands on the current line are executed. This allows multilevel breaks to be written as a list of **break** commands, all on one line.

breaksw Break from a **switch**, resuming after the **endsw**.

case label: A label in a **switch** statement.

cd [*dir*]

chdir [*dir*]

Change the shell's working directory to directory *dir*. If no argument is given, change to the home directory of the user. If *dir* is a relative pathname not found in the current directory, check for it in those directories listed in the **cdpath** variable. If *dir* is the name of a shell variable whose value starts with a /, change to the directory named by that value.

continue Continue execution of the next iteration of the nearest enclosing **while** or **foreach** loop.

- default:** Labels the default case in a **switch** statement. The default should come after all **case** labels. Any remaining commands on the command line are first executed.
- dirs** [**-l**]
Print the directory stack, most recent to the left; the first directory shown is the current directory. With the **-l** argument, produce an unabbreviated printout; use of the **~** notation is suppressed.
- echo** [**-n**] *list*
The words in *list* are written to the shell's standard output, separated by space characters. The output is terminated with a newline unless the **-n** option is used.
- eval** *argument* ...
Reads the arguments as input to the shell and executes the resulting command(s). This is usually used to execute commands generated as the result of command or variable substitution, since parsing occurs before these substitutions. See **tset**(1B) for an example of how to use **eval**.
- exec** *command*
Execute *command* in place of the current shell, which terminates.
- exit** [(*expr*)]
The shell exits, either with the value of the status variable or with the value specified by the expression *expr* .
- fg** [%*job*]
Bring the current or specified *job* into the foreground.
- foreach** *var* (*wordlist*)
...
end The variable *var* is successively set to each member of *wordlist*. The sequence of commands between this command and the matching **end** is executed for each new value of *var*. Both **foreach** and **end** must appear alone on separate lines.

The built-in command **continue** may be used to terminate the execution of the current iteration of the loop and the built-in command **break** may be used to terminate execution of the **foreach** command. When this command is read from the terminal, the loop is read once prompting with **?** before any statements in the loop are executed.
- glob** *wordlist*
Perform filename expansion on *wordlist*. Like **echo**, but no **** escapes are recognized. Words are delimited by NULL characters in the output.
- goto** *label* The specified *label* is a filename and a command expanded to yield a label. The shell rewinds its input as much as possible and searches for a line of the form *label*: possibly preceded by space or tab characters. Execution continues after the indicated line. It is an error to jump to a label that occurs between a **while** or **for** built-in command and its corresponding **end**.

hashstat Print a statistics line indicating how effective the internal hash table has been at locating commands (and avoiding **execs**). An **exec** is attempted for each component of the *path* where the hash function indicates a possible hit and in each component that does not begin with a *'*.

history [**-hr**] [*n*]

Display the history list; if *n* is given, display only the *n* most recent events.

-r Reverse the order of printout to be most recent first rather than oldest first.

-h Display the history list without leading numbers. This is used to produce files suitable for sourcing using the **-h** option to *source*.

if (*expr*) *command*

If the specified expression evaluates to true, the single *command* with arguments is executed. Variable substitution on *command* happens early, at the same time it does for the rest of the *if* command. *command* must be a simple command, not a pipeline, a command list, or a parenthesized command list. Note: I/O redirection occurs even if *expr* is false, when *command* is *not* executed (this is a bug).

if (*expr*) **then**

...

else if (*expr2*) **then**

...

else

...

endif If *expr* is true, commands up to the first **else** are executed. Otherwise, if *expr2* is true, the commands between the **else if** and the second **else** are executed. Otherwise, commands between the **else** and the **endif** are executed. Any number of **else if** pairs are allowed, but only one **else**. Only one **endif** is needed, but it is required. The words **else** and **endif** must be the first nonwhite characters on a line. The **if** must appear alone on its input line or after an **else**.

jobs[-l] List the active jobs under job control.

-l List process IDs, in addition to the normal information.

kill [*-sig*] [*pid*] [*%job*] ...

kill **-l** Send the **TERM** (terminate) signal, by default, or the signal specified, to the specified process ID, the *job* indicated, or the current *job*. Signals are either given by number or by name. There is no default. Typing **kill** does not send a signal to the current job. If the signal being sent is **TERM** (terminate) or **HUP** (hangup), then the job or process is sent a **CONT** (continue) signal as well.

-l List the signal names that can be sent.

limit [**-h**] [*resource* [*max-use*]]

Limit the consumption by the current process or any process it spawns, each not to exceed *max-use* on the specified *resource*. If *max-use* is omitted, print the current limit; if *resource* is omitted, display all limits.

-h Use hard limits instead of the current limits. Hard limits impose a ceiling on the values of the current limits. Only the privileged user may raise the hard limits.

resource is one of:

cputime	Maximum CPU seconds per process.
filesize	Largest single file allowed.
datasize	Maximum data size (including stack) for the process.
stacksize	Maximum stack size for the process.
coredumpsize	Maximum size of a core dump (file).
descriptors	Maximum number of file descriptors.
memorysize	Maximum size of virtual memory.

max-use is a number, with an optional scaling factor, as follows:

nh	Hours (for cputime).
nk	<i>n</i> kilobytes. This is the default for all but cputime .
nm	<i>n</i> megabytes or minutes (for cputime).
mm:ss	Minutes and seconds (for cputime).

login [*username* | **-p**]

Terminate a login shell and invoke **login**(1). The **.logout** file is not processed. If *username* is omitted, **login** prompts for the name of a user.

-p Preserve the current environment (variables).

logout Terminate a login shell.

nice [*+n* | *-n*] [*command*]

Increase the process priority value for the shell or for *command* by *n*. The higher the priority value, the lower the priority of a process, and the slower it runs. When given, *command* is always run in a subshell, and the restrictions placed on commands in simple **if** commands apply. If *command* is omitted, **nice** increments the value for the current shell. If no increment is specified, **nice** sets the process priority value to 4. The range of process priority values is from -20 to 20. Values of *n* outside this range set the value to the lower, or to the higher boundary, respectively.

+n Increment the process priority value by *n*.

-n Decrement by *n*. This argument can be used only by the privileged user.

nohup [*command*]

Run *command* with HUPs ignored. With no arguments, ignore HUPs throughout the remainder of a script. When given, *command* is always run in a subshell, and the restrictions placed on commands in simple **if** statements apply. All processes detached with **&** are effectively **nohup**'d.

notify [*%job*] ...
 Notify the user asynchronously when the status of the current job or specified jobs changes.

onintr [- | *label*]
 Control the action of the shell on interrupts. With no arguments, **onintr** restores the default action of the shell on interrupts. (The shell terminates shell scripts and returns to the terminal command input level). With the - argument, the shell ignores all interrupts. With a *label* argument, the shell executes a **goto label** when an interrupt is received or a child process terminates because it was interrupted.

popd [+*n*] Pop the directory stack and **cd** to the new top directory. The elements of the directory stack are numbered from 0 starting at the top.
 +*n* Discard the *n*'th entry in the stack.

pushd [+*n* | *dir*]
 Push a directory onto the directory stack. With no arguments, exchange the top two elements.
 +*n* Rotate the *n*'th entry to the top of the stack and **cd** to it.
dir Push the current working directory onto the stack and change to *dir*.

rehash Recompute the internal hash table of the contents of directories listed in the *path* variable to account for new commands added.

repeat *count command*
 Repeat *command* *count* times. *command* is subject to the same restrictions as with the one-line **if** statement.

set [*var* [= *value*]]
set *var[n]* = *word*
 With no arguments, **set** displays the values of all shell variables. Multiword values are displayed as a parenthesized list. With the *var* argument alone, **set** assigns an empty (null) value to the variable *var*. With arguments of the form *var* = *value* **set** assigns *value* to *var*, where *value* is one of:

<i>word</i>	A single word (or quoted string).
<i>(wordlist)</i>	A space-separated list of words enclosed in parentheses.

Values are command and filename expanded before being assigned. The form **set** *var[n]* = *word* replaces the *n*'th word in a multiword value with *word*.

setenv [*VAR* [*word*]]
 With no arguments, **setenv** displays all environment variables. With the *VAR* argument, **setenv** sets the environment variable *VAR* to have an empty (null) value. (By convention, environment variables are normally given upper-case names.) With both *VAR* and *word* arguments, **setenv** sets the environment variable *NAME* to the value *word*, which must be either a single word or a quoted string. The most commonly used environment variables, **USER**, **TERM**, and **PATH**, are automatically imported to and exported from the **csh** variables **user**, **term**, and **path**; there is no need to use **setenv** for these. In

addition, the shell sets the **PWD** environment variable from the **csh** variable **cwd** whenever the latter changes.

The environment variables **LC_CTYPE**, **LC_MESSAGES**, **LC_TIME**, **LC_COLLATE**, **LC_NUMERIC**, and **LC_MONETARY** take immediate effect when changed within the C shell.

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S./English).

LC_NUMERIC

Determines the value of the radix character (decimal point (".") in the "C" locale) and thousand separator (empty string ("") in the "C" locale).

shift [*variable*]

The components of **argv**, or *variable*, if supplied, are shifted to the left, discarding the first component. It is an error for the variable not to be set or to have a null value.

source [**-h**] *name*
 Reads commands from *name*. **source** commands may be nested, but if they are nested too deeply the shell may run out of file descriptors. An error in a sourced file at any level terminates all nested **source** commands.

-h Place commands from the file *name* on the history list without executing them.

stop %*jobid* ...
 Stop the current or specified background job.

stop *pid* ...
 Stop the specified process, *pid*. (see **ps**(1)).

suspend Stop the shell in its tracks, much as if it had been sent a stop signal with **^Z**. This is most often used to stop shells started by **su**.

switch (*string*)
case *label*:
 ...
breaksw
 ...
default:
 ...
breaksw
endsw

Each *label* is successively matched, against the specified *string*, which is first command and filename expanded. The file metacharacters *, ? and [...] may be used in the case labels, which are variable expanded. If none of the labels match before a "default" label is found, execution begins after the default label. Each **case** statement and the **default** statement must appear at the beginning of a line. The command **breaksw** continues execution after the **endsw**. Otherwise control falls through subsequent **case** and **default** statements as with C. If no label matches and there is no default, execution continues after the **endsw**.

time [*command*]
 With no argument, print a summary of time used by this C shell and its children. With an optional *command*, execute *command* and print a summary of the time it uses.

umask [*value*]
 Display the file creation mask. With *value*, set the file creation mask. With *value* given in octal, the user can turn-off any bits, but cannot turn-on bits to allow new permissions. Common values include 077, restricting all permissions from everyone else; 002, giving complete access to the group, and read (and directory search) access to others; or 022, giving read (and directory search) but not write permission to the group and others.

unalias *pattern*
 Discard aliases that match (filename substitution) *pattern*. All aliases are

removed by **unalias ***.

unhash Disable the internal hash table.

unlimit [**-h**] [*resource*]

Remove a limitation on *resource*. If no *resource* is specified, then all resource limitations are removed. See the description of the **limit** command for the list of resource names.

-h Remove corresponding hard limits. Only the privileged user may do this.

unset *pattern*

Remove variables whose names match (filename substitution) *pattern*. All variables are removed by **unset ***; this has noticeably distasteful side effects.

unsetenv *variable*

Remove *variable* from the environment. As with **unset**, pattern matching is not performed.

wait Wait for background jobs to finish (or for an interrupt) before prompting.

while (*expr*)

...

end While *expr* is true (evaluates to nonzero), repeat commands between the **while** and the matching **end** statement. **break** and **continue** may be used to terminate or continue the loop prematurely. The **while** and **end** must appear alone on their input lines. If the shell's input is a terminal, it prompts for commands with a question-mark until the **end** command is entered and then performs the commands in the loop.

%[*job*] [**&]**

Bring the current or indicated *job* to the foreground. With the ampersand, continue running *job* in the background.

@ [*var* =*expr*]

@ [*var*[*n*] =*expr*]

With no arguments, display the values for all shell variables. With arguments, set the variable *var*, or the *n*'th word in the value of *var*, to the value that *expr* evaluates to. (If [*n*] is supplied, both *var* and its *n*'th component must already exist.)

If the expression contains the characters **>**, **<**, **&**, or **|**, then at least this part of *expr* must be placed within parentheses.

The operators ***=**, **+=**, and so forth, are available as in C. The space separating the name from the assignment operator is optional. Spaces are, however, mandatory in separating components of *expr* that would otherwise be single words.

Special postfix operators, **++** and **--**, increment or decrement *name*, respectively.

**Environment
Variables and
Predefined Shell
Variables**

Unlike the Bourne shell, the C shell maintains a distinction between environment variables, which are automatically exported to processes it invokes, and shell variables, which are not. Both types of variables are treated similarly under variable substitution. The shell sets the variables **argv**, **cwd**, **home**, **path**, **prompt**, **shell**, and **status** upon initialization. The shell copies the environment variable **USER** into the shell variable **user**, **TERM** into **term**, and **HOME** into **home**, and copies each back into the respective environment variable whenever the shell variables are reset. **PATH** and **path** are similarly handled. You need only set **path** once in the **.cshrc** or **.login** file. The environment variable **PWD** is set from **cwd** whenever the latter changes. The following shell variables have predefined meanings:

argv	Argument list. Contains the list of command line arguments supplied to the current invocation of the shell. This variable determines the value of the positional parameters \$1 , \$2 , and so on.
cdpath	Contains a list of directories to be searched by the cd , chdir , and popd commands, if the directory argument each accepts is not a subdirectory of the current directory.
cwd	The full pathname of the current directory.
echo	Echo commands (after substitutions) just before execution.
ignore	A list of filename suffixes to ignore when attempting filename completion. Typically the single word ‘.o’ .
filec	Enable filename completion, in which case the CTRL-d character EOT and the ESC character have special significance when typed in at the end of a terminal input line: <ul style="list-style-type: none"> EOT Print a list of all filenames that start with the preceding string. ESC Replace the preceding string with the longest unambiguous extension.
hardpaths	If set, pathnames in the directory stack are resolved to contain no symbolic-link components.
histchars	A two-character string. The first character replaces ! as the history-substitution character. The second replaces the carat (^) for quick substitutions.
history	The number of lines saved in the history list. A very large number may use up all of the C shell's memory. If not set, the C shell saves only the most recent command.
home	The user's home directory. The filename expansion of ~ refers to the value of this variable.
ignoreeof	If set, the shell ignores EOF from terminals. This protects against accidentally killing a C shell by typing a CTRL-d.

mail	A list of files where the C shell checks for mail. If the first word of the value is a number, it specifies a mail checking interval in seconds (default 5 minutes).
nobeep	Suppress the bell during command completion when asking the C shell to extend an ambiguous filename.
noclobber	Restrict output redirection so that existing files are not destroyed by accident. > redirections can only be made to new files. >> redirections can only be made to existing files.
noglob	Inhibit filename substitution. This is most useful in shell scripts once filenames (if any) are obtained and no further expansion is desired.
nonomatch	Returns the filename substitution pattern, rather than an error, if the pattern is not matched. Malformed patterns still result in errors.
notify	If set, the shell notifies you immediately as jobs are completed, rather than waiting until just before issuing a prompt.
path	The list of directories in which to search for commands. path is initialized from the environment variable PATH , which the C shell updates whenever path changes. A null word specifies the current directory. The default is typically: (/bin /usr/bin /usr/ucb /etc .). If path becomes unset only full pathnames will execute. An interactive C shell will normally hash the contents of the directories listed after reading .cshrc , and whenever path is reset. If new commands are added, use the rehash command to update the table.
prompt	The string an interactive C shell prompts with. Noninteractive shells leave the prompt variable unset. Aliases and other commands in the .cshrc file that are only useful interactively, can be placed after the following test: ' if (\$?prompt == 0) exit ', to reduce startup time for noninteractive shells. A ! in the prompt string is replaced by the current event number. The default prompt is <i>hostname%</i> for mere mortals, or <i>hostname#</i> for the privileged user.
savehist	The number of lines from the history list that are saved in ~/.history when the user logs out. Large values for savehist slow down the C shell during startup.
shell	The file in which the C shell resides. This is used in forking shells to interpret files that have execute bits set, but that are not executable by the system.
status	The status returned by the most recent command. If that command terminated abnormally, 0200 is added to the status. Built-in commands that fail return exit status 1; all other built-in commands set status to 0.
time	Control automatic timing of commands. Can be supplied with one or two values. The first is the reporting threshold in CPU seconds. The second is a string of tags and text indicating which resources to

report on. A tag is a percent sign (%) followed by a single upper-case letter (unrecognized tags print as text):

- %D** Average amount of unshared data space used in Kilobytes.
- %E** Elapsed (wallclock) time for the command.
- %F** Page faults.
- %I** Number of block input operations.
- %K** Average amount of unshared stack space used in Kilobytes.
- %M** Maximum real memory used during execution of the process.
- %O** Number of block output operations.
- %P** Total CPU time — U (user) plus S (system) — as a percentage of E (elapsed) time.
- %S** Number of seconds of CPU time consumed by the kernel on behalf of the user's process.
- %U** Number of seconds of CPU time devoted to the user's process.
- %W** Number of swaps.
- %X** Average amount of shared memory used in Kilobytes.

The default summary display outputs from the %U, %S, %E, %P, %X, %D, %I, %O, %F, and %W tags, in that order.

verbose Display each command after history substitution takes place.

FILES	<p>~/.cshrc Read at beginning of execution by each shell.</p> <p>~/.login Read by login shells after .cshrc at login.</p> <p>~/.logout Read by login shells at logout.</p> <p>~/.history Saved history for use at next login.</p> <p>/usr/bin/sh The Bourne shell, for shell scripts not starting with a '#'. Temporary file for '<<'.</p> <p>/etc/passwd Source of home directories for '~name'.</p>
--------------	--

SEE ALSO **login(1), ps(1), sh(1), shell_builtins(1), tset(1B), access(2), exec(2), fork(2), pipe(2), a.out(4), environ(4), environ(5), ascii(5), termio(7)**

DIAGNOSTICS **You have stopped jobs.**
You attempted to exit the C shell with stopped jobs under job control. An immediate second attempt to exit will succeed, terminating the stopped jobs.

NOTES Words can be no longer than 1024 characters. The system limits argument lists to 1,048,576 characters. However, the maximum number of arguments to a command for which filename expansion applies is 1706. Command substitutions may expand to no more characters than are allowed in the argument list. To detect looping, the shell restricts the number of **alias** substitutions on a single line to 20.

When a command is restarted from a stop, the shell prints the directory it started in if this is different from the current directory; this can be misleading (that is, wrong) as the job may have changed directories internally.

Shell built-in functions are not stoppable/restartable. Command sequences of the form *a ; b ; c* are also not handled gracefully when stopping is attempted. If you suspend *b*, the shell never executes *c*. This is especially noticeable if the expansion results from an alias. It can be avoided by placing the sequence in parentheses to force it into a subshell.

Control over terminal output after processes are started is primitive; use the Sun Window system if you need better output control.

Multiline shell procedures should be provided, as they are with the Bourne shell.

Commands within loops, prompted for by *?*, are not placed in the *history* list.

Control structures should be parsed rather than being recognized as built-in commands. This would allow control commands to be placed anywhere, to be combined with *|*, and to be used with *&* and *;* metasyntax.

It should be possible to use the *:* modifiers on the output of command substitutions.

There are two problems with *:* modifier usage on variable substitutions: not all of the modifiers are available, and only one modifier per substitution is allowed.

The *g* (global) flag in history substitutions applies only to the first match in each word, rather than all matches in all words. The common text editors consistently do the latter when given the *g* flag in a substitution command.

Quoting conventions are confusing. Overriding the escape character to force variable substitutions within double quotes is counterintuitive and inconsistent with the Bourne shell.

Symbolic links can fool the shell. Setting the **hardpaths** variable alleviates this.

'**set path**' should remove duplicate pathnames from the pathname list. These often occur because a shell script or a **.cshrc** file does something like '**set path=(/usr/local /usr/hosts **\$path**)**' to ensure that the named directories are in the pathname list.

The only way to direct the standard output and standard error separately is by invoking a subshell, as follows:

example% (*command* > *outfile*) >& *errorfile*

Although robust enough for general use, adventures into the esoteric periphery of the C shell may reveal unexpected quirks.

If you start **csh** as a login shell and you do not have a **.login** in your home directory, then the **csh** reads in the **/etc/.login**.

NAME	csplit – split a file with respect to a given context
SYNOPSIS	csplit [-s] [-k] [-f prefix] <i>filename</i> <i>argument1</i> [... <i>argumentn</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	<p>csplit reads <i>filename</i> and separates it into <i>n</i>+1 sections, defined by the arguments <i>argument1</i> ... <i>argumentn</i>. By default the sections are placed in xx00 ... xxn (<i>n</i> may not be greater than 99). These sections get the following pieces of <i>filename</i>:</p> <p>00: From the start of <i>filename</i> up to (but not including) the line referenced by <i>argument1</i>.</p> <p>01: From the line referenced by <i>argument1</i> up to the line referenced by <i>argument2</i>.</p> <p>⋮</p> <p><i>n</i>: From the line referenced by <i>argumentn</i> to the end of <i>filename</i>.</p> <p>If the <i>filename</i> argument is a -, then standard input is used.</p> <p>The arguments (<i>argument1</i> ... <i>argumentn</i>) to csplit can be a combination of the following:</p> <p><i>/rexp/</i> A file is to be created for the section from the current line up to (but not including) the line containing the regular expression <i>rexp</i>. The current line becomes the line containing <i>rexp</i>. This argument may be followed by an optional + or - some number of lines (for example, /Page/-5). See ed(1) for a description of how to specify a regular expression.</p> <p>%rexp% This argument is the same as <i>/rexp/</i>, except that no file is created for the section.</p> <p><i>Inno</i> A file is to be created from the current line up to (but not including) <i>Inno</i>. The current line becomes <i>Inno</i>.</p> <p>{num} Repeat argument. This argument may follow any of the above arguments. If it follows a <i>rexp</i> type argument, that argument is applied <i>num</i> more times. If it follows <i>Inno</i>, the file will be split every <i>Inno</i> lines (<i>num</i> times) from that point.</p> <p>Enclose all <i>rexp</i> type arguments that contain blanks or other characters meaningful to the shell in the appropriate quotes. Regular expressions may not contain embedded newlines. csplit does not affect the original file; it is the user's responsibility to remove it if it is no longer wanted.</p>
OPTIONS	<p>-s csplit normally prints the character counts for each file created. If the -s option is present, csplit suppresses the printing of all character counts.</p> <p>-k csplit normally removes created files if an error occurs. If the -k option is present, csplit leaves previously created files intact.</p> <p>-f prefix If the -f option is used, the created files are named <i>prefix00</i> ... <i>prefixn</i>. The default is xx00 ... xxn.</p>

EXAMPLES

This example creates four files, **cobol00...cobol03**.

```
example% csplit -f cobol filename '/procedure division' /par5./ /par16./
```

After editing the "split" files, they can be recombined as follows:

```
example% cat cobol0[0-3] > filename
```

Note: This example overwrites the original file.

This example splits the file at every 100 lines, up to 10,000 lines. The **-k** option causes the created files to be retained if there are less than 10,000 lines; however, an error message would still be printed.

```
example% csplit -k filename 100 {99}
```

If **prog.c** follows the normal C coding convention (the last line of a routine consists only of a **}** in the first character position), this example creates a file for each separate C routine (up to 21) in **prog.c**.

```
example% csplit -k prog.c '%main(%`'~/+1' {20}
```

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

ed(1), **sh(1)**, **environ(5)**

DIAGNOSTICS

Self-explanatory except for:

```
arg - out of range
```

which means that the given argument did not reference a line between the current position and the end of the file.

NAME	ct – spawn login to a remote terminal
SYNOPSIS	ct [<i>options</i>] <i>telno</i> . . .
AVAILABILITY	SUNWbnuu
DESCRIPTION	<p>ct dials the telephone number of a modem that is attached to a terminal and spawns a login process to that terminal. The <i>telno</i> is a telephone number, with equal signs for secondary dial tones and minus signs for delays at appropriate places. (The set of legal characters for <i>telno</i> is 0 through 9, -, =, *, and #. The maximum length <i>telno</i> is 31 characters). If more than one telephone number is specified, ct will try each in succession until one answers; this is useful for specifying alternate dialing paths.</p> <p>ct will try each line listed in the file /etc/uucp/Devices until it finds an available line with appropriate attributes, or runs out of entries.</p> <p>After the user on the destination terminal logs out, there are two things that could occur depending on what type of port monitor is monitoring the port. In the case of no port monitor, ct prompts: Reconnect? If the response begins with the letter n, the line will be dropped; otherwise, <i>ttymon</i> will be started again and the login: prompt will be printed. In the second case, where a port monitor is monitoring the port, the port monitor reissues the login: prompt.</p> <p>The user should log out properly before disconnecting.</p>
OPTIONS	<p>-h Normally, ct will hang up the current line so that it can be used to answer the incoming call. The -h option will prevent this action. The -h option will also wait for the termination of the specified ct process before returning control to the user's terminal.</p> <p>-speed The data rate may be set with the -s option. <i>speed</i> is expressed in baud rates. The default baud rate is 1200.</p> <p>-v If the -v (verbose) option is used, ct will send a running narrative to the standard error output stream.</p> <p>-wn If there are no free lines ct will ask if it should if so, for how many minutes it should wait before it gives up. ct will continue to try to open the dialers at one-minute intervals until the specified limit is exceeded. This dialogue may be overridden by specifying the -wn option. <i>n</i> is the maximum number of minutes that ct is to wait for a line.</p> <p>-xn This option is used for debugging; it produces a detailed output of the program execution on stderr. <i>n</i> is a single number between 0 and 9. As <i>n</i> increases to 9, more detailed debugging information is given.</p>
FILES	<p>/etc/uucp/Devices</p> <p>/var/adm/ctlog</p>

SEE ALSO `cu(1C)`, `login(1)`, `uucp(1C)`, `ttymon(1M)`

NOTES The `ct` program will not work with a DATAKIT Multiplex interface.
For a shared port, one used for both dial-in and dial-out, the `ttymon` program running on the line must have the `-r` and `-b` options specified (see `ttymon(1M)`).

NAME	ctags – create a tags file for use with ex and vi
SYNOPSIS	ctags [-aBFtuvwx] [-f tagsfile] <i>filename</i> . . .
DESCRIPTION	<p>ctags makes a tags file for ex(1) from the specified C, C++, Pascal, FORTRAN, YACC, and LEX sources. A tags file gives the locations of specified objects (in this case functions and typedefs) in a group of files. Each line of the tags file contains the object name, the file in which it is defined, and an address specification for the object definition. Functions are searched with a pattern, typedefs with a line number. Specifiers are given in separate fields on the line, separated by SPACE or TAB characters. Using the tags file, ex can quickly find these objects definitions.</p> <p>Normally ctags places the tag descriptions in a file called tags; this may be overridden with the -f option.</p> <p>Files with names ending in .c or .h are assumed to be either C or C++ source files and are searched for C/C++ routine and macro definitions. Files with names ending in .cc, .C, or .cxx, are assumed to be C++ source files. Files with names ending in .y are assumed to be YACC source files. Files with names ending in .l are assumed to be LEX files. Others are first examined to see if they contain any Pascal or FORTRAN routine definitions; if not, they are processed again looking for C definitions.</p> <p>The tag main is treated specially in C or C++ programs. The tag formed is created by prepending M to <i>filename</i>, with a trailing .c, .cc, .C, or .cxx removed, if any, and leading pathname components also removed. This makes use of ctags practical in directories with more than one program.</p>
OPTIONS	<p>-a Append output to an existing tags file.</p> <p>-B Use backward searching patterns (?...?).</p> <p>-F Use forward searching patterns (/.../) (default).</p> <p>-t Create tags for typedefs.</p> <p>-u Update the specified files in tags, that is, all references to them are deleted, and the new values are appended to the file. Beware: this option is implemented in a way which is rather slow; it is usually faster to simply rebuild the tags file.</p> <p>-v Produce on the standard output an index listing the function name, file name, and page number (assuming 64 line pages). Since the output will be sorted into lexicographic order, it may be desired to run the output through sort -f.</p> <p>-w Suppress warning diagnostics.</p> <p>-x Produce a list of object names, the line number and file name on which each is defined, as well as the text of that line and prints this on the standard output. This is a simple index which can be printed out as an off-line readable function index.</p> <p>-f tagsfile Places the tag descriptions in a file called <i>tagsfile</i> instead of tags.</p>

USAGE	The -v option is mainly used with vgrind which will be part of the optional BSD Compatibility Package.
EXAMPLES	<p>Using ctags with the -v option produces entries in an order which may not always be appropriate for vgrind. To produce results in alphabetical order, you may want to run the output through 'sort -f'.</p> <pre>example% ctags -v filename.c filename.h sort -f > index example% vgrind -x index</pre> <p>To build a tags file for C sources in a directory hierarchy rooted at <i>sourcedir</i>, first create an empty tags file, and then run find(1):</p> <pre>example% cd sourcedir ; rm -f tags ; touch tags example% find . \(-name SCCS -prune -name \ '*.c' -o -name '*.h' \) -exec ctags -u {} \;</pre> <p>Note that spaces must be entered exactly as shown.</p>
FILES	tags output tags file
SEE ALSO	ex(1) , vgrind(1) , vi(1)
NOTES	<p>Recognition of functions, subroutines and procedures for FORTRAN and Pascal is done in a very simpleminded way. No attempt is made to deal with block structure; if you have two Pascal procedures in different blocks with the same name you lose.</p> <p>The method of deciding whether to look for C or Pascal and FORTRAN functions is a hack.</p> <p>ctags does not know about #ifdefs.</p> <p>ctags should know about Pascal types. Relies on the input being well formed to detect typedefs. Use of -tx shows only the last line of typedefs.</p>

NAME	cu – call another UNIX system
SYNOPSIS	cu [<i>-c device</i> <i>-l line</i>] [<i>-s speed</i>] [<i>-b bits</i>] [<i>-h</i>] [<i>-n</i>] [<i>-t</i>] [<i>-d</i>] [<i>-o</i> <i>-e</i>] [<i>-L</i>] [<i>-C</i>] [<i>-H</i>] <i>telno</i> <i>systemname</i> [<i>local-cmd</i>]
AVAILABILITY	SUNWbnuu
DESCRIPTION	cu calls up another UNIX system, a terminal, or possibly a non-UNIX system. It manages an interactive conversation with possible transfers of files. It is convenient to think of cu as operating in two phases. The first phase is the connection phase in which the connection is established. cu then enters the conversation phase. The -d option is the only one that applies to both phases.
OPTIONS	cu accepts many options. The -c , -l , and -s options play a part in selecting the medium; the remaining options are used in configuring the line. <ul style="list-style-type: none"> -c device The first field in the /etc/uucp/Devices file is the "Type" field. The -c option forces cu to only use entries in the "Type" field that match the user specified <i>device</i>. The specified <i>device</i> is usually the name of a local area network. -s speed Specifies the transmission speed (300, 1200, 2400 9600, 19200, 38400). The default value is "Any" speed which will depend on the order of the lines in the /etc/uucp/Devices file. -l line Specifies a device name to use as the communication line. This can be used to override the search that would otherwise take place for the first available line having the right speed. When the -l option is used without the -s option, the speed of a line is taken from the /etc/uucp/Devices file record in which line matches the second field (the Line field). When the -l and -s options are both used together, cu will search the /etc/uucp/Devices file to check if the requested speed for the requested line is available. If so, the connection will be made at the requested speed, otherwise, an error message will be printed and the call will not be made. In the general case where a specified device is a directly connected asynchronous line (for instance, /dev/term/a), a telephone number (<i>telno</i>) is not required. The specified device need not be in the /dev directory. If the specified device is associated with an auto dialer, a telephone number must be provided. -b bits Forces <i>bits</i> to be the number of bits processed on the line. <i>bits</i> is either 7 or 8. This allows connection between systems with different character sizes. By default, the character size of the line is set to the same as the current local terminal. -h Set communication mode to half-duplex. This option emulates local echo in order to support calls to other computer systems that expect terminals to be set to half-duplex mode.

- n Request user prompt for telephone number. For added security, this option will prompt the user to provide the telephone number to be dialed, rather than taking it from the command line.
- t Used to dial a terminal which has been set to auto answer. Appropriate mapping of carriage-return to carriage-return-line-feed pairs is set.
- d Diagnostic traces are printed.
- o Set an ODD data parity. This option designates that ODD parity is to be generated for data sent to the remote system.
- e Set an EVEN data parity. This option designates that EVEN parity is to be generated for data sent to the remote system.
- L will cause **cu** to go through the login chat sequence specified in the **/etc/uucp/Systems** file. For more information about the chat sequence, see *TCP/IP Network Administration Guide*
- C will cause **cu** to run the local command that is specified at the end of the command line, instead of entering interactive mode. The **stdin** and **stdout** of the command that is run refer to the remote connection.
- H ignore one hangup. This allows the user to remain in **cu** while the remote machine disconnects and places a call back to the local machine. This option should be used when connecting to systems with callback or dialback modems. Once the callback occurs subsequent hangups will cause **cu** to terminate. This option can be specified more than once. For more information about dialback configuration, see **remote(4)** and *TCP/IP Network Administration Guide*

USAGE

Connection Phase

cu uses the same mechanism that **uucp(1C)** does to establish a connection. This means that it will use the **uucp** control files **/etc/uucp/Devices** and **/etc/uucp/Systems**. This gives **cu** the ability to choose from several different media to establish the connection. The possible media include telephone lines, direct connections, and local area networks (LAN). The **/etc/uucp/Devices** file contains a list of media that are available on your system. The **/etc/uucp/Systems** file contains information for connecting to remote systems, but it is not generally readable.

Note: **cu** determines which **/etc/uucp/Systems** and **/etc/uucp/Devices** files to use based upon the name used to invoke **cu**. In the simple case, this name will be "cu", but you could also have created a link to **cu** with another name, e.g., "pppcu", in which case **cu** would then look for a "service=pppcu" entry in the **/etc/uucp/Sysfiles** file, which will tell it which **/etc/uucp/Systems** file to use.

The *telno* or *systemname* parameter from the command line is used to tell **cu** what system you wish to connect to. This parameter can be blank, a telephone number, a system name, or a LAN specific address.

telephone number A telephone number is a string consisting of the tone dial characters (the digits **0** through **9**, *, and #) plus the special characters = and -. The equal sign designates a secondary dial tone and the

minus sign creates a 4 second delay.

system name A system name is the name of any computer that **uucp** can call; the **uname(1C)** command prints a list of these names.

LAN address The documentation for your LAN will show the form of the LAN specific address.

If **cu**'s default behavior is invoked (not using the **-c** or **-l** options), **cu** will use the *telno* or *systemname* parameter to determine which medium to use. If a telephone number is specified, **cu** will assume that you wish to use a telephone line and it will select an automatic call unit (ACU). Otherwise, **cu** will assume that it is a system name. **cu** will follow the **uucp** calling mechanism and use the **/etc/uucp/Systems** and **/etc/uucp/Devices** files to obtain the best available connection. Since **cu** will choose a speed that is appropriate for the medium that it selects, you may not use the **-s** option when this parameter is a system name.

The **-c** and **-l** options modify this default behavior. **-c** is most often used to select a LAN by specifying a Type field from the **/etc/uucp/Devices** file. You must include either a *telno* or *systemname* value when using the **-c** option. If the connection to *systemname* fails, a connection will be attempted using *systemname* as a LAN specific address. The **-l** option is used to specify a device associated with a direct connection. If the connection is truly a direct connection to the remote machine, then there is no need to specify a *systemname*. This is the only case where a *telno* or *systemname* parameter is unnecessary. On the other hand, there may be cases in which the specified device connects to a dialer, so it is valid to specify a telephone number. The **-c** and **-l** options should not be specified on the same command line.

Conversation Phase

After making the connection, **cu** runs as two processes: the *transmit* process reads data from the standard input and, except for lines beginning with `~`, passes it to the remote system; the *receive* process accepts data from the remote system and, except for lines beginning with `~`, passes it to the standard output. Normally, an automatic DC3/DC1 protocol is used to control input from the remote so the buffer is not overrun. Lines beginning with `~` have special meanings.

Commands

The *transmit* process interprets the following user initiated commands:

`~.` Terminate the conversation.

`~!` Escape to an interactive shell on the local system.

`~!cmd...` Run *cmd* on the local system (via **sh -c**).

`~$cmd...` Run *cmd* locally and send its output to the remote system.

`~%cd` Change the directory on the local system. Note: `~!cd` will cause the command to be run by a sub-shell, probably not what was intended.

`~%take from [to]` Copy file *from* (on the remote system) to file *to* on the local system. If *to* is omitted, the *from* argument is used in both places.

<code>~%put from [to]</code>	Copy file <i>from</i> (on local system) to file <i>to</i> on remote system. If <i>to</i> is omitted, the <i>from</i> argument is used in both places.
<code>~ line</code>	Send the line <code>~ line</code> to the remote system.
<code>~%break</code>	Transmit a BREAK to the remote system (which can also be specified as <code>~%b</code>).
<code>~%debug</code>	Toggles the <code>-d</code> debugging option on or off (which can also be specified as <code>~%d</code>).
<code>~t</code>	Prints the values of the termio structure variables for the user's terminal (useful for debugging).
<code>~l</code>	Prints the values of the termio structure variables for the remote communication line (useful for debugging).
<code>~%ifc</code>	Toggles between DC3/DC1 input control protocol and no input control. This is useful when the remote system does not respond properly to the DC3 and DC1 characters (can also be specified as <code>~%nostop</code>).
<code>~%ofc</code>	Toggles the output flow control setting. When enabled, outgoing data may be flow controlled by the remote host (can also be specified as <code>~%noostop</code>).
<code>~%divert</code>	Allow/disallow unsolicited diversions. That is, diversions not specified by <code>~%take</code> .
<code>~%old</code>	Allow/disallow old style syntax for received diversions.

The *receive* process normally copies data from the remote system to the standard output of the local system. It may also direct the output to local files.

The use of `~%put` requires `stty(1)` and `cat(1)` on the remote side. It also requires that the current erase and kill characters on the remote system be identical to these current control characters on the local system. Backslashes are inserted at appropriate places.

The use of `~%take` requires the existence of `echo(1)` and `cat(1)` on the remote system, and that the remote system must be using the Bourne shell, `sh`. Also, `tabs` mode (see `stty(1)`) should be set on the remote system if tabs are to be copied without expansion to spaces.

When `cu` is used on system X to connect to system Y and subsequently used on system Y to connect to system Z, commands on system Y can be executed by using `~`. Executing a tilde command reminds the user of the local system `uname`. For example, `uname` can be executed on Z, X, and Y as follows:

```

uname
Z
~[X]!uname
X
~~[Y]!uname
Y

```

In general, `~` causes the command to be executed on the original machine. `~~` causes the

command to be executed on the next machine in the chain.

EXAMPLES

To dial a system whose telephone number is **9 1 201 555 1234** using **1200** baud (where dialtone is expected after the **9**):

```
example% cu -s 1200 9=12015551234
```

If the speed is not specified, "Any" is the default value.

To login to a system connected by a direct line:

```
example% cu -l /dev/term/b
```

or

```
example% cu -l term/b
```

To dial a system with a specific line and speed:

```
example% cu -s 1200 -l term/b
```

To use a system name:

```
example% cu systemname
```

RETURN VALUES

If successful, **cu** returns **0**; otherwise, **-1** is returned.

FILES

```
/etc/uucp/Devices  
/etc/uucp/Sysfiles  
/etc/uucp/Systems  
/var/spool/locks/*
```

SEE ALSO

cat(1), **stty(1)**, **uname(1)**, **ct(1C)**, **uname(1C)**, **uucp(1C)**
TCP/IP Network Administration Guide

NOTES

The **cu** command does not do any integrity checking on data it transfers. Data fields with special **cu** characters may not be transmitted properly. Depending on the interconnection hardware, it may be necessary to use a **~**. to terminate the conversion, even if **stty 0** has been used. Non-printing characters are not dependably transmitted using either the **~%put** or **~%take** commands. **~%put** and **~%take** cannot be used over multiple links. Files must be moved one link at a time.

There is an artificial slowing of transmission by **cu** during the **~%put** operation so that loss of data is unlikely. Files transferred using **~%take** or **~%put** must contain a trailing newline, otherwise, the operation will hang. Entering a CTRL-D command usually clears the hang condition.

NAME	cut – cut out selected fields of each line of a file
SYNOPSIS	<pre>cut -b list [-n] [filename ...] cut -c list [filename ...] cut -f list [-d delim] [-s] [filename]</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>Use cut to cut out columns from a table or fields from each line of a file; in data base parlance, it implements the projection of a relation. The fields as specified by <i>list</i> can be fixed length, that is, character positions as on a punched card (-c option) or the length can vary from line to line and be marked with a field delimiter character like <i>tab</i> (-f option). cut can be used as a filter; if no files are given, the standard input is used. In addition, a file name of “-” explicitly refers to standard input.</p> <p>Either the -b, -c, or -f option must be specified.</p> <p>Use grep(1) to make horizontal “cuts” (by context) through a file, or paste(1) to put files together column-wise (that is, horizontally). To reorder columns in a table, use cut and paste.</p>
OPTIONS	<p><i>list</i> A comma-separated list of integer field numbers (in increasing order), with optional – to indicate ranges (for instance, 1,4,7; 1–3,8; –5,10 (short for 1–5,10); or 3– (short for third through last field)).</p> <p>-b list The <i>list</i> following -b specifies byte positions (for instance, -b1–72 would pass the first 72 bytes of each line). When -b and -n are used together, <i>list</i> is adjusted so that no multi-byte character is split. If -b is used, the input line should contain 1023 bytes or less.</p> <p>-c list The <i>list</i> following -c specifies character positions (for instance, -c1–72 would pass the first 72 characters of each line).</p> <p>-d delim The character following -d is the field delimiter (-f option only). Default is <i>tab</i>. Space or other characters with special meaning to the shell must be quoted. <i>delim</i> can be a multi-byte character.</p> <p>-f list The <i>list</i> following -f is a list of fields assumed to be separated in the file by a delimiter character (see -d); for instance, -f1,7 copies the first and seventh field only. Lines with no field delimiters will be passed through intact (useful for table subheadings), unless -s is specified. If -f is used, the input line should contain 1023 characters or less.</p> <p>-n Do not split characters. When -b list and -n are used together, <i>list</i> is adjusted so that no multi-byte character is split.</p> <p>-s Suppresses lines with no delimiter characters in case of -f option. Unless specified, lines with no delimiters will be passed through untouched.</p>

EXAMPLES

A mapping of user IDs to names follows:

```
example% cut -d: -f1,5 /etc/passwd
```

To set **name** to current login name follows:

```
example% name=`who am i` | cut -f1 -d `
```

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

grep(1), **paste(1)**, **environ(5)**

DIAGNOSTICS**ERROR: line too long**

A line can have no more than 1023 characters or fields, or there is no new-line character.

ERROR: bad list for c / f option

Missing **-c** or **-f** option or incorrectly specified *list*. No error occurs if a line has fewer fields than the *list* calls for.

ERROR: no fields

The *list* is empty.

ERROR: no delimiter

Missing *char* on **-d** option.

ERROR: no newline in the last

The input file ends without a newline

ERROR: invalid multibyte character

An error in the input data.

WARNING: cannot open <filename>

Either *filename* cannot be read or does not exist. If multiple filenames are present, processing continues.

NAME	date – print and set the date
SYNOPSIS	<pre>date [-u] [+ format] date [-a [-]sss.fff] date [-u] [[mdd] HHMM mddHHMM [cc] yy]</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>If no argument is given, or if the argument begins with +, the current date and time are printed. Otherwise, the current date is set (only by super-user).</p> <p>Specifications of native language translations of month and weekday names are supported. The month and weekday names used for a language are based on the locale specified by the environment variables LC_TIME and LANG (see environ(5)).</p> <p>The month and weekday names used for a language are taken from a file whose format is specified in strftime(4). This file also defines country-specific date and time formats such as %C, which specifies the default date format. The following form is the default for %C:</p> <pre>%a %b %e %T %Z %Y</pre> <p>for example,</p> <pre>Fri Dec 23 10:10:42 EST 1988</pre> <p>Field Descriptors (must be preceded by a %):</p> <ul style="list-style-type: none"> a abbreviated weekday name A full weekday name b abbreviated month name B full month name c locale's appropriate date and time representation C default date and time format d day of month – 01 to 31 D date as %m/%d/%y e day of month – 1 to 31 (single digits are preceded by a blank) h abbreviated month name (alias for %b) H hour – 00 to 23 I hour – 01 to 12 j day of year – 001 to 366 m month of year – 01 to 12 M minute – 00 to 59 n insert a new-line character p string containing ante-meridiem or post-meridiem indicator (by default, AM or PM) r time as %I:%M:%S %p R time as %H:%M S second – 00 to 61, allows for leap seconds t insert a tab character T time as %H:%M:%S

U week number of year (Sunday as the first day of the week) – 00 to 53
w day of week – Sunday = 0
W week number of year (Monday as the first day of the week) – 00 to 53
x Country-specific date format
X Country-specific time format
y year within century – 00 to 99
Y year as *ccyy* (4 digits)
Z timezone name

OPTIONS

-a [-]sss.fff Slowly adjust the time by *sss.fff* seconds (*fff* represents fractions of a second). This adjustment can be positive or negative. The system's clock will be sped up or slowed down until it has drifted by the number of seconds specified.

-u Display (or set) the date in Greenwich Mean Time (GMT—universal time), bypassing the normal conversion to (or from) local time.

mm is the month number

dd is the day number in the month

HH is the hour number (24 hour system)

MM is the minute number

cc is the century minus one

yy is the last 2 digits of the year number

The month, day, year, and century may be omitted; the current values are applied as defaults. For example:

date 10080045

sets the date to Oct 8, 12:45 AM. The current year is the default because no year is supplied. The system operates in GMT. **date** takes care of the conversion to and from local standard and daylight time. Only the super-user may change the date. After successfully setting the date and time, **date** displays the new date according to the default format. The **date** command uses **TZ** to determine the correct time zone information (see **environ(5)**).

+ *format* If the argument begins with +, the output of **date** is under the control of the user. Each Field Descriptor, described above, is preceded by % and is replaced in the output by its corresponding value. A single % is encoded by %%. All other characters are copied to the output without change. The string is always terminated with a new-line character. If the argument contains embedded blanks it must be quoted (see the **EXAMPLES** section).

EXAMPLES

The command

```
example% date '+DATE: %m/%d/%y%nTIME: %H:%M:%S'
```

generates as output:

```
DATE: 08/01/76  
TIME: 14:45:05
```

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

LC_TIME

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

SEE ALSO

strftime(4), **environ(5)**

DIAGNOSTICS

No permission You are not the super-user and you try to change the date.
bad conversion The date set is syntactically incorrect.

NOTES

If you attempt to set the current date to one of the dates that the standard and alternate time zones change (for example, the date that daylight time is starting or ending), and you attempt to set the time to a time in the interval between the end of standard time and the beginning of the alternate time (or the end of the alternate time and the beginning of standard time), the results are unpredictable.

NAME	dc – desk calculator
SYNOPSIS	dc [<i>filename</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	dc is an arbitrary precision arithmetic package. Ordinarily it operates on decimal integers, but one may specify an input base, output base, and a number of fractional digits to be maintained. bc is a preprocessor for dc that provides infix notation and a C-like syntax that implements functions. bc also provides reasonable control structures for programs. See bc (1). The overall structure of dc is a stacking (reverse Polish) calculator. If an argument is given, input is taken from that file until its end, then from the standard input.
USAGE	The following constructions are recognized: <i>number</i> The value of the number is pushed on the stack. A number is an unbroken string of the digits 0–9. It may be preceded by an underscore (<code>_</code>) to input a negative number. Numbers may contain decimal points. + - / * % ^ The top two values on the stack are added (+), subtracted (-), multiplied (*), divided (/), remaindered (%), or exponentiated (^). The two entries are popped off the stack; the result is pushed on the stack in their place. Any fractional part of an exponent is ignored. sx The top of the stack is popped and stored into a register named <i>x</i> , where <i>x</i> may be any character. If the s is capitalized, <i>x</i> is treated as a stack and the value is pushed on it. lx The value in register <i>x</i> is pushed on the stack. The register <i>x</i> is not altered. All registers start with zero value. If the l is capitalized, register <i>x</i> is treated as a stack and its top value is popped onto the main stack. d The top value on the stack is duplicated. p The top value on the stack is printed. The top value remains unchanged. P Interprets the top of the stack as an ASCII string, removes it, and prints it. f All values on the stack are printed. q Exits the program. If executing a string, the recursion level is popped by two. Q Exits the program. The top value on the stack is popped and the string execution level is popped by that value. x Treats the top element of the stack as a character string and executes it as a string of dc commands. X Replaces the number on the top of the stack with its scale factor. [...] Puts the bracketed ASCII string onto the top of the stack.

<x >x =x The top two elements of the stack are popped and compared. Register *x* is evaluated if they obey the stated relation.

v Replaces the top element on the stack by its square root. Any existing fractional part of the argument is taken into account, but otherwise the scale factor is ignored.

! Interprets the rest of the line as a shell command.

c All values on the stack are popped.

i The top value on the stack is popped and used as the number radix for further input. **I** Pushes the input base on the top of the stack.

o The top value on the stack is popped and used as the number radix for further output.

O Pushes the output base on the top of the stack.

k The top of the stack is popped, and that value is used as a non-negative scale factor: the appropriate number of places are printed on output, and maintained during multiplication, division, and exponentiation. The interaction of scale factor, input base, and output base will be reasonable if all are changed together.

z The stack level is pushed onto the stack.

Z Replaces the number on the top of the stack with its length.

? A line of input is taken from the input source (usually the terminal) and executed.

; are used by **bc(1)** for array operations.

EXAMPLES

This example prints the first ten values of *n!*:

```
[!a1+dsa*pla10>y]sy
0sa1
lyx
```

SEE ALSO

bc(1)

DIAGNOSTICS

***x* is unimplemented** *x* is an octal number.

stack empty Not enough elements on the stack to do what was asked.

Out of space The free list is exhausted (too many digits).

Out of headers Too many numbers being kept around.

Out of pushdown Too many items on the stack.

Nesting Depth Too many levels of nested execution.

NAME	deroff – remove nroff/troff, tbl, and eqn constructs
SYNOPSIS	deroff [-m [m s l] [-w] [-i] [<i>filename</i> . . .]
AVAILABILITY	SUNWdoc
DESCRIPTION	deroff reads each of the <i>filenames</i> in sequence and removes all troff (1) requests, macro calls, backslash constructs, eqn (1) constructs (between .EQ and .EN lines, and between delimiters), and tbl (1) descriptions, perhaps replacing them with white space (blanks and blank lines), and writes the remainder of the file on the standard output. deroff follows chains of included files (.so and .nx troff commands); if a file has already been included, a .so naming that file is ignored and a .nx naming that file terminates execution. If no input file is given, deroff reads the standard input.
OPTIONS	<p>-m The -m option may be followed by an m, s, or l. The -mm option causes the macros to be interpreted so that only running text is output (that is, no text from macro lines.) The -ml option forces the -mm option and also causes deletion of lists associated with the mm macros.</p> <p>-w If the -w option is given, the output is a word list, one “word” per line, with all other characters deleted. Otherwise, the output follows the original, with the deletions mentioned above. In text, a “word” is any string that <i>contains</i> at least two letters and is composed of letters, digits, ampersands (&), and apostrophes ('); in a macro call, however, a “word” is a string that <i>begins</i> with at least two letters and contains a total of at least three letters. Delimiters are any characters other than letters, digits, apostrophes, and ampersands. Trailing apostrophes and ampersands are removed from “words.”</p> <p>-i The -i option causes deroff to ignore .so and .nx commands.</p>
SEE ALSO	eqn (1), nroff (1), tbl (1), troff (1)
NOTES	deroff is not a complete troff interpreter, so it can be confused by subtle constructs. Most such errors result in too much rather than too little output. The -ml option does not handle nested lists correctly.

NAME	df – display status of disk space on file systems
SYNOPSIS	<code>/usr/ucb/df [-a] [-i] [-t type] [filesystem...] [filename...]</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>df displays the amount of disk space occupied by currently mounted file systems, the amount of used and available space, and how much of the file system's total capacity has been used.</p> <p>If arguments to df are path names, df produces a report on the file system containing the named file. Thus 'df .' shows the amount of space on the file system containing the current directory.</p>
OPTIONS	<p>-a Report on all filesystems including the uninteresting ones which have zero total blocks. (that is, auto- mounter)</p> <p>-i Report the number of used and free inodes. Print '*' if no information is available.</p> <p>-t type Report on filesystems of a given type (for example, nfs or ufs).</p>
EXAMPLES	<p>A sample of output for df looks like:</p> <pre> example% df Filesystem kbytes used avail capacity Mounted on sparky:/ 7445 4714 1986 70% / sparky:/usr 42277 35291 2758 93% /usr </pre> <p>Note: used+avail is less than the amount of space in the file system (kbytes); this is because the system reserves a fraction of the space in the file system to allow its file system allocation routines to work well. The amount reserved is typically about 10%; this may be adjusted using tunefs. When all the space on a file system except for this reserve is in use, only the super-user can allocate new files and data blocks to existing files. When a file system is overallocated in this way, df may report that the file system is more than 100% utilized.</p>
FILES	<p>/etc/mnttab list of file systems currently mounted</p> <p>/etc/vfstab list of default parameters for each file system</p>
SEE ALSO	du(1M) , quot(1M) , tunefs(1M) , mnttab(4)

NAME	diff – display line-by-line differences between pairs of text files
SYNOPSIS	<pre>diff [-bitw] [-c -e -f -h -n] filename1 filename2 diff [-bitw] [-C number] filename1 filename2 diff [-bitw] [-D string] filename1 filename2 diff [-bitw] [-c -e -f -h -n] [-l] [-r] [-s] [-S name] directory1 directory2</pre>
AVAILABILITY	SUNWdoc
DESCRIPTION	<p>diff tells what lines must be changed in two files to bring them into agreement. If <i>filename1</i> (<i>filename2</i>) is <code>-</code>, the standard input is used. If <i>filename1</i> (<i>filename2</i>) is a directory, then a file in that directory with the name <i>filename2</i> (<i>filename1</i>) is used. The normal output contains lines of these forms:</p> <pre> n1 a n3,n4 n1,n2 d n3 n1,n2 c n3,n4</pre> <p>These lines resemble ed commands to convert <i>filename1</i> into <i>filename2</i>. The numbers after the letters pertain to <i>filename2</i>. In fact, by exchanging a for d and reading backward one may ascertain equally how to convert <i>filename2</i> into <i>filename1</i>. As in ed, identical pairs, where $n1 = n2$ or $n3 = n4$, are abbreviated as a single number.</p> <p>Following each of these lines come all the lines that are affected in the first file flagged by <code><</code>, then all the lines that are affected in the second file flagged by <code>></code>.</p>
OPTIONS	<p>-b Ignores trailing blanks (spaces and tabs) and treats other strings of blanks as equivalent.</p> <p>-i Ignores the case of letters; for example, 'A' will compare equal to 'a'.</p> <p>-t Expands TAB characters in output lines. Normal or -c output adds character(s) to the front of each line that may adversely affect the indentation of the original source lines and make the output lines difficult to interpret. This option will preserve the original source's indentation.</p> <p>-w Ignores all blanks (SPACE and TAB characters) and treats all other strings of blanks as equivalent; for example, 'if (a == b)' will compare equal to 'if(a= b)'. The following options are mutually exclusive:</p> <p>-c Produces a listing of differences with three lines of context. With this option output format is modified slightly: output begins with identification of the files involved and their creation dates, then each change is separated by a line with a dozen *'s. The lines removed from <i>filename1</i> are marked with '-'; those added to <i>filename2</i> are marked '+'. Lines that are changed from one file to the other are marked in both files with '!'. The following options are mutually exclusive:</p>

- C number** Produces a listing of differences identical to that produced by **-c** with *number* lines of context.
- e** Produces a script of *a*, *c*, and *d* commands for the editor **ed**, which will recreate *filename2* from *filename1*. In connection with **-e**, the following shell program may help maintain multiple versions of a file. Only an ancestral file (\$1) and a chain of version-to-version **ed** scripts (\$2,\$3,...) made by **diff** need be on hand. A "latest version" appears on the standard output.
- (shift; cat \$*; echo '1,\$p') | ed - \$1**
- Except in rare circumstances, **diff** finds a smallest sufficient set of file differences.
- f** Produces a similar script, not useful with **ed**, in the opposite order.
- h** Does a fast, half-hearted job. It works only when changed stretches are short and well separated, but does work on files of unlimited length. Options **-e** and **-f** are unavailable with **-h**.
- n** Produces a script similar to **-e**, but in the opposite order and with a count of changed lines on each insert or delete command.
- D string** Creates a merged version of *filename1* and *filename2* with C preprocessor controls included so that a compilation of the result without defining *string* is equivalent to compiling *filename1*, while defining *string* will yield *filename2*.
- The following options are used for comparing directories:
- l** Produce output in long format. Before the **diff**, each text file is piped through **pr(1)** to paginate it. Other differences are remembered and summarized after all text file differences are reported.
- r** Applies **diff** recursively to common subdirectories encountered.
- s** Reports files that are the identical; these would not otherwise be mentioned.
- S name** Starts a directory **diff** in the middle, beginning with the file *name*.

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

FILES

/tmp/d?????
/usr/lib/diffh for **-h**
/usr/bin/pr

SEE ALSO

bdiff(1), **cmp(1)**, **comm(1)**, **ed(1)**, **pr(1)**, **sdiff(1)**, **environ(5)**

DIAGNOSTICS

Exit status is **0** for no differences, **1** for some differences, **2** for trouble.

NOTES

Editing scripts produced under the **-e** or **-f** option are naive about creating lines consisting of a single period (.).

Missing newline at end of file indicates that the last line of the file in question did not have a new-line. If the lines are different, they will be flagged and output; although the output will seem to indicate they are the same.

NAME	diff3 – 3-way differential file comparison
SYNOPSIS	diff3 [-exEX3] <i>filename1 filename2 filename3</i>
AVAILABILITY	SUNWesu
DESCRIPTION	<p>diff3 compares three versions of a file, and publishes disagreeing ranges of text flagged with these codes:</p> <pre> ==== all three files differ ====1 <i>filename1</i> is different ====2 <i>filename2</i> is different ====3 <i>filename3</i> is different </pre> <p>The type of change suffered in converting a given range of a given file to some other is indicated in one of these ways:</p> <pre> <i>f</i>: <i>n1</i> a Text is to be appended after line number <i>n1</i> in file <i>f</i>, where <i>f</i> = 1, 2, or 3. <i>f</i>: <i>n1</i>, <i>n2</i> c Text is to be changed in the range line <i>n1</i> to line <i>n2</i>. If <i>n1</i> = <i>n2</i>, the range may be abbreviated to <i>n1</i>. </pre> <p>The original contents of the range follows immediately after a c indication. When the contents of two files are identical, the contents of the lower-numbered file is suppressed. The following command will apply the resulting script to <i>filename1</i>.</p> <pre>(cat script; echo '1,\$p') ed - <i>filename1</i></pre>
OPTIONS	<p>-e Produce a script for the editor ed(1) that will incorporate into <i>filename1</i> all changes between <i>filename2</i> and <i>filename3</i>, i.e., the changes that normally would be flagged ==== and ====3.</p> <p>-x Produce a script to incorporate only changes flagged ====.</p> <p>-3 Produce a script to incorporate only changes flagged ====3.</p> <p>-E Produce a script that will incorporate all changes between <i>filename2</i> and <i>filename3</i>, but treat overlapping changes (that is, changes that would be flagged with ==== in the normal listing) differently. The overlapping lines from both files will be inserted by the edit script, bracketed by <<<<<< and >>>>>> lines.</p> <p>-X Produce a script that will incorporate only changes flagged ====, but treat these changes in the manner of the -E option.</p>
FILES	<p>/tmp/d3* /usr/lib/diff3prog</p>

SEE ALSO `diff(1)`

NOTES Text lines that consist of a single '.' will defeat `-e`.
Files longer than 64 Kbytes will not work.

NAME	diffmk – mark differences between versions of a troff input file
SYNOPSIS	diffmk <i>oldfile newfile markedfile</i>
AVAILABILITY	SUNWdoc
DESCRIPTION	diffmk compares two versions of a file and creates a third version that includes “change mark” (.mc) commands for nroff(1) and troff(1) . <i>oldfile</i> and <i>newfile</i> are the old and new versions of the file. diffmk generates <i>markedfile</i> , which, contains the text from <i>newfile</i> with troff(1) “change mark” requests (.mc) inserted where <i>newfile</i> differs from <i>oldfile</i> . When <i>markedfile</i> is formatted, changed or inserted text is shown by at the right margin of each line. The position of deleted text is shown by a single *.
EXAMPLES	<p>diffmk can also be used in conjunction with the proper troff requests to produce program listings with marked changes. In the following command line:</p> <pre>example% diffmk old.c new.c marked.c ; nroff reqs marked.c pr</pre> <p>the file reqs contains the following troff requests:</p> <pre>.pl 1 .ll 77 .nf .eo .nh</pre> <p>which eliminate page breaks, adjust the line length, set no-fill mode, ignore escape characters, and turn off hyphenation, respectively.</p> <p>If the characters and * are inappropriate, you might run <i>markedfile</i> through sed(1) to globally change them.</p>
SEE ALSO	diff(1) , nroff(1) , sed(1) , troff(1)
BUGS	Aesthetic considerations may dictate manual adjustment of some output. File differences involving only formatting requests may produce undesirable output, that is, replacing .sp 2 will produce a “change mark” on the preceding or following line of output.

NAME	dircmp – directory comparison
SYNOPSIS	dircmp [-d] [-s] [-wn] <i>dir1 dir2</i>
AVAILABILITY	SUNWesu
DESCRIPTION	dircmp examines <i>dir1</i> and <i>dir2</i> and generates various tabulated information about the contents of the directories. Listings of files that are unique to each directory are generated for all the options. If no option is entered, a list is output indicating whether the file names common to both directories have the same contents.
OPTIONS	-d Compare the contents of files with the same name in both directories and output a list telling what must be changed in the two files to bring them into agreement. The list format is described in diff(1) . -s Suppress messages about identical files. -wn Change the width of the output line to <i>n</i> characters. The default width is 72.
SEE ALSO	cmp(1) , diff(1)

NAME	dis – object code disassembler
SYNOPSIS	dis [-C] [[-o] [-V] [-d sec] [-D sec] [-F function] [-I string] [-t sec] <i>filename</i> ...
DESCRIPTION	The dis command produces an assembly language listing of <i>file</i> , which may be an object file or an archive of object files. The listing includes assembly statements and an octal or hexadecimal representation of the binary that produced those statements.
OPTIONS	<p>The following <i>options</i> are interpreted by the disassembler and may be specified in any order.</p> <p>-C Display demangled C++ symbol names in the disassembly.</p> <p>-o Print numbers in octal. The default is hexadecimal.</p> <p>-V Print, on standard error, the version number of the disassembler being executed.</p> <p>-d sec Disassemble the named section as data, printing the offset of the data from the beginning of the section.</p> <p>-D sec Disassemble the named section as data, printing the actual address of the data.</p> <p>-F function Disassemble only the named function in each object file specified on the command line. The -F option may be specified multiple times on the command line.</p> <p>-I string Disassemble the archive file specified by <i>string</i>. For example, one would issue the command dis -I x -I z to disassemble libx.a and libz.a, which are assumed to be in LIBDIR.</p> <p>-t sec Disassemble the named section as text.</p> <p>If the -d, -D or -t options are specified, only those named sections from each user-supplied file name will be disassembled. Otherwise, all sections containing text will be disassembled.</p> <p>On output, a number enclosed in brackets at the beginning of a line, such as [5], indicates that the break-pointable line number starts with the following instruction. These line numbers will be printed only if the file was compiled with additional debugging information, for example, the -g option of cc(1B). An expression such as <40> in the operand field or in the symbolic disassembly, following a relative displacement for control transfer instructions, is the computed address within the section to which control will be transferred. A function name will appear in the first column, followed by () if the object file contains a symbol table.</p>
ENVIRONMENT	<p>LIBDIR If this environment variable contains a value, use this as the path to search for the library. If the variable contains a null value, or is not set, it defaults to searching for the library under /usr/ccs/lib.</p>

FILES /usr/ccs/lib default **LIBDIR**

SEE ALSO **as(1)**, **cc(1B)**, **ld(1)**, **a.out(4)**

DIAGNOSTICS The self-explanatory diagnostics indicate errors in the command line or problems encountered with the specified files.

NAME	dispgid – displays a list of all valid group names
SYNOPSIS	dispgid
AVAILABILITY	SUNWcsu
DESCRIPTION	dispgid displays a list of all group names on the system (one group per line).
EXIT CODES	0 Successful execution 1 Cannot read the group file

NAME	dispuid – displays a list of all valid user names
SYNOPSIS	dispuid
AVAILABILITY	SUNWcsu
DESCRIPTION	dispuid displays a list of all user names on the system (one line per name).
EXIT CODES	0 Successful execution 1 Cannot read the password file

NAME	dos2unix – convert text file from DOS format to ISO format
SYNOPSIS	dos2unix [-ascii] [-iso] [-7] <i>originalfile convertedfile</i>
AVAILABILITY	SUNWesu
DESCRIPTION	<p>dos2unix converts characters in the DOS extended character set to the corresponding ISO standard characters.</p> <p>This command can be invoked from either DOS or SunOS. However, the filenames must conform to the conventions of the environment in which the command is invoked.</p> <p>If the original file and the converted file are the same, dos2unix will rewrite the original file after converting it.</p>
OPTIONS	<p>-ascii Removes extra carriage returns and converts end of file characters in DOS format text files to conform to SunOS requirements.</p> <p>-iso This is the default. It converts characters in the DOS extended character set to the corresponding ISO standard characters.</p> <p>-7 Convert 8 bit DOS graphics characters to 7 bit space characters so that SunOS can read the file.</p>
SEE ALSO	unix2dos(1)
DIAGNOSTICS	<p>File <i>filename</i> not found, or no read permission The input file you specified does not exist, or you do not have read permission (check with the SunOS ls -l command).</p> <p>Bad output filename <i>filename</i>, or no write permission The output file you specified is either invalid, or you do not have write permission for that file or the directory that contains it. Check also that the drive or diskette is not write-protected.</p> <p>Error while writing to temporary file An error occurred while converting your file, possibly because there is not enough space on the current drive. Check the amount of space on the current drive using the DIR command. Also be certain that the default diskette or drive is write-enabled (not write-protected). Note that when this error occurs, the original file remains intact.</p> <p>Could not rename temporary file to Translated temporary file name = <i>filename</i>. The program could not perform the final step in converting your file. Your converted file is stored under the name indicated on the second line of this message.</p>

NAME	download – host resident PostScript font downloader
SYNOPSIS	download [-f] [-p <i>printer</i>] [-m <i>name</i>] [-H <i>directory</i>] [<i>file...</i>] /usr/lib/lp/postscript/download
DESCRIPTION	<p>download prepends host resident fonts to <i>files</i> and writes the results on the standard output. If no <i>files</i> are specified, or if – is one of the input <i>files</i>, the standard input is read. <i>download</i> assumes the input <i>files</i> make up a single PostScript job and that requested fonts can be included at the start of each input <i>file</i>.</p> <p>Requested fonts are named in a comment (marked with %%DocumentFonts:) in the input <i>files</i>. Available fonts are the ones listed in the map table selected using the –m option.</p> <p>The map table consists of fontname–file pairs. The fontname is the full name of the PostScript font, exactly as it would appear in a %%DocumentFonts: comment. The file is the pathname of the host resident font. A file that begins with a / is used as is. Otherwise the pathname is relative to the host font directory. Comments are introduced by % (as in PostScript) and extend to the end of the line.</p> <p>The only candidates for downloading are fonts listed in the map table that point download to readable files. A font is downloaded once, at most. Requests for unlisted fonts or inaccessible files are ignored. All requests are ignored if the map table can not be read.</p>
OPTIONS	<p>–f Force a complete scan of each input <i>file</i>. In the absence of an explicit comment pointing <i>download</i> to the end of the file, the default scan stops immediately after the PostScript header comments.</p> <p>–p <i>printer</i> Check the list of printer-resident fonts in /etc/lp/printers/<i>printer/residentfonts</i> before downloading.</p> <p>–m <i>name</i> Use <i>name</i> as the font map table. A <i>name</i> that begins with / is the full pathname of the map table and is used as is. Otherwise <i>name</i> is appended to the pathname of the host font directory.</p> <p>–H <i>directory</i> Use <i>dir</i> as the host font directory. The default is /usr/lib/lp/postscript.</p>
EXAMPLES	<p>The following map table could be used to control the downloading of the Bookman font family:</p> <pre> % % The first string is the full PostScript font name. The second string % is the file name - relative to the host font directory unless it begins % with a /. % Bookman-Light bookman/light Bookman-LightItalic bookman/lightitalic Bookman-Demi bookman/demi </pre>

Bookman-DemiItalic **bookman/demiitalic**

Using the file **myprinter/map** (in the default host font directory) as the map table, you could download fonts by issuing the following command:

example% download -m myprinter/map file

SEE ALSO

dpost(1), postdaisy(1), postdmd(1), postio(1), postmd(1), postprint(1), posttek(1)

DIAGNOSTICS

An exit status of **0** is returned if *files* were successfully processed.

NOTES

The **download** program should be part of a more general program.

download does not look for %%PageFonts: comments and there is no way to force multiple downloads of a particular font.

We do not recommend the use of full pathnames in either map tables or the names of map tables.

NAME	dpost – troff postprocessor for PostScript printers
SYNOPSIS	dpost [-c <i>num</i>] [-e <i>num</i>] [-m <i>num</i>] [-n <i>num</i>] [-o <i>list</i>] [-w <i>num</i>] [-x <i>num</i>] [-y <i>num</i>] [-F <i>dir</i>] [-H <i>dir</i>] [-L <i>file</i>] [-O] [-T <i>name</i>] [<i>file</i> ...] /usr/lib/lp/postscript/dpost
DESCRIPTION	dpost translates <i>files</i> created by troff (1) into PostScript and writes the results on the standard output. If no <i>files</i> are specified, or if – is one of the input <i>files</i> , the standard input is read. The <i>files</i> should be prepared by troff . The default font files in /usr/lib/font/devpost produce the best and most efficient output. They assume a resolution of 720 dpi, and can be used to format files by adding the -Tpost option to the troff call. Older versions of the eqn and pic preprocessors need to know the resolution that troff will be using to format the <i>files</i> . If those are the versions installed on your system, use the -r720 option with eqn and -T720 with pic . dpost makes no assumptions about resolutions. The first x res command sets the resolution used to translate the input <i>files</i> , the DESC.out file, usually /usr/lib/font/devpost/DESC.out , defines the resolution used in the binary font files, and the PostScript prologue is responsible for setting up an appropriate user coordinate system.
OPTIONS	<p>-c <i>num</i> Print <i>num</i> copies of each page. By default only one copy is printed.</p> <p>-e <i>num</i> Sets the text encoding level to <i>num</i>. The recognized choices are 0, 1, and 2. The size of the output file and print time should decrease as <i>num</i> increases. Level 2 encoding will typically be about 20 percent faster than level 0, which is the default and produces output essentially identical to previous versions of dpost.</p> <p>-m <i>num</i> Magnify each logical page by the factor <i>num</i>. Pages are scaled uniformly about the origin, which is located near the upper left corner of each page. The default magnification is 1.0.</p> <p>-n <i>num</i> Print <i>num</i> logical pages on each piece of paper, where <i>num</i> can be any positive integer. By default, <i>num</i> is set to 1.</p> <p>-o <i>list</i> Print those pages for which numbers are given in the comma-separated <i>list</i>. The list contains single numbers <i>N</i> and ranges <i>N1–N2</i>. A missing <i>N1</i> means the lowest numbered page, a missing <i>N2</i> means the highest. The page range is an expression of logical pages rather than physical sheets of paper. For example, if you are printing two logical pages to a sheet, and you specified a range of 4, then two sheets of paper would print, containing four page layouts. If you specified a page range of 3-4, when requesting two logical pages to a sheet; then <i>only</i> page 3 and page 4 layouts would print, and they would appear on one physical sheet of paper.</p>

- p mode** Print *files* in either portrait or landscape *mode*. Only the first character of *mode* is significant. The default *mode* is portrait.
- w num** Set the line width used to implement *troff* graphics commands to *num* points, where a point is approximately 1/72 of an inch. By default, *num* is set to **0.3** points.
- x num** Translate the origin *num* inches along the positive x axis. The default coordinate system has the origin fixed near the upper left corner of the page, with positive x to the right and positive y down the page. Positive *num* moves everything right. The default offset is **0** inches.
- y num** Translate the origin *num* inches along the positive y axis. Positive *num* moves text up the page. The default offset is **0**.
- F dir** Use *dir* as the font directory. The default *dir* is **/usr/lib/font**, and **dpost** reads binary font files from directory **/usr/lib/font/devpost**.
- H dir** Use *dir* as the host resident font directory. Files in this directory should be complete PostScript font descriptions, and must be assigned a name that corresponds to the appropriate two-character **troff** font name. Each font file is copied to the output file only when needed and at most once during each job. There is no default directory.
- L file** Use *file* as the PostScript prologue which, by default, is **/usr/lib/lp/postscript/dpost.ps**.
- O** Disables PostScript picture inclusion. A recommended option when **dpost** is run by a spooler in a networked environment.
- T name** Use font files for device *name* as the best description of available PostScript fonts. By default, *name* is set to **post** and **dpost** reads binary files from **/usr/lib/font/devpost**.

EXAMPLES

If the old versions of **eqn** and **pic** are installed on your system, you can obtain the best possible looking output by issuing a command line such as the following:

```
example% pic -T720 file | tbl | eqn -r720 | troff -mm -Tpost | dpost
```

Otherwise,

```
example% pic file | tbl | eqn | troff -mm -Tpost | dpost
```

should give the best results.

FILES

```
/usr/lib/font/devpost/*.out
/usr/lib/font/devpost/charlib/*
/usr/lib/lp/postscript/color.ps
/usr/lib/lp/postscript/draw.ps
/usr/lib/lp/postscript/forms.ps
/usr/lib/lp/postscript/ps.requests
/usr/lib/macros/pictures
/usr/lib/macros/color
```

SEE ALSO **download(1), postdaisy(1), postdmd(1), postio(1), postmd(1), postprint(1), postreverse(1), posttek(1), troff(1)**

DIAGNOSTICS An exit status of **0** is returned if *files* have been translated successfully, while 2 often indicates a syntax error in the input *files*.

NOTES Output files often do not conform to Adobe's file structuring conventions. Piping the output of **dpost** through **postreverse(1)** should produce a minimally conforming PostScript file.

Although **dpost** can handle files formatted for any device, emulation is expensive and can easily double the print time and the size of the output file. No attempt has been made to implement the character sets or fonts available on all devices supported by **troff**. Missing characters will be replaced by white space, and unrecognized fonts will usually default to one of the Times fonts (that is, **R**, **I**, **B**, or **BI**).

An **x res** command must precede the first **x init** command, and all the input *files* should have been prepared for the same output device.

Use of the **-T** option is not encouraged. Its only purpose is to enable the use of other PostScript font and device description files, that perhaps use different resolutions, character sets, or fonts.

Although level 0 encoding is the only scheme that has been thoroughly tested, level 2 is fast and may be worth a try.

NAME	du – display the number of disk blocks used per directory or file
SYNOPSIS	<code>/usr/ucb/du</code> <code>/usr/ucb/du [-a] [-s] [<i>filename</i>]</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	du gives the number of kilobytes contained in all files and, recursively, directories within each specified directory or file <i>filename</i> . If <i>filename</i> is missing, '.' (the current directory) is used. A file which has multiple links to it is only counted once.
OPTIONS	-a Generate an entry for each file. -s Only display the grand total for each of the specified <i>filenames</i> . Entries are generated only for each directory in the absence of options.
EXAMPLES	Here is an example of using du in a directory. We used the pwd(1) command to identify the directory, then used du to show the usage of all the subdirectories in that directory. The grand total for the directory is the last entry in the display: <pre>example% pwd /usr/ralph/misc example% du 5 ./jokes 33 ./squash 44 ./tech.papers/lpr.document 217 ./tech.papers/new.manager 401 ./tech.papers 144 ./memos 80 ./letters 388 ./window 93 ./messages 15 ./useful.news 1211 . example%</pre>
ENVIRONMENT	If any of the LC_* variables (LC_CTYPE , LC_MESSAGES , LC_TIME , LC_COLLATE , LC_NUMERIC , and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of du for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how du behaves.

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

pwd(1), **df(1M)**, **quot(1M)**, **environ(5)**

NOTES

Filename arguments that are not directory names are ignored, unless you use **-a**.
If there are too many distinct linked files, **du** will count the excess files more than once.

NAME	dump – dump selected parts of an object file
SYNOPSIS	dump [-aCcDfghLlorstV] [-T <i>index</i> [, <i>indexn</i>]] <i>filename</i> ... dump [-afhorstL [<i>v</i>]] <i>filename</i> ... dump [-hsr [-d <i>number</i> [, <i>numbern</i>]]] <i>filename</i> ... dump [-hsrt [-n <i>name</i>]] <i>filename</i> ...
DESCRIPTION	The dump command dumps selected parts of each of its object <i>file</i> arguments.
OPTIONS	<p>This command will accept both object files and archives of object files. It processes each file argument according to one or more of the following options:</p> <ul style="list-style-type: none"> -a Dump the archive header of each member of an archive. -c Dump the string table(s). -C Dump decoded C++ symbol table names. -D Dump debugging information. -f Dump each file header. -g Dump the global symbols in the symbol table of an archive. -h Dump the section headers. -l Dump line number information. -L Dump dynamic linking information and static shared library information, if available. -o Dump each program execution header. -r Dump relocation information. -s Dump section contents in hexadecimal. -t Dump symbol table entries. -T <i>index</i> or -T <i>index1,index2</i> Dump only the indexed symbol table entry defined by <i>index</i> or a range of entries defined by <i>index1,index2</i>. -V Print version information. <p>The following modifiers are used in conjunction with the options listed above to modify their capabilities.</p> <ul style="list-style-type: none"> -d <i>number</i> or -d <i>number1,number2</i> Dump the section number indicated by <i>number</i> or the range of sections starting at <i>number1</i> and ending at <i>number2</i>. This modifier can be used with -h, -s, and -r. When -d is used with -h or -s, the argument is treated as the number of a section or range of sections. When -d is used with -r, the argument is treated as the number of the section or range of sections to which the relocation applies. For example, to print out all relocation entries associated with the .text section, specify the number of the section as the argument to -d. If .text is section number 2 in the file, dump -r -d 2 will print all

associated entries. To print out a specific relocation section use **dump -s -n *name*** for raw data output, or **dump -sv -n *name*** for interpreted output.

- n *name*** Dump information pertaining only to the named entity. This modifier can be used with **-h**, **-s**, **-r**, and **-t**. When **-n** is used with **-h** or **-s**, the argument will be treated as the name of a section. When **-n** is used with **-t** or **-r**, the argument will be treated as the name of a symbol. For example, **dump -t -n .text** will dump the symbol table entry associated with the symbol whose name is **dump -h -n .text** will dump the section header information for the **section**.
- p** Suppress printing of the headings.
- v** Dump information in symbolic representation rather than numeric. This modifier can be used with **-a** (date, user id, group id), **-f** (class, data, type, machine, version, flags), **-h** (type, flags), **-o** (type, flags), **-r** (name, type), **-s** (interpret section contents wherever possible), **-t** (type, bind), and **-L** (value). When **-v** is used with **-s**, all sections that can be interpreted, such as the string table or symbol table, will be interpreted. For example, **dump -sv -n .symtab filename...** will produce the same formatted output as **dump -tv filename...**, but **dump -s -n .symtab filename...** will print raw data in hexadecimal. Without additional modifiers, **dump -sv filename...** will dump all sections in the files interpreting all those that it can and dumping the rest (such as **.text** or **.data**) as raw data.

The **dump** command attempts to format the information it dumps in a meaningful way, printing certain information in character, hexadecimal, octal or decimal representation as appropriate.

SEE ALSO

nm(1), a.out(4), ar(4)

NAME	dumpcs – show codeset table for the current locale
SYNOPSIS	dumpcs [-0123vw]
DESCRIPTION	<p>dumpcs shows a list of printable characters for the user's current locale, along with their hexadecimal code values. The display device is assumed to be capable of displaying characters for a given locale. With no option, dumpcs displays the entire list of printable characters for the current locale.</p> <p>With one or more numeric options specified, it shows EUC codeset(s) for the current locale according to the numbers specified, and in order of codeset number. Each non-printable character is represented by an asterisk "*" and enough ASCII space character(s) to fill that codeset's column width.</p>
OPTIONS	<p>-0 Show ASCII (or EUC primary) codeset.</p> <p>-1 Show EUC codeset 1, if used for the current locale.</p> <p>-2 Show EUC codeset 2, if used for the current locale.</p> <p>-3 Show EUC codeset 3, if used for the current locale.</p> <p>-v Verbose. Normally, ranges of non-printable characters are collapsed into a single line. This option produces one line for each non-printable character.</p> <p>-w Replace code values with corresponding wide character values (process codes).</p>
ENVIRONMENT	The environment variables LC_CTYPE and LANG control the character classification throughout dumpcs . On entry to dumpcs , these environment variables are checked in that order. This implies that a new setting for LANG does not override the setting of LC_CTYPE . When none of the values is valid, the character classification defaults to the POSIX.1 "C" locale.
FILES	/usr/lib/locale/locale-name/LC_CTYPE/ctype data file containing character classification, conversion, and character set width information
SEE ALSO	chrtbl(1M)

NAME	echo – echo arguments
SYNOPSIS	echo [<i>argument</i>] ... echo [-n] [<i>argument</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>echo writes its arguments separated by blanks and terminated by a new-line on the standard output.</p> <p>/usr/bin/sh (Bourne shell) understands the following C-like escape conventions, however, beware of conflicts with the shell's use of \ :</p> <p>Note that the escape conventions are not available if the BSD version of the Bourne shell echo command is used. (See OPTIONS section)</p> <ul style="list-style-type: none"> \ b backspace \ c print line without new-line \ f form-feed \ n new-line \ r carriage return \ t tab \ v vertical tab \\ backslash \ 0n where <i>n</i> is the 8-bit character whose ASCII code is the 1-, 2- or 3-digit octal number representing that character. <p>echo is useful for producing diagnostics in command files, for sending known data into a pipe, and for displaying the contents of environment variables.</p>
OPTIONS	<p>The -n option is available to /usr/bin/sh users only if /usr/ucb precedes /usr/bin in the user's PATH, however, then the backslash characters, described above, are NOT available. -n is available to /usr/bin/csh users, regardless of PATH:</p> <p>-n Do not add the newline to the output.</p>
SEE ALSO	sh(1) , shell_builtins(1) , ascii(5)
NOTES	<p>The -n option is a transition aid for BSD applications, and may not be supported in future releases.</p> <p>When representing an 8-bit character by using the escape convention \ 0n, the <i>n</i> must always be preceded by the digit zero (0).</p> <p>For example, typing: echo `WARNING:\07` will print the phrase WARNING: and sound the "bell" on your terminal. The use of single (or double) quotes (or two backslashes) is required to protect the "`" that precedes the "07".</p> <p>Following the \ 0, up to three digits are used in constructing the octal output character. If, following the \ 0n, you want to echo additional digits that are not part of the octal representation, you must use the full 3-digit <i>n</i>. For example, if you want to echo "ESC 7"</p>

you must use the three digits "033" rather than just the two digits "33" after the \0.

2 digits	Incorrect:	echo "\0337" od -xc	
	produces:	df0a	(hex)
		337	(ascii)

3 digits	Correct:	echo "\00337" od -xc	
	produces:	1b37 0a00	(hex)
		033 7	(ascii)

For the octal equivalents of each character, see **ascii(5)**.

NAME	echo – echo arguments to standard output
SYNOPSIS	<code>/usr/ucb/echo [<i>argument</i>] ...</code> <code>/usr/ucb/echo [-n] [<i>argument</i>]</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>echo writes its arguments separated by blanks and terminated by a new-line on the standard output.</p> <p>The /usr/bin/sh version understands the following C-like escape conventions; beware of conflicts with the shell's use of <code>\</code>:</p> <ul style="list-style-type: none"><code>\b</code> backspace<code>\c</code> print line without new-line<code>\f</code> form-feed<code>\n</code> new-line<code>\r</code> carriage return<code>\t</code> tab<code>\v</code> vertical tab<code>\\</code> backslash<code>\0n</code> where <i>n</i> is the 8-bit character whose ASCII code is the 1-, 2- or 3-digit octal number representing that character. <p>echo is useful for producing diagnostics in command files and for sending known data into a pipe.</p>
OPTIONS	<code>-n</code> Do not add the newline to the output.
SEE ALSO	<code>sh(1)</code>
NOTES	<p>The <code>-n</code> option is a transition aid for BSD applications, and may not be supported in future releases.</p> <p>When representing an 8-bit character by using the escape convention <code>\0n</code>, the <i>n</i> must always be preceded by the digit zero (0).</p> <p>For example, typing: <code>echo 'WARNING:\07'</code> will print the phrase WARNING: and sound the "bell" on your terminal. The use of single (or double) quotes (or two backslashes) is required to protect the <code>"\"</code> that precedes the <code>"07"</code>.</p> <p>For the octal equivalents of each character, see <code>ascii(5)</code>.</p>

NAME	echo – put string on virtual output
SYNOPSIS	echo [<i>string</i> ...]
DESCRIPTION	The echo function directs each string it is passed to the standard output. If no argument is given, echo looks to the standard input for input. It is often used in conditional execution or for passing a string to another command.
EXAMPLES	Set the done descriptor to help if a test fails: <pre>done=`if [-s \$F1]; then echo close; else echo help; fi`</pre>
SEE ALSO	echo(1)

NAME	ed, red – text editor
SYNOPSIS	ed [- -s] [-p <i>string</i>] [-x] [-C] [<i>filename</i>] red [- -s] [-p <i>string</i>] [-x] [-C] [<i>filename</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>ed is the standard text editor. If the <i>filename</i> argument is given, ed simulates an e command (see below) on the named file; that is to say, the file is read into ed's buffer so that it can be edited.</p> <p>ed operates on a copy of the file it is editing; changes made to the copy have no effect on the file until a w (write) command is given. The copy of the text being edited resides in a temporary file called the <i>buffer</i>. There is only one buffer.</p> <p>red is a restricted version of ed. It will only allow editing of files in the current directory. It prohibits executing shell commands via <i>!shell command</i>. Attempts to bypass these restrictions result in an error message (<i>restricted shell</i>).</p> <p>Both ed and red support the fspec(4) formatting capability. After including a format specification as the first line of <i>filename</i> and invoking ed with your terminal in stty –tabs or stty tab3 mode (see stty(1)), the specified tab stops will automatically be used when scanning <i>filename</i>. For example, if the first line of a file contained:</p> <p style="padding-left: 40px;"><:t5,10,15 s72:></p> <p>tab stops would be set at columns 5, 10, and 15, and a maximum line length of 72 would be imposed. Note: when you are entering text into the file, this format is not in effect; instead, because of being in stty –tabs or stty tab3 mode, tabs are expanded to every eighth column.</p> <p>Commands to ed have a simple and regular structure: zero, one, or two <i>addresses</i> followed by a single-character <i>command</i>, possibly followed by parameters to that command. These addresses specify one or more lines in the buffer. Every command that requires addresses has default addresses, so that the addresses can very often be omitted.</p> <p>In general, only one command may appear on a line. Certain commands allow the input of text. This text is placed in the appropriate place in the buffer. While ed is accepting text, it is said to be in <i>input mode</i>. In this mode, <i>no</i> commands are recognized; all input is merely collected. Leave input mode by typing a period (.) at the beginning of a line, followed immediately by a carriage return.</p> <p>ed supports a limited form of <i>regular expression</i> notation; regular expressions are used in addresses to specify lines and in some commands (for example, s) to specify portions of a line that are to be substituted. A regular expression (RE) specifies a set of character strings. A member of this set of strings is said to be <i>matched</i> by the RE. The REs allowed by ed are constructed as follows:</p> <p>The following <i>one-character REs</i> match a <i>single</i> character:</p> <ol style="list-style-type: none"> 1.1 An ordinary character (<i>not</i> one of those discussed in 1.2 below) is a one-character RE that matches itself.

- 1.2 A backslash (\) followed by any special character is a one-character RE that matches the special character itself. The special characters are:
- ., *, [, and \ (period, asterisk, left square bracket, and backslash, respectively), which are always special, *except* when they appear within square brackets ([]); see 1.4 below).
 - ^ (caret or circumflex), which is special at the *beginning* of an *entire* RE (see 4.1 and 4.3 below), or when it immediately follows the left of a pair of square brackets ([]) (see 1.4 below).
 - \$ (dollar sign), which is special at the **end** of an *entire* RE (see 4.2 below).
 - The character used to bound (that is, delimit) an entire RE, which is special for that RE (for example, see how slash (/) is used in the **g** command, below.)
- 1.3 A period (.) is a one-character RE that matches any character except new-line.
- 1.4 A non-empty string of characters enclosed in square brackets ([]) is a one-character RE that matches *any one* character in that string. If, however, the first character of the string is a circumflex (^), the one-character RE matches any character *except* new-line and the remaining characters in the string. The ^ has this special meaning *only* if it occurs first in the string. The minus (-) may be used to indicate a range of consecutive characters; for example, [0-9] is equivalent to [0123456789]. The - loses this special meaning if it occurs first (after an initial ^, if any) or last in the string. The right square bracket (]) does not terminate such a string when it is the first character within it (after an initial ^, if any); for example, []a-f] matches either a right square bracket (]) or one of the ASCII letters **a** through **f** inclusive. The four characters listed in 1.2.a above stand for themselves within such a string of characters.

The following rules may be used to construct RE s from one-character REs:

- 2.1 A one-character RE is a RE that matches whatever the one-character RE matches.
- 2.2 A one-character RE followed by an asterisk (*) is a RE that matches *zero* or more occurrences of the one-character RE. If there is any choice, the longest leftmost string that permits a match is chosen.
- 2.3 A one-character RE followed by \{m\}, \{m,\}, or \{m,n\} is a RE that matches a *range* of occurrences of the one-character RE. The values of *m* and *n* must be non-negative integers less than 256; \{m\} matches *exactly* *m* occurrences; \{m,\} matches *at least* *m* occurrences; \{m,n\} matches *any number* of occurrences *between* *m* and *n* inclusive. Whenever a choice exists, the RE matches as many occurrences as possible.
- 2.4 The concatenation of REs is a RE that matches the concatenation of the strings matched by each component of the RE.
- 2.5 A RE enclosed between the character sequences \ (and \) is a RE that matches whatever the unadorned RE matches.

2.6 The expression `\n` matches the same string of characters as was matched by an expression enclosed between `\(` and `\)` *earlier* in the same RE. Here *n* is a digit; the sub-expression specified is that beginning with the *n*-th occurrence of `\(` (counting from the left). For example, the expression `^\(.*\)1$` matches a line consisting of two repeated appearances of the same string.

A RE may be constrained to match words.

3.1 `\<` constrains a RE to match the beginning of a string or to follow a character that is not a digit, underscore, or letter. The first character matching the RE must be a digit, underscore, or letter.

3.2 `\>` constrains a RE to match the end of a string or to precede a character that is not a digit, underscore, or letter.

An *entire RE* may be constrained to match only an initial segment or final segment of a line (or both).

4.1 A circumflex (`^`) at the beginning of an entire RE constrains that RE to match an *initial* segment of a line.

4.2 A dollar sign (`$`) at the end of an entire RE constrains that RE to match a *final* segment of a line.

4.3 The construction `^entire RE$` constrains the entire RE to match the entire line.

The null RE (for example, `/`) is equivalent to the last RE encountered. See also the last paragraph before FILES below.

To understand addressing in `ed` it is necessary to know that at any time there is a *current line*. Generally speaking, the current line is the last line affected by a command; the exact effect on the current line is discussed under the description of each command. *Addresses* are constructed as follows:

1. The character `.` addresses the current line.
2. The character `$` addresses the last line of the buffer.
3. A decimal number *n* addresses the *n*-th line of the buffer.
4. `'x` addresses the line marked with the mark name character *x*, which must be an ASCII lower-case letter (`a-z`). Lines are marked with the `k` command described below.
5. A RE enclosed by slashes (`/`) addresses the first line found by searching *forward* from the line *following* the current line toward the end of the buffer and stopping at the first line containing a string matching the RE. If necessary, the search wraps around to the beginning of the buffer and continues up to and including the current line, so that the entire buffer is searched. See also the last paragraph before FILES below.
6. A RE enclosed in question marks (`?`) addresses the first line found by searching *backward* from the line *preceding* the current line toward the beginning of the buffer and stopping at the first line containing a string matching the RE. If necessary, the search wraps around to the end of the buffer and continues up to and including the current line. See also the last paragraph before FILES below.

7. An address followed by a plus sign (+) or a minus sign (–) followed by a decimal number specifies that address plus (respectively minus) the indicated number of lines. A shorthand for .+5 is .5.
8. If an address begins with + or –, the addition or subtraction is taken with respect to the current line; for example, –5 is understood to mean .–5.
9. If an address ends with + or –, then 1 is added to or subtracted from the address, respectively. As a consequence of this rule and of Rule 8, immediately above, the address – refers to the line preceding the current line. (To maintain compatibility with earlier versions of the editor, the character ^ in addresses is entirely equivalent to –.) Moreover, trailing + and – characters have a cumulative effect, so — refers to the current line less 2.
10. For convenience, a comma (,) stands for the address pair 1,\$, while a semicolon (;) stands for the pair .,\$.

Commands may require zero, one, or two addresses. Commands that require no addresses regard the presence of an address as an error. Commands that accept one or two addresses assume default addresses when an insufficient number of addresses is given; if more addresses are given than such a command requires, the last one(s) are used.

Typically, addresses are separated from each other by a comma (,). They may also be separated by a semicolon (;). In the latter case, the first address is calculated, the current line (.) is set to that value, and then the second address is calculated. This feature can be used to determine the starting line for forward and backward searches (see Rules 5 and 6, above). The second address of any two-address sequence must correspond to a line in the buffer that follows the line corresponding to the first address.

In the following list of **ed** commands, the parentheses shown prior to the command are *not* part of the address; rather they show the default address(es) for the command.

It is generally illegal for more than one command to appear on a line. However, any command (except **e**, **f**, **r**, or **w**) may be suffixed by **l**, **n**, or **p** in which case the current line is either listed, numbered or printed, respectively, as discussed below under the **l**, **n**, and **p** commands.

(.)a
<text>

The **append** command accepts zero or more lines of text and appends it after the addressed line in the buffer. The current line (.) is left at the last inserted line, or, if there were none, at the addressed line. Address 0 is legal for this command: it causes the “appended” text to be placed at the beginning of the buffer. The maximum number of characters that may be entered from a terminal is 256 per line (including the new-line character).

(.)c
 <text>
 .

The **change** command deletes the addressed lines from the buffer, then accepts zero or more lines of text that replaces these lines in the buffer. The current line (.) is left at the last line input, or, if there were none, at the first line that was not deleted.

C

Same as the **X** command, described later, except that **ed** assumes all text read in for the **e** and **r** commands is encrypted unless a null key is typed in.

(..)d

The **delete** command deletes the addressed lines from the buffer. The line after the last line deleted becomes the current line; if the lines deleted were originally at the end of the buffer, the new last line becomes the current line.

e filename

The **edit** command deletes the entire contents of the buffer and then reads the contents of *filename* into the buffer. The current line (.) is set to the last line of the buffer. If *filename* is not given, the currently remembered file name, if any, is used (see the **f** command). The number of characters read in is printed; *filename* is remembered for possible use as a default file name in subsequent **e**, **r**, and **w** commands. If *filename* is replaced by **!**, the rest of the line is taken to be a shell (**sh(1)**) command whose output is to be read in. Such a shell command is *not* remembered as the current file name. See also DIAGNOSTICS below.

E filename

The **Edit** command is like **e**, except that the editor does not check to see if any changes have been made to the buffer since the last **w** command.

f filename

If *filename* is given, the **file-name** command changes the currently remembered file name to *filename*; otherwise, it prints the currently remembered file name.

(1, \$)g/RE/command list

In the **global** command, the first step is to mark every line that matches the given RE. Then, for every such line, the given *command list* is executed with the current line (.) initially set to that line. A single command or the first of a list of commands appears on the same line as the global command. All lines of a multi-line list except the last line must be ended with a ****; **a**, **i**, and **c** commands and associated input are permitted. The **.** terminating input mode may be omitted if it would be the last line of the *command list*. An empty *command list* is equivalent to the **p** command. The **g**, **G**, **v**, and **V** commands are *not* permitted in the *command list*. See also the NOTES and the last paragraph before FILES below.

(1, \$)G/RE/

In the interactive **G**lobal command, the first step is to mark every line that matches the given **RE**. Then, for every such line, that line is printed, the current line (.) is changed to that line, and any *one* command (other than one of the **a**, **c**, **i**, **g**, **G**, **v**, and **V** commands) may be input and is executed. After the execution of that command, the next marked line is printed, and so on; a new-line acts as a null command; an **&** causes the re-execution of the most recent command executed within the current invocation of **G**. Note: The commands input as part of the execution of the **G** command may address and affect *any* lines in the buffer. The **G** command can be terminated by an interrupt signal (ASCII DEL or BREAK).

h

The **help** command gives a short error message that explains the reason for the most recent ? diagnostic.

H

The **Help** command causes **ed** to enter a mode in which error messages are printed for all subsequent ? diagnostics. It will also explain the previous ? if there was one. The **H** command alternately turns this mode on and off; it is initially off.

(.)i

<text>

.

The **insert** command accepts zero or more lines of text and inserts it before the addressed line in the buffer. The current line (.) is left at the last inserted line, or, if there were none, at the addressed line. This command differs from the **a** command only in the placement of the input text. Address 0 is not legal for this command. The maximum number of characters that may be entered from a terminal is 256 per line (including the new-line character).

(.,.+1)j

The **join** command joins contiguous lines by removing the appropriate new-line characters. If exactly one address is given, this command does nothing.

(.)kx

The **mark** command marks the addressed line with name **x**, which must be an ASCII lower-case letter (**a-z**). The address 'x then addresses this line; the current line (.) is unchanged.

(.,.)l

The **list** command prints the addressed lines in an unambiguous way: a few non-printing characters (for example, *tab*, *backspace*) are represented by visually mnemonic overstrikes. All other non-printing characters are printed in octal, and long lines are folded. An **l** command may be appended to any command other than **e**, **f**, **r**, or **w**.

(...)*ma*

The **move** command repositions the addressed line(s) after the line addressed by *a*. Address **0** is legal for *a* and causes the addressed line(s) to be moved to the beginning of the file. It is an error if address *a* falls within the range of moved lines; the current line (.) is left at the last line moved.

(...)*n*

The **number** command prints the addressed lines, preceding each line by its line number and a tab character; the current line (.) is left at the last line printed. The **n** command may be appended to any command other than **e**, **f**, **r**, or **w**.

(...)*p*

The **print** command prints the addressed lines; the current line (.) is left at the last line printed. The **p** command may be appended to any command other than **e**, **f**, **r**, or **w**. For example, **dp** deletes the current line and prints the new current line.

P

The editor will prompt with a * for all subsequent commands. The **P** command alternately turns this mode on and off; it is initially off.

q

The **quit** command causes **ed** to exit. No automatic write of a file is done; however, see **DIAGNOSTICS**, below.

Q

The editor exits without checking if changes have been made in the buffer since the last **w** command.

(\$)*r filename*

The **read** command reads the contents of *filename* into the buffer. If *filename* is not given, the currently remembered file name, if any, is used (see the **e** and **f** commands). The currently remembered file name is *not* changed unless *filename* is the very first file name mentioned since **ed** was invoked. Address **0** is legal for **r** and causes the file to be read in at the beginning of the buffer. If the read is successful, the number of characters read in is printed; the current line (.) is set to the last line read in. If *filename* is replaced by **!**, the rest of the line is taken to be a shell (see **sh(1)**) command whose output is to be read in. For example, **\$r !ls** appends current directory to the end of the file being edited. Such a shell command is *not* remembered as the current file name.

(...)*s/RE/replacement/* or

(...)*s/RE/replacement/g* or

(...)*s/RE/replacement/n* *n* = 1-512

The **substitute** command searches each addressed line for an occurrence of the specified RE. In each line in which a match is found, all (non-overlapped) matched strings are replaced by the *replacement* if the global replacement indicator **g** appears after the command. If the global indicator does not appear, only the first occurrence of the matched string is replaced. If a number *n*, appears after the command, only the *n*-th occurrence of the matched string on each addressed line is replaced. It is an error if the substitution fails on *all* addressed

lines. Any character other than space or new-line may be used instead of / to delimit the RE and the *replacement*; the current line (.) is left at the last line on which a substitution occurred. See also the last paragraph before FILES below.

An ampersand (&) appearing in the *replacement* is replaced by the string matching the RE on the current line. The special meaning of & in this context may be suppressed by preceding it by \. As a more general feature, the characters \n, where n is a digit, are replaced by the text matched by the n-th regular subexpression of the specified RE enclosed between \ (and \). When nested parenthesized subexpressions are present, n is determined by counting occurrences of \ (starting from the left. When the character % is the only character in the *replacement*, the *replacement* used in the most recent substitute command is used as the *replacement* in the current substitute command. The % loses its special meaning when it is in a replacement string of more than one character or is preceded by a \.

A line may be split by substituting a new-line character into it. The new-line in the *replacement* must be escaped by preceding it by \. Such substitution cannot be done as part of a g or v command list.

(.,.)ta

This command acts just like the m command, except that a *copy* of the addressed lines is placed after address a (which may be 0); the current line (.) is left at the last line copied.

u

The **undo** command nullifies the effect of the most recent command that modified anything in the buffer, namely the most recent a, c, d, g, i, j, m, r, s, t, v, G, or V command.

(1, \$)v/RE/command list

This command is the same as the global command g, except that the lines marked during the first step are those that do *not* match the RE.

(1, \$)V/RE/

This command is the same as the interactive global command G, except that the lines that are marked during the first step are those that do *not* match the RE.

(1, \$)w filename

The **write** command writes the addressed lines into *filename*. If *filename* does not exist, it is created with mode **666** (readable and writable by everyone), unless your file creation mask dictates otherwise; see the description of the **umask** special command on **sh(1)**. The currently remembered file name is *not* changed unless *filename* is the very first file name mentioned since **ed** was invoked. If no file name is given, the currently remembered file name, if any, is used (see the e and f commands); the current line (.) is unchanged. If the command is successful, the number of characters written is printed. If *filename* is replaced by !, the rest of the line is taken to be a shell (see **sh(1)**) command whose standard input is the addressed lines. Such a shell command is *not* remembered as the current file name.

(1, \$)W *filename*

This command is the same as the write command above, except that it appends the addressed lines to the end of *filename* if it exists. If *filename* does not exist, it is created as described above for the *w* command.

X

An educated guess is made to determine whether text read in for the *e* and *r* commands is encrypted. A null key turns off encryption. Subsequent *e*, *r*, and *w* commands will use this key to encrypt or decrypt the text. An explicitly empty key turns off encryption. Also, see the *-x* option of *ed*.

(\$)=

The line number of the addressed line is typed; the current line (.) is unchanged by this command.

!shell *command*

The remainder of the line after the ! is sent to the UNIX system shell (see *sh*(1)) to be interpreted as a command. Within the text of that command, the unescaped character % is replaced with the remembered file name; if a ! appears as the first character of the shell command, it is replaced with the text of the previous shell command. Thus, !! will repeat the last shell command. If any expansion is performed, the expanded line is echoed; the current line (.) is unchanged.

(.+1)<new-line>

An address alone on a line causes the addressed line to be printed. A new-line alone is equivalent to *+.1p*; it is useful for stepping forward through the buffer.

If an interrupt signal (ASCII DEL or BREAK) is sent, *ed* prints a ? and returns to *its* command level.

Some size limitations: 512 characters in a line, 256 characters in a global command list, and 255 characters in the pathname of a file (counting slashes). The limit on the number of lines depends on the amount of user memory: each line takes 1 word.

When reading a file, *ed* discards ASCII NUL characters.

If a file is not terminated by a new-line character, *ed* adds one and puts out a message explaining what it did.

If the closing delimiter of a RE or of a replacement string (for example, /) would be the last character before a new-line, that delimiter may be omitted, in which case the addressed line is printed. The following pairs of commands are equivalent:

s/s1/s2	s/s1/s2/p
g/s1	g/s1/p
?s1	?s1?

OPTIONS

- C** Encryption option; the same as the **-x** option, except that **ed** simulates a **C** command. The **C** command is like the **X** command, except that all text read in is assumed to have been encrypted.
- p** Allows the user to specify a prompt string.
- s** Suppresses the printing of character counts by **e**, **r**, and **w** commands, of diagnostics from **e** and **q** commands, and of the **!** prompt after a *!shell command*.
- x** Encryption option; when used, **ed** simulates an **X** command and prompts the user for a key. The **X** command makes an educated guess to determine whether text read in is encrypted or not.
The temporary buffer file is encrypted also, using a transformed version of the key typed in for the **-x** option. Also, see the NOTES section at the end of this manual page.

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

FILES

- \$TMPDIR** If this environment variable is not null, its value is used in place of **/var/tmp** as the directory name for the temporary work file.
- /var/tmp** If **/var/tmp** exists, it is used as the directory name for the temporary work file.
- /tmp** If the environment variable **TMPDIR** does not exist or is null, and if **/var/tmp** does not exist, then **/tmp** is used as the directory name for the temporary work file.
- ed.hup** Work is saved here if the terminal is hung up.

SEE ALSO

edit(1), ex(1), grep(1), sed(1), sh(1), stty(1), umask(1), vi(1), fspec(4), environ(5), regexp(5)

DIAGNOSTICS

? for command errors.

?filename for an inaccessible file.

(use the **help** and **Help** commands for detailed explanations).

If changes have been made in the buffer since the last **w** command that wrote the entire buffer, **ed** warns the user if an attempt is made to destroy **ed**'s buffer via the **e** or **q** commands. It prints **?** and allows one to continue editing. A second **e** or **q** command at this point will take effect. The **-s** command-line option inhibits this feature.

NOTES

The **-** option, although it continues to be supported, has been replaced in the documentation by the **-s** option that follows the Command Syntax Standard (see **intro(1)**).

The encryption options and commands are provided with the Security Administration Utilities package, which is available only in the United States.

A **!** command cannot be subject to a **g** or a **v** command.

The **!** command and the **!** escape from the **e**, **r**, and **w** commands cannot be used if the editor is invoked from a restricted shell (see **sh(1)**).

The sequence **\n** in a RE does not match a new-line character.

If the editor input is coming from a command file (for example, **ed filename < ed_cmd_file**), the editor exits at the first failure.

NAME	edit – text editor (variant of ex for casual users)								
SYNOPSIS	edit [- -s] [-I] [-L] [-R] [-r [<i>filename</i>]] [-t <i>tag</i>] [-v] [-V] [-x] [-wn] [-C] [+ <i>command</i> -c <i>command</i>] <i>filename</i> . . .								
AVAILABILITY	SUNWcsu								
DESCRIPTION	<p>edit is a variant of the text editor ex recommended for new or casual users who wish to use a command-oriented editor. It operates precisely as ex with the following options automatically set:</p> <table border="0" style="margin-left: 40px;"> <tr> <td>novice</td> <td>ON</td> </tr> <tr> <td>report</td> <td>ON</td> </tr> <tr> <td>showmode</td> <td>ON</td> </tr> <tr> <td>magic</td> <td>OFF</td> </tr> </table> <p>The following brief introduction should help you get started with edit. If you are using a CRT terminal you may want to learn about the display editor vi.</p> <p>To edit the contents of an existing file you begin with the command edit name to the shell. edit makes a copy of the file that you can then edit, and tells you how many lines and characters are in the file. To create a new file, you also begin with the command edit with a filename: edit name; the editor will tell you it is a [New File].</p> <p>The edit command prompt is the colon (:), which you should see after starting the editor. If you are editing an existing file, then you will have some lines in edit's buffer (its name for the copy of the file you are editing). When you start editing, edit makes the last line of the file the current line. Most commands to edit use the current line if you do not tell them which line to use. Thus if you say print (which can be abbreviated p) and type carriage return (as you should after all edit commands), the current line will be printed. If you delete (d) the current line, edit will print the new current line, which is usually the next line in the file. If you delete the last line, then the new last line becomes the current one.</p> <p>If you start with an empty file or wish to add some new lines, then the append (a) command can be used. After you execute this command (typing a carriage return after the word append), edit will read lines from your terminal until you type a line consisting of just a dot (.); it places these lines after the current line. The last line you type then becomes the current line. The insert (i) command is like append, but places the lines you type before, rather than after, the current line.</p> <p>edit numbers the lines in the buffer, with the first line having number 1. If you execute the command 1, then edit will type the first line of the buffer. If you then execute the command d, edit will delete the first line, line 2 will become line 1, and edit will print the current line (the new line 1) so you can see where you are. In general, the current line will always be the last line affected by a command.</p>	novice	ON	report	ON	showmode	ON	magic	OFF
novice	ON								
report	ON								
showmode	ON								
magic	OFF								

You can make a change to some text within the current line by using the **substitute (s)** command: **s/old/new/** where *old* is the string of characters you want to replace and *new* is the string of characters you want to replace *old* with.

The **filename (f)** command will tell you how many lines there are in the buffer you are editing and will say **[Modified]** if you have changed the buffer. After modifying a file, you can save the contents of the file by executing a **write (w)** command. You can leave the editor by issuing a **quit (q)** command. If you run **edit** on a file, but do not change it, it is not necessary (but does no harm) to **write** the file back. If you try to **quit** from **edit** after modifying the buffer without writing it out, you will receive the message **No write since last change (:quit! overrides)**, and **edit** will wait for another command. If you do not want to write the buffer out, issue the **quit** command followed by an exclamation point (**q!**). The buffer is then irretrievably discarded and you return to the shell.

By using the **d** and **a** commands and giving line numbers to see lines in the file, you can make any changes you want. You should learn at least a few more things, however, if you will use **edit** more than a few times.

The **change (c)** command changes the current line to a sequence of lines you supply (as in **append**, you type lines up to a line consisting of only a dot (.). You can tell **change** to change more than one line by giving the line numbers of the lines you want to change, that is, **3,5c**. You can print lines this way too: **1,23p** prints the first 23 lines of the file.

The **undo (u)** command reverses the effect of the last command you executed that changed the buffer. Thus if you execute a **substitute** command that does not do what you want, type **u** and the old contents of the line will be restored. You can also **undo** an **undo** command. **edit** will give you a warning message when a command affects more than one line of the buffer. Note that commands such as **write** and **quit** cannot be undone.

To look at the next line in the buffer, type carriage return. To look at a number of lines, type **^D** (while holding down the control key, press **d**) rather than carriage return. This will show you a half-screen of lines on a CRT or 12 lines on a hardcopy terminal. You can look at nearby text by executing the **z** command. The current line will appear in the middle of the text displayed, and the last line displayed will become the current line; you can get back to the line where you were before you executed the **z** command by typing **``**. The **z** command has other options: **z-** prints a screen of text (or 24 lines) ending where you are; **z+** prints the next screenful. If you want less than a screenful of lines, type **z.11** to display five lines before and five lines after the current line. (Typing **z.n**, when *n* is an odd number, displays a total of *n* lines, centered about the current line; when *n* is an even number, it displays *n*-1 lines, so that the lines displayed are centered around the current line.) You can give counts after other commands; for example, you can delete 5 lines starting with the current line with the command **d5**.

To find things in the file, you can use line numbers if you happen to know them; since the line numbers change when you insert and delete lines this is somewhat unreliable. You can search backwards and forwards in the file for strings by giving commands of the form **/text/** to search forward for *text* or **?text?** to search backward for *text*. If a search reaches the end of the file without finding *text*, it wraps around and continues to search

back to the line where you are. A useful feature here is a search of the form `/^text/` which searches for *text* at the beginning of a line. Similarly `/text$/` searches for *text* at the end of a line. You can leave off the trailing `/` or `?` in these commands.

The current line has the symbolic name `dot` (`.`); this is most useful in a range of lines as in `.,$p` which prints the current line plus the rest of the lines in the file. To move to the last line in the file, you can refer to it by its symbolic name `$`. Thus the command `$d` deletes the last line in the file, no matter what the current line is. Arithmetic with line references is also possible. Thus the line `$-5` is the fifth before the last and `+.20` is 20 lines after the current line.

You can find out the current line by typing `='.='.` This is useful if you wish to move or copy a section of text within a file or between files. Find the first and last line numbers you wish to copy or move. To move lines 10 through 20, type `10,20d a` to delete these lines from the file and place them in a buffer named **a**. `edit` has 26 such buffers named **a** through **z**. To put the contents of buffer **a** after the current line, type `put a`. If you want to move or copy these lines to another file, execute an `edit` (**e**) command after copying the lines; following the **e** command with the name of the other file you wish to edit, that is, `edit chapter2`. To copy lines without deleting them, use `yank` (**y**) in place of **d**. If the text you wish to move or copy is all within one file, it is not necessary to use named buffers. For example, to move lines 10 through 20 to the end of the file, type `10,20m $`.

OPTIONS

These options can be turned on or off using the `set` command in `ex(1)`.

- `- | -s` Suppress all interactive user feedback. This is useful when processing editor scripts.
- `-l` Set up for editing LISP programs.
- `-L` List the name of all files saved as the result of an editor or system crash.
- `-R` **Readonly** mode; the **readonly** flag is set, preventing accidental overwriting of the file.
- `-r filename` Edit *filename* after an editor or system crash. (Recovers the version of *filename* that was in the buffer when the crash occurred.)
- `-t tag` Edit the file containing the *tag* and position the editor at its definition.
- `-v` Start up in display editing state using `vi`. You can achieve the same effect by simply typing the `vi` command itself.
- `-V` Verbose. Any non-tty input will be echoed on standard error. This may be useful when processing editor commands within shell scripts.
- `-x` Encryption option; when used, `edit` simulates the **X** command of `ex` and prompts the user for a key. This key is used to encrypt and decrypt text using the algorithm of the `crypt` command. The **X** command makes an educated guess to determine whether text read in is encrypted or not. The temporary buffer file is encrypted also, using a transformed version of the key typed in for the `-x` option.
- `-wn` Set the default window size to *n*. This is useful when using the editor over a slow speed line.

-C Encryption option; same as the **-x** option, except that **vi** simulates the **C** command of **ex**. The **C** command is like the **X** command of **ex**, except that all text read in is assumed to have been encrypted.

+command | **-c command**

Begin editing by executing the specified editor *command* (usually a search or positioning command).

The *filename* argument indicates one or more files to be edited.

SEE ALSO **ed(1)**, **ex(1)**, **vi(1)**

NOTES The encryption options are provided with the Security Administration Utilities package, which is available only in the United States.

NAME	egrep – search a file for a pattern using full regular expressions
SYNOPSIS	egrep [-bchilnsv] [-e <i>special-expression</i>] [-f <i>filename</i>] [<i>strings</i>] [<i>filename ...</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>egrep (<i>expression grep</i>) searches files for a pattern of characters and prints all lines that contain that pattern. egrep uses full regular expressions (expressions that have string values that use the full set of alphanumeric and special characters) to match the patterns. It uses a fast deterministic algorithm that sometimes needs exponential space.</p> <p>egrep accepts full regular expressions as in ed(1), except for \ (and \), with the addition of:</p> <ol style="list-style-type: none"> 1. A full regular expression followed by + that matches one or more occurrences of the full regular expression. 2. A full regular expression followed by ? that matches 0 or 1 occurrences of the full regular expression. 3. Full regular expressions separated by or by a new-line that match strings that are matched by any of the expressions. 4. A full regular expression that may be enclosed in parentheses () for grouping. <p>Be careful using the characters \$, *, [, ^, , (,), and \ in <i>full regular expression</i>, because they are also meaningful to the shell. It is safest to enclose the entire <i>full regular expression</i> in single quotes '... '.</p> <p>The order of precedence of operators is [], then *?+, then concatenation, then and new-line.</p> <p>If no files are specified, egrep assumes standard input. Normally, each line found is copied to the standard output. The file name is printed before each line found if there is more than one input file.</p>
OPTIONS	<p>-b Precede each line by the block number on which it was found. This can be useful in locating block numbers by context (first block is 0).</p> <p>-c Print only a count of the lines that contain the pattern.</p> <p>-h Suppress printing of filenames when searching multiple files.</p> <p>-i Ignore upper/lower case distinction during comparisons.</p> <p>-l Print the names of files with matching lines once, separated by new-lines. Does not repeat the names of files when the pattern is found more than once.</p> <p>-n Precede each line by its line number in the file (first line is 1).</p> <p>-s Work silently, that is, display nothing except error messages. This is useful for checking the error status.</p> <p>-v Print all lines except those that contain the pattern.</p> <p>-e <i>special-expression</i> Search for a <i>special expression</i> (<i>full regular expression</i> that begins with a -).</p>

-f filename Take the list of *full regular expressions* from *filename*.

SEE ALSO **ed(1), fgrep(1), grep(1), sed(1), sh(1)**

DIAGNOSTICS Exit status is 0 if any matches are found, 1 if none, 2 for syntax errors or inaccessible files (even if matches were found).

NOTES Ideally there should be only one **grep** command, but there is not a single algorithm that spans a wide enough range of space-time tradeoffs. Lines are limited to BUFSIZ characters; longer lines are truncated. BUFSIZ is defined in **<stdio.h>**.

NAME	eject – eject media such as CD-ROM and floppy from drive
SYNOPSIS	eject [-dfnq] [<i>device</i> <i>nickname</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>eject is used for those removable media devices that do not have a manual eject button, or for those that do, but are managed by Volume Management. The device may be specified by its name or by a nickname; if no device is specified the default device is used.</p> <p>Only devices that support eject under program control respond to this command. When eject is used on media that can only be ejected manually, it will do everything except remove the media — including unmounting the file system if it is mounted (see rmount(1M)). In this case, eject displays a message that the media can now be manually ejected. If a window system is running, the message is displayed as a pop-up window. If no window system is running, a message is displayed both to stderr and to the system console that the media can now be physically removed.</p> <p>It is not recommended to physically eject media from a device which contains mounted filesystems. eject automatically searches for any mounted filesystems which reside on the device and attempts to umount them prior to ejecting the media (see mount(1M)). If the unmount operation fails, eject prints a warning message and exits. The -f flag may be used to specify an eject <i>even</i> if the device contains mounted partitions.</p> <p>eject can also display its default device and a list of nicknames.</p> <p>If you have inserted a floppy diskette, you must use volcheck(1) before ejecting the media to inform Volume Management of the floppy's presence.</p>
OPTIONS	<p>-d Display the name of the default device to be ejected.</p> <p>-f Force the device to eject even if it is busy.</p> <p>-n Display the nickname to device name translation table.</p> <p>-q Query to see if the media is present.</p> <p><i>device</i> Specify which device to eject, by the name it appears in the directory /dev.</p> <p><i>nickname</i> Specify which device to eject, by its nickname as known to this command.</p>
EXIT CODES	<p>eject returns the following exit codes:</p> <p>0 If the operation was successful or, with the -q option, the media <i>is</i> in the drive.</p> <p>1 If the operation was unsuccessful or, with the -q option, the media is <i>not</i> in the drive.</p> <p>2 If invalid flags were specified.</p> <p>3 If an ioctl() request failed.</p>

4 Manually ejectable media is now okay to remove.

FILES

/dev/rfd0c
/dev/rsr0 raw files
/dev/sr0 block files
/usr/lib/vold/eject_popup popup used for manually ejected media

SEE ALSO

volcancel(1), **volcheck(1)**, **volmissing(1)**, **mount(1M)**, **rmmount(1M)**, **vold(1M)**,
rmmount.conf(4), **vold.conf(4)**, **volfs(7)**

DIAGNOSTICS

A short help message is printed if an unknown flag is specified. A diagnostic is printed if the device name cannot be opened or does not support **eject**.

Device Busy An attempt was made to eject a device that has a mounted filesystem. A warning message is printed when doing a forced eject of a mounted device.

BUGS

There should be a way to change the default on a per-user basis.

NAME	enable, disable – enable/disable LP printers
SYNOPSIS	<code>/usr/bin/enable printer ...</code> <code>/usr/bin/disable [-c -W] [-r [reason]] printer ...</code>
AVAILABILITY	SUNWlpu
DESCRIPTION	<p>The enable command activates the named <i>printers</i>, enabling them to print requests submitted by the lp command. If the printer is remote, the command will only enable the transfer of requests to the remote system; the enable command must be run again, on the remote system, to activate the printer. (Run lpstat -p to get the status of <i>printers</i>.)</p> <p>The disable command deactivates the named <i>printer</i>, disabling it from printing requests submitted by lp. By default, any requests that are currently printing on the designated printer(s) will be reprinted in their entirety either on the same printer or on another member of the same class of printers. If the printer is remote, this command will only stop the transmission of jobs to the remote system. The disable command must be run on the remote system to disable the printer. (Run lpstat -p to get the status of <i>printers</i>.)</p>
OPTIONS	<p>Options for use with disable are:</p> <p>-c Cancel any requests that are currently printing on the designated printer(s). This option cannot be used with the -W option. If the printer is remote, the -c option will be silently ignored.</p> <p>-W Wait until the request currently being printed is finished before disabling the specified printer. This option cannot be used with the -c option. If the printer is remote, the -W option will be silently ignored.</p> <p>-r reason Assign a <i>reason</i> for the disabling of the printer(s). This <i>reason</i> applies to all <i>printers</i> specified. This <i>reason</i> is reported by lpstat -p. <i>reason</i> must be enclosed in quotes if it contains blanks. The default reason is unknown reason for the existing printer, and new printer for a printer just added to the system but not yet enabled.</p>
FILES	<code>/var/spool/lp/*</code>
SEE ALSO	lp(1) , lpstat(1)

NAME	env – obtain or alter environment variables for command execution
SYNOPSIS	env [-] [<i>name=value</i>] ... [<i>command args</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>env obtains the current <i>environment</i>, modifies it according to its arguments, then executes the command with the modified environment. Arguments of the form <i>name=value</i> are merged into the inherited environment before the command is executed. The - flag causes the inherited environment to be ignored completely, so that the command is executed with exactly the environment specified by the arguments.</p> <p>If no command is specified, the resulting environment is printed, one name-value pair per line.</p>
OPTIONS	<p>- Ignore the environment that would otherwise be inherited from the current shell. Restricts the environment for <i>command</i> to that specified by the arguments.</p> <p><i>name=value</i> Set the environment variable <i>filename</i> to <i>value</i> and add it to the environment.</p>
SEE ALSO	sh(1), exec(2), profile(4), environ(5)

NAME	eqn, neqn, checkeq – typeset mathematics test
SYNOPSIS	eqn [-dxy] [-fn] [-pn] [-sn] [<i>filename</i>] ... neqn [<i>filename</i>] ... checkeq [<i>filename</i>] ...
AVAILABILITY	SUNWdoc
DESCRIPTION	eqn and neqn are language processors to assist in describing equations. eqn is a preprocessor for troff (1) and is intended for devices that can print troff 's output. neqn is a preprocessor for nroff (1) and is intended for use with terminals. Usage is almost always: <pre style="margin-left: 4em;">example% eqn filename ... troff example% neqn filename ... nroff</pre> <p>If no <i>filenames</i> are specified, eqn and neqn read from the standard input. A line beginning with .EQ marks the start of an equation; the end of an equation is marked by a line beginning with .EN. Neither of these lines is altered, so they may be defined in macro packages to get centering, numbering, etc. It is also possible to set two characters as “delimiters”; subsequent text between delimiters is also treated as eqn input.</p> <p>checkeq reports missing or unbalanced delimiters and .EQ/.EN pairs.</p>
OPTIONS	-dxy Set equation delimiters set to characters <i>x</i> and <i>y</i> with the command-line argument. The more common way to do this is with delimxy between .EQ and .EN . The left and right delimiters may be identical. Delimiters are turned off by delim off appearing in the text. All text that is neither between delimiters nor between .EQ and .EN is passed through untouched.
	-fn Change font to <i>n</i> globally in the document. The font can also be changed globally in the body of the document by using the gfont directive.
	-pn Reduce subscripts and superscripts by <i>n</i> point sizes from the previous size. In the absence of the -p option, subscripts and superscripts are reduced by 3 point sizes from the previous size.
	-sn Change point size to <i>n</i> globally in the document. The point size can also be changed globally in the body of the document by using the gsize directive.
EQN LANGUAGE	Tokens within eqn are separated by braces, double quotes, tildes, circumflexes, SPACE, TAB, or NEWLINE characters. Braces { } are used for grouping; generally speaking, anywhere a single character like <i>x</i> could appear, a complicated construction enclosed in braces may be used instead. Tilde (~) represents a full SPACE in the output, circumflex (^) half as much.
	Subscripts and superscripts are produced with the keywords sub and sup . Thus ‘ x sub i ’ makes x_i , ‘ a sub i sup 2 ’ produces a_i^2 , and ‘ e sup {x sup 2 + y sup 2} ’ gives $e^{x^2+y^2}$.

Fractions are made with **over**: 'a over b' yields $\frac{a}{b}$.

sqrt makes square roots: '1 over sqrt {ax sup 2 +bx+c}' results in $\frac{1}{\sqrt{ax^2+bx+c}}$.

The keywords **from** and **to** introduce lower and upper limits on arbitrary things: $\lim_{n \rightarrow \infty} \sum_0^n x_i$ is made with 'lim from {n → inf } sum from 0 to n x sub i'.

Left and right brackets, braces, etc., of the right height are made with **left** and **right**: 'left [x sup 2 + y sup 2 over alpha right] ~ ~ 1' produces $\left[x^2 + \frac{y^2}{\alpha} \right] = 1$. The **right** clause is optional. Legal characters after **left** and **right** are braces, brackets, bars, **c** and **f** for ceiling and floor, and "" for nothing at all (useful for a right-side-only bracket).

Vertical piles of things are made with **pile**, **lpile**, **cpile**, and **rpile**: 'pile {a above b above c}' produces $\begin{matrix} a \\ b \\ c \end{matrix}$. There can be an arbitrary number of elements in a pile. **lpile** left-justifies, **cpile** and **cpile** center, with different vertical spacing, and **rpile** right justifies.

Matrices are made with **matrix**: 'matrix { lcol { x sub i above y sub 2 } ccol { 1 above 2 } }' produces $\begin{matrix} x_1 & 1 \\ y_2 & 2 \end{matrix}$. In addition, there is **rcol** for a right-justified column.

Diacritical marks are made with **dot**, **dotdot**, **hat**, **tilde**, **bar**, **vec**, **dyad**, and **under**: 'x dot = f(t) bar' is $\dot{x} = f(t)$, 'y dotdot bar ~ ~ n under' is $\ddot{y} = \underline{n}$, and 'x vec ~ ~ y dyad' is $\vec{x} = \vec{y}$.

Sizes and font can be changed with **size n** or **size ±n**, **roman**, **italic**, **bold**, and **font n**. Size and fonts can be changed globally in a document by **gsize n** and **gfont n**, or by the command-line arguments **-sn** and **-fn**.

Successive display arguments can be lined up. Place **mark** before the desired lineup point in the first equation; place **lineup** at the place that is to line up vertically in subsequent equations.

Shorthands may be defined or existing keywords redefined with **define**:

define thing % replacement %

defines a new token called *thing* which will be replaced by *replacement* whenever it appears thereafter. The % may be any character that does not occur in *replacement*.

Keywords like **sum** (Σ), **int** (\int), **inf** (∞), and shorthands like **>=** (\geq), **→** (\rightarrow), and **!=** (\neq) are recognized. Greek letters are spelled out in the desired case, as in **alpha** or **GAMMA**. Mathematical words like **sin**, **cos**, and **log** are made Roman automatically. **troff(1)** four-character escapes like **\(bu** (\bullet) can be used anywhere. Strings enclosed in double quotes **".."** are passed through untouched; this permits keywords to be entered as text, and can be used to communicate with **troff** when all else fails.

SEE ALSO **tbl(1), troff(1), ms(5)**

BUGS To embolden digits, parens, etc., it is necessary to quote them, as in '**bold "12.3"**'.

NAME	error – insert compiler error messages at right source lines
SYNOPSIS	error [-n] [-q] [-s] [-v] [-t <i>suffixlist</i>] [-I <i>ignorefile</i>] [<i>filename</i>]
DESCRIPTION	<p>error analyzes error messages produced by a number of compilers and language processors. It replaces the painful, traditional methods of scribbling abbreviations of errors on paper, and permits error messages and source code to be viewed simultaneously.</p> <p>error looks at error messages, either from the specified file <i>filename</i> or from the standard input, and:</p> <ul style="list-style-type: none"> • Determines which language processor produced each error message. • Determines the file name and line number of the erroneous line. • Inserts the error message into the source file immediately preceding the erroneous line. <p>Error messages that can't be categorized by language processor or content are not inserted into any file, but are sent to the standard output. error touches source files only after all input has been read.</p> <p>error is intended to be run with its standard input connected with a pipe to the error message source. Some language processors put error messages on their standard error file; others put their messages on the standard output. Hence, both error sources should be piped together into error. For example, when using the csh syntax, the following command analyzes all the error messages produced by whatever programs make(1S) runs when making lint:</p> <pre style="margin-left: 40px;">example% make -s lint & error</pre> <p>error knows about the error messages produced by: as(1), cpp(1), ld(1), cc(1B), make(1S) and other compilers. For all languages except Pascal, error messages are restricted to one line. Some error messages refer to more than one line in more than one file, in which case error duplicates the error message and inserts it in all the appropriate places.</p>
OPTIONS	<p>-n Do <i>not</i> touch any files; all error messages are sent to the standard output.</p> <p>-q error asks whether the file should be touched. A 'y' or 'n' to the question is necessary to continue. Absence of the -q option implies that all referenced files (except those referring to discarded error messages) are to be touched.</p> <p>-s Print out statistics regarding the error categorization.</p> <p>-v After all files have been touched, overlay the visual editor vi with it set up to edit all files touched, and positioned in the first touched file at the first error. If vi(1) can't be found, try ex(1) or ed(1) from standard places.</p>

-t suffixlist

Take the following argument as a suffix list. Files whose suffices do not appear in the suffix list are not touched. The suffix list is dot separated, and '*' wildcards work. Thus the suffix list:

.c.y.f*.h

allows **error** to touch files ending with '.c', '.y', '.f*' and '.h'.

error catches interrupt and terminate signals, and terminates in an orderly fashion.

EXAMPLES

In the following C shell (**/usr/bin/csh**) example, **error** takes its input from the FORTRAN compiler:

```
example% f77 -c any.f | & error" options
```

Here is the same example using the Korn shell (**/usr/bin/ksh**):

```
example% f77 -c any.f 2>&1 | error" options
```

USAGE

error does one of six things with error messages.

synchronize Some language processors produce short errors describing which file they are processing. **error** uses these to determine the file name for languages that do not include the file name in each error message. These synchronization messages are consumed entirely by **error**.

discard Error messages from **lint** that refer to one of the two **lint** libraries, **/usr/lib/lint/lib-lc** and **/usr/lib/lint/lib-port** are discarded, to prevent accidentally touching these libraries. Again, these error messages are consumed entirely by **error**.

nullify Error messages from **lint** can be nullified if they refer to a specific function, which is known to generate diagnostics which are not interesting. Nullified error messages are not inserted into the source file, but are written to the standard output. The names of functions to ignore are taken from either the file named **.errorrc** in the user's home directory, or from the file named by the **-I** option. If the file does not exist, no error messages are nullified. If the file does exist, there must be one function name per line.

not file specific Error messages that can't be intuited are grouped together, and written to the standard output before any files are touched. They are not inserted into any source file.

file specific Error messages that refer to a specific file but to no specific line are written to the standard output when that file is touched.

true errors Error messages that can be intuited are candidates for insertion into the file to which they refer.

Only true error messages are inserted into source files. Other error messages are consumed entirely by **error** or are written to the standard output. **error** inserts the error messages into the source file on the line preceding the line number in the error message.

Each error message is turned into a one line comment for the language, and is internally flagged with the string `###` at the beginning of the error, and `%%%` at the end of the error. This makes pattern searching for errors easier with an editor, and allows the messages to be easily removed. In addition, each error message contains the source line number for the line the message refers to. A reasonably formatted source program can be recompiled with the error messages still in it, without having the error messages themselves cause future errors. For poorly formatted source programs in free format languages, such as C or Pascal, it is possible to insert a comment into another comment, which can wreak havoc with a future compilation. To avoid this, format the source program so there are no language statements on the same line as the end of a comment.

FILES `~/errorrc` function names to ignore for **lint** error messages
`/dev/tty` user's teletype

SEE ALSO `as(1)`, `cpp(1)`, `cs(1)`, `ed(1)`, `ex(1)`, `ld(1)`, `vi(1)`, `cc(1B)`, `make(1S)`

BUGS Opens the tty-device directly for user input.
 Source files with links make a new copy of the file with only one link to it.
 Changing a language processor's error message format may cause **error** to not understand the error message.
error, since it is purely mechanical, will not filter out subsequent errors caused by "floodgating" initiated by one syntactically trivial error. Humans are still much better at discarding these related errors.
 Pascal error messages belong after the lines affected, **error** puts them before. The alignment of the `|` marking the point of error is also disturbed by **error**.
error was designed for work on CRT's at reasonably high speed. It is less pleasant on slow speed terminals, and was not designed for use on hardcopy terminals.

NAME	ex – text editor
SYNOPSIS	ex [- -s] [-I] [-L] [-R] [-r [<i>filename</i>]] [-t <i>tag</i>] [-v] [-V] [-x] [-wn] [-C] [+ <i>command</i> -c <i>command</i>] <i>filename</i> . . .
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>ex is the root of a family of editors: ex and vi. ex is a superset of ed, with the most notable extension being a display editing facility. Display based editing is the focus of vi. If you have a CRT terminal, you may wish to use a display based editor; in this case see vi(1), which is a command which focuses on the display-editing portion of ex.</p>
For ed Users	<p>If you have used ed you will find that, in addition to having all of the ed commands available, ex has a number of additional features useful on CRT terminals. Intelligent terminals and high speed terminals are very pleasant to use with vi. Generally, the ex editor uses far more of the capabilities of terminals than ed does, and uses the terminal capability data base (see terminfo(4)) and the type of the terminal you are using from the environment variable TERM to determine how to drive your terminal efficiently. The editor makes use of features such as insert and delete character and line in its visual command (which can be abbreviated vi) and which is the central mode of editing when using the vi command.</p> <p>ex contains a number of features for easily viewing the text of the file. The z command gives easy access to windows of text. Typing ^D (control-D) causes the editor to scroll a half-window of text and is more useful for quickly stepping through a file than just typing return. Of course, the screen-oriented visual mode gives constant access to editing context.</p> <p>ex gives you help when you make mistakes. The undo (u) command allows you to reverse any single change which goes astray. ex gives you a lot of feedback, normally printing changed lines, and indicates when more than a few lines are affected by a command so that it is easy to detect when a command has affected more lines than it should have.</p> <p>The editor also normally prevents overwriting existing files, unless you edited them, so that you do not accidentally overwrite a file other than the one you are editing. If the system (or editor) crashes, or you accidentally hang up the telephone, you can use the editor recover command (or -r filename option) to retrieve your work. This will get you back to within a few lines of where you left off.</p> <p>ex has several features for dealing with more than one file at a time. You can give it a list of files on the command line and use the next (n) command to deal with each in turn. The next command can also be given a list of file names, or a pattern as used by the shell to specify a new set of files to be dealt with. In general, file names in the editor may be formed with full shell metasyntax. The metacharacter '%' is also available in forming file names and is replaced by the name of the current file.</p>

The editor has a group of buffers whose names are the ASCII lower-case letters (a-z). You can place text in these named buffers where it is available to be inserted elsewhere in the file. The contents of these buffers remain available when you begin editing a new file using the **edit** (e) command.

There is a command **&** in **ex** which repeats the last **substitute** command. In addition, there is a confirmed substitute command. You give a range of substitutions to be done and the editor interactively asks whether each substitution is desired.

It is possible to ignore the case of letters in searches and substitutions. **ex** also allows regular expressions which match words to be constructed. This is convenient, for example, in searching for the word "edit" if your document also contains the word "editor."

ex has a set of options which you can set to tailor it to your liking. One option which is very useful is the **autoindent** option that allows the editor to supply leading white space to align text automatically. You can then use **^D** as a backtab and space or tab to move forward to align new code easily.

Miscellaneous useful features include an intelligent **join** (j) command that supplies white space between joined lines automatically, commands **<** and **>** which shift groups of lines, and the ability to filter portions of the buffer through commands such as **sort**.

OPTIONS

- | -s** Suppress all interactive user feedback. This is useful when processing editor scripts.
- I** Set up for editing LISP programs.
- L** List the name of all files saved as the result of an editor or system crash.
- R** **Readonly** mode; the **readonly** flag is set, preventing accidental overwriting of the file.
- r filename** Edit *filename* after an editor or system crash. (Recovers the version of *filename* that was in the buffer when the crash occurred.)
- t tag** Edit the file containing the *tag* and position the editor at its definition.
- v** Start up in display editing state using **ex**. You can achieve the same effect by simply typing the **-ex** command itself.
- V** Verbose. Any non-tty input will be echoed on standard error. This may be useful when processing editor commands within shell scripts.
- x** Encryption option. Simulates the **X** command and prompts the user for a key. This key is used to encrypt and decrypt text using the algorithm of the **crypt** command. The **X** command makes an educated guess to determine whether text read in is encrypted or not. The temporary buffer file is encrypted also, using a transformed version of the key typed in for the **-x** option.
- wn** Set the default window size to *n*. This is useful when using the editor over a slow speed line.
- C** Encryption option. Same as the **-x** option, except simulates the **C** command. The **C** command is like the **X** command, except that all text read

in is assumed to have been encrypted.

+command | **-c command**

Begin editing by executing the specified editor *command* (usually a search or positioning command).

ex States

Command Normal and initial state. Input prompted for by “:”. Your line kill character cancels a partial command.

Insert Entered by **a**, **i**, or **c**. Arbitrary text may be entered. Insert state normally is terminated by a line having only "." on it, or, abnormally, with an interrupt.

Visual Entered by typing **vi**; terminated by typing **Q** or **^\
(control- \backslash).**

ex Command Names and Abbreviations

abbrev	ab	map		set	se
append	a	mark	ma	shell	sh
args	ar	move	m	source	so
change	c	next	n	substitute	s
copy	co	number	nu	unabbrev	unab
delete	d	preserve	pre	undo	u
edit	e	print	p	unmap	unm
file	f	put	pu	version	ve
global	g	quit	q	visual	vi
insert	i	read	r	write	w
join	j	recover	rec	xit	x
list	l	rewind	rew	yank	ya

ex Commands

forced encryption	C	heuristic encryption	X
resubst	&	print next	CR
rshift	>	lshift	<
scroll	^D	window	z
shell escape	!		

ex Command Addresses

<i>n</i>	line <i>n</i>	<i>/pat</i>	next with <i>pat</i>
.	current	<i>?pat</i>	previous with <i>pat</i>
\$	last	<i>x-n</i>	<i>n</i> before <i>x</i>
+	next	<i>x,y</i>	<i>x</i> through <i>y</i>
-	previous	<i>^x</i>	marked with <i>x</i>
+n	<i>n</i> forward	<i>``</i>	previous context
%	1, <i>\$</i>		

Initializing options	EXINIT		place set 's here in environment variable
	\$HOME/.exrc		editor initialization file
	./exrc		editor initialization file
	set x		enable option <i>x</i>
	set nox		disable option <i>x</i>
	set x=val		give value <i>val</i> to option <i>x</i>
	set		show changed options
Most useful options and their abbreviations	set all		show all options
	set x?		show value of option <i>x</i>
	autoindent	ai	supply indent
	autowrite	aw	write before changing files
	directory		pathname of directory for temporary work files
	exrc	ex	allow vi/ex to read the .exrc in the current directory. This option is set in the EXINIT shell variable or in the .exrc file in the \$HOME directory.
	ignorecase	ic	ignore case of letters in scanning
	list		print ^I for tab, \$ at end
	magic		treat . [* special in patterns
	modelines		first five lines and last five lines executed as vi/ex commands if they are of the form ex:command: or vi:command:
	number	nu	number lines
	paragraphs	para	macro names that start paragraphs
	redraw		simulate smart terminal
	report		informs you if the number of lines modified by the last command is greater than the value of the report variable
	scroll		command mode lines
	sections	sect	macro names that start sections
	shiftwidth	sw	for < , > , and input ^D
	showmatch	sm	to) and } as typed
	showmode	smd	show insert mode in vi
	slowopen	slow	stop updates during insert
term		specifies to vi the type of terminal being used (the default is the value of the environment variable TERM)	
window		visual mode lines	
wrapmargin	wm	automatic line splitting	
wrapsan	ws	search around end (or beginning) of buffer	

Scanning pattern formation

<code>^</code>	beginning of line
<code>\$</code>	end of line
<code>.</code>	any character
<code>\<</code>	beginning of word
<code>\></code>	end of word
<code>[str]</code>	any character in <i>str</i>
<code>[^str]</code>	any character not in <i>str</i>
<code>[x-y]</code>	any character between <i>x</i> and <i>y</i>
<code>*</code>	any number of preceding characters

AUTHOR

vi and **ex** are based on software developed by The University of California, Berkeley California, Computer Science Division, Department of Electrical Engineering and Computer Science.

FILES

<code>/tmp/Exnnnnn</code>	editor temporary
<code>/tmp/Rxnnnnn</code>	named buffer temporary
<code>/usr/lib/exprserve</code>	preserve command
<code>/usr/lib/exrecover</code>	recover command
<code>/usr/lib/exstrings</code>	error messages
<code>/usr/share/lib/terminfo/*</code>	describes capabilities of terminals
<code>/var/preserve/login</code>	preservation directory (where login is the user's login)
<code>\$HOME/.exrc</code>	editor startup file
<code>./exrc</code>	editor startup file

SEE ALSO

ed(1), **edit(1)**, **grep(1)**, **sed(1)**, **sort(1)**, **vi(1)**, **curses(3X)**, **term(4)**, **terminfo(4)**

Solaris Advanced User's Guide

NOTES

Several options, although they continue to be supported, have been replaced in the documentation by options that follow the Command Syntax Standard (see **intro(1)**). The `-` option has been replaced by `-s`, a `-r` option that is not followed with an option-argument has been replaced by `-L`, and `+command` has been replaced by `-c command`.

The message **file too large to recover with -r option**, which is seen when a file is loaded, indicates that the file can be edited and saved successfully, but if the editing session is lost, recovery of the file with the `-r` option will not be possible.

The encryption options and commands are provided with the Security Administration Utilities package, which is available only in the United States.

The `z` command prints the number of logical rather than physical lines. More than a screen full of output may result if long lines are present.

File input/output errors do not print a name if the command line `-s` option is used.

The editing environment defaults to certain configuration options. When an editing session is initiated, **ex** attempts to read the **EXINIT** environment variable. If it exists, the editor uses the values defined in **EXINIT**, otherwise the values set in **\$HOME/.exrc** are used. If **\$HOME/.exrc** does not exist, the default values are used.

To use a copy of **.exrc** located in the current directory other than **\$HOME**, set the *exrc* option in **EXINIT** or **\$HOME/.exrc**. Options set in **EXINIT** can be turned off in a local **.exrc** only if *exrc* is set in **EXINIT** or **\$HOME/.exrc**.

There is no easy way to do a single scan ignoring case.

The editor does not warn if text is placed in named buffers and not used before exiting the editor.

Null characters are discarded in input files and cannot appear in resultant files.

NAME	exec, eval, source – shell built-in functions to execute other commands
SYNOPSIS	
sh	exec [<i>argument...</i>] eval [<i>argument...</i>]
csh	exec <i>command</i> eval <i>argument...</i> source [-h] <i>name</i>
ksh	† exec [<i>arg...</i>] † eval [<i>arg...</i>]
DESCRIPTION	
sh	The exec command specified by the arguments is executed in place of this shell without creating a new process. Input/output arguments may appear and, if no other arguments are given, cause the shell input/output to be modified. The <i>arguments</i> to the eval built-in are read as input to the shell and the resulting command(s) executed.
csh	exec executes <i>command</i> in place of the current shell, which terminates. eval reads its <i>arguments</i> as input to the shell and executes the resulting command(s). This is usually used to execute commands generated as the result of command or variable substitution, since parsing occurs before these substitutions. source reads commands from <i>name</i> . source commands may be nested, but if they are nested too deeply the shell may run out of file descriptors. An error in a sourced file at any level terminates all nested source commands. -h Place commands from the file <i>name</i> on the history list without executing them.
ksh	With the exec built-in, if <i>arg</i> is given, the command specified by the arguments is executed in place of this shell without creating a new process. Input/output arguments may appear and affect the current process. If no arguments are given the effect of this command is to modify file descriptors as prescribed by the input/output redirection list. In this case, any file descriptor numbers greater than 2 that are opened with this mechanism are closed when invoking another program. The arguments to eval are read as input to the shell and the resulting command(s) executed. On this man page, ksh (1) commands that are preceded by one or two † (daggers) are treated specially in the following ways: 1. Variable assignment lists preceding the command remain in effect when the command completes. 2. I/O redirections are processed after variable assignments.

3. Errors cause a script that contains them to abort.
4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.

SEE ALSO**cs****h**(1), **ks****h**(1), **sh**(1)

NAME	exit, return, goto – shell built-in functions to enable the execution of the shell to advance beyond its sequence of steps
SYNOPSIS	
sh	exit [<i>n</i>]
csh	exit [(<i>expr</i>)] return [<i>n</i>] goto <i>label</i>
ksh	† exit [<i>n</i>] † return [<i>n</i>]
DESCRIPTION	
sh	exit will cause a shell to exit with the exit status specified by <i>n</i> . If <i>n</i> is omitted the exit status is that of the last command executed (an EOF will also cause the shell to exit.) return causes a function to exit with the return value specified by <i>n</i> . If <i>n</i> is omitted, the return status is that of the last command executed.
csh	exit will cause the shell to exit, either with the value of the status variable or with the value specified by the expression <i>expr</i> . The goto built-in uses a specified <i>label</i> as a search string amongst commands. The shell rewinds its input as much as possible and searches for a line of the form <i>label</i> : possibly preceded by space or tab characters. Execution continues after the indicated line. It is an error to jump to a label that occurs between a while or for built-in command and its corresponding end .
ksh	exit will cause the shell to exit with the exit status specified by <i>n</i> . The value will be the least significant 8 bits of the specified status. If <i>n</i> is omitted then the exit status is that of the last command executed. When exit occurs when executing a trap, the last command refers to the command that executed before the trap was invoked. An end-of-file will also cause the shell to exit except for a shell which has the ignoreeof option (See set below) turned on. return causes a shell function or '.' script to return to the invoking script with the return status specified by <i>n</i> . The value will be the least significant 8 bits of the specified status. If <i>n</i> is omitted then the return status is that of the last command executed. If return is invoked while not in a function or a '.' script, then it is the same as an exit . On this man page, ksh(1) commands that are preceded by one or two † (daggers) are treated specially in the following ways: <ol style="list-style-type: none"> 1. Variable assignment lists preceding the command remain in effect when the command completes. 2. I/O redirections are processed after variable assignments. 3. Errors cause a script that contains them to abort.

4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.

SEE ALSO **break(1), csh(1), ksh(1), sh(1)**

NAME	expand, unexpand – expand TAB characters to SPACE characters, and vice versa
SYNOPSIS	expand [<i>-tabstop</i>] [<i>-tab1, tab2, . . . , tabn</i>] [<i>filename. . .</i>] unexpand [<i>-a</i>] [<i>filename. . .</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	expand copies <i>filenames</i> (or the standard input) to the standard output, with TAB characters expanded to SPACE characters. BACKSPACE characters are preserved into the output and decrement the column count for TAB calculations. expand is useful for pre-processing character files (before sorting, looking at specific columns, and so forth) that contain TAB characters. unexpand copies <i>filenames</i> (or the standard input) to the standard output, putting TAB characters back into the data. By default, only leading SPACE and TAB characters are converted to strings of tabs, but this can be overridden by the -a option (see the OPTIONS section below).
OPTIONS	expand options are: <i>-tabstop</i> Specify as a single argument, sets TAB characters <i>tabstop</i> SPACE characters apart instead of the default 8. <i>-tab1, tab2, . . . , tabn</i> Set TAB characters at the columns specified by <i>tab1. . .</i> unexpand options are: -a Insert TAB characters when replacing a run of two or more SPACE characters would produce a smaller output file.
ENVIRONMENT	LC_CTYPE determines how expand handles characters. When LC_CTYPE is set to a valid value, expand can display and handle text and filenames containing valid characters for that locale. expand can display and handle Extended Unix Code (EUC) characters where any character can be 1, 2, or 3 bytes wide. expand can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.

NAME	exportfs – translates exportfs options to share/unshare commands
SYNOPSIS	<code>/usr/sbin/exportfs [-aiuv] [-o options] [pathname]</code>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>exportfs translates SunOS 4.x exportfs options to the corresponding share/unshare options and invokes share/unshare with the translated options.</p> <p>With no options or arguments, exportfs invokes share to print out the list of all currently shared NFS filesystems.</p> <p>exportfs is the BSD/Compatibility Package command of share(1M) and unshare(1M). Use share(1M)/ unshare(1M) whenever possible.</p>
OPTIONS	<p>-a Invokes shareall(1M), or if -u is specified, invokes unshareall(1M).</p> <p>-i Ignore options in /etc/dfs/dfstab.</p> <p>-u Invokes unshare(1M) on <i>pathname</i>.</p> <p>-v Verbose.</p> <p>-o options Specify a comma-separated list of optional characteristics for the filesystems being exported. exportfs translates <i>options</i> to share-equivalent options. (see share(1M) for information about individual options).</p>
SEE ALSO	share(1M) , shareall(1M) , unshare(1M) , unshareall(1M)

NAME	expr – evaluate arguments as an expression
SYNOPSIS	expr <i>arguments</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	The <i>arguments</i> are taken as an expression. After evaluation, the result is written on the standard output. Terms of the expression must be separated by blanks. Characters special to the shell must be escaped. Note that 0 is returned to indicate a zero value, rather than the null string. Strings containing blanks or other special characters should be quoted. Integer-valued arguments may be preceded by a unary minus sign. Internally, integers are treated as 32-bit, 2s complement numbers. The length of the expression is limited to 512 characters.
USAGE	<p>The operators and keywords are listed below. Characters that need to be escaped in the shell (see sh(1)) are preceded by \. The list is in order of increasing precedence, with equal precedence operators grouped within {} symbols.</p> <p><i>expr \ expr</i> returns the first <i>expr</i> if it is neither null nor 0, otherwise returns the second <i>expr</i>.</p> <p><i>expr \& expr</i> returns the first <i>expr</i> if neither <i>expr</i> is null or 0, otherwise returns 0.</p> <p><i>expr { =, \>, \>=, \<, \<=, != } expr</i> returns the result of an integer comparison if both arguments are integers, otherwise returns the result of a lexical comparison.</p> <p><i>expr { +, - } expr</i> addition or subtraction of integer-valued arguments.</p> <p><i>expr { *, /, % } expr</i> multiplication, division, or remainder of the integer-valued arguments.</p> <p><i>expr : expr</i> The matching operator : compares the first argument with the second argument, which must be a regular expression. Regular expression syntax is the same as that of ed(1), except that all patterns are “anchored” (that is, begin with ^) and, therefore, ^ is not a special character, in that context. Normally, the matching operator returns the number of bytes matched (0 on failure). Alternatively, the \(...\) pattern symbols can be used to return a portion of the first argument.</p>
EXAMPLES	<p>Add 1 to the shell variable a:</p> <pre>a=expr \$a + 1</pre>

The following example emulates **basename(1)** — it returns the last segment of the path name **\$a**. For **\$a** equal to either **/usr/abc/file** or just **file**, the example returns **file**. (Watch out for **/** alone as an argument: **expr** takes it as the division operator; see the NOTES below.)

```
expr $a : ^.*\/(.*\)^ \ | $a
```

Here is a better version of the previous example. The addition of the **//** characters eliminates any ambiguity about the division operator and simplifies the whole expression.

```
expr // $a : ^.*\/(.*\)^
```

Return the number of characters in **\$VAR**:

```
expr $VAR : ^.*^
```

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

ed(1), **sh(1)**, **environ(5)**

DIAGNOSTICS

As a side effect of expression evaluation, **expr** returns the following exit values:

- 0 if the expression is neither null nor 0
- 1 if the expression is null or 0
- 2 for invalid expressions.

syntax error for operator/operand errors

non-numeric argument if arithmetic is attempted on such a string

NOTES

After argument processing by the shell, **expr** cannot tell the difference between an operator and an operand except by the value. If **\$a** is an =, the command:

```
expr $a = '= '
```

looks like:

```
expr = = =
```

as the arguments are passed to **expr** (and they are all taken as the = operator). The following works:

```
expr X$a = X=
```

NAME	expr – evaluate arguments as a logical, arithmetic, or string expression
SYNOPSIS	<i>/usr/ucb/expr argument. . .</i>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>expr evaluates expressions as specified by its arguments. After evaluation, the result is written on the standard output. Each token of the expression is a separate argument, so terms of the expression must be separated by blanks. Characters special to the shell must be escaped. Note: 0 is returned to indicate a zero value, rather than the null string. Strings containing blanks or other special characters should be quoted. Integer-valued arguments may be preceded by a unary minus sign. Internally, integers are treated as 32-bit, two's-complement numbers.</p> <p>The operators and keywords are listed below. Characters that need to be escaped are preceded by '\'. The list is in order of increasing precedence, with equal precedence operators grouped within {} symbols.</p> <p><i>expr \ expr</i> Return the first <i>expr</i> if it is neither NULL nor 0, otherwise returns the second <i>expr</i>.</p> <p><i>expr \ & expr</i> Return the first <i>expr</i> if neither <i>expr</i> is NULL or 0, otherwise returns 0.</p> <p><i>expr { =, \>, \>=, \<, \<=, != } expr</i> Return the result of an integer comparison if both arguments are integers, otherwise returns the result of a lexical comparison.</p> <p><i>expr { +, - } expr</i> Addition or subtraction of integer-valued arguments.</p> <p><i>expr { *, /, % } expr</i> Multiplication, division, or remainder of the integer-valued arguments.</p> <p><i>string : regular-expression</i> match <i>string regular-expression</i></p> <p>The two forms of the matching operator above are synonymous. The matching operators : and match compare the first argument with the second argument which must be a regular expression. Regular expression syntax is the same as that of ed(1), except that all patterns are “anchored” (treated as if they begin with ^) and, therefore, ^ is not a special character, in that context. Normally, the matching operator returns the number of characters matched (0 on failure). Alternatively, the \(...\) pattern symbols can be used to return a portion of the first argument.</p>

substr *string integer-1 integer-2*

Extract the substring of *string* starting at position *integer-1* and of length *integer-2* characters. If *integer-1* has a value greater than the length of *string*, **expr** returns a null string. If you try to extract more characters than there are in *string*, **expr** returns all the remaining characters from *string*. Beware of using negative values for either *integer-1* or *integer-2* as **expr** tends to run forever in these cases.

index *string character-list*

Report the first position in *string* at which any one of the characters in *character-list* matches a character in *string*.

length *string*

Return the length (that is, the number of characters) of *string*.

(**expr**) Parentheses may be used for grouping.

EXAMPLES

1. **a='expr \$a + 1'**

Adds 1 to the shell variable **a**.

2. **# 'For \$a equal to either "/usr/abc/file" or just "file"'**
expr \$a : '.*\/(.*)' \| \$a

Returns the last segment of a path name (that is, the filename part). Watch out for / alone as an argument: *expr* will take it as the division operator (see BUGS below).

3. **# A better representation of example 2.**
expr // \$a : '.*\/(.*)'

The addition of the // characters eliminates any ambiguity about the division operator and simplifies the whole expression.

4. **expr \$VAR : '.*'**

Returns the number of characters in **\$VAR**.

SEE ALSO

ed(1), **sh(1)**, **test(1)**

EXIT CODES

expr returns the following exit codes:

- 0 if the expression is neither null nor 0
- 1 if the expression is null or 0
- 2 for invalid expressions.

DIAGNOSTICS

syntax error for operator/operand errors

non-numeric argument

if arithmetic is attempted on such a string

division by zero

if an attempt to divide by zero is made

BUGS

After argument processing by the shell, **expr** cannot tell the difference between an operator and an operand except by the value. If **\$a** is an =, the command:

```
expr $a = '='
```

looks like:

```
expr = = =
```

as the arguments are passed to **expr** (and they will all be taken as the = operator). The following works:

```
expr X$a = X=
```

Note: the **match**, **substr**, **length**, and **index** operators cannot themselves be used as ordinary strings. That is, the expression:

```
example% expr index expurgatorious length
```

```
syntax error
```

```
example%
```

generates the '**syntax error**' message as shown instead of the value **1** as you might expect.

NAME	exstr – extract strings from source files
SYNOPSIS	<pre> exstr <i>filename</i> ... exstr -e <i>filename</i> ... exstr -r [-d] <i>filename</i> ... </pre>
DESCRIPTION	<p>The exstr utility is used to extract strings from C-language source files and replace them by calls to the message retrieval function (see gettext(3C)). This utility will extract all character strings surrounded by double quotes, not just strings used as arguments to the printf command or the printf routine. In the first form, exstr finds all strings in the source files and writes them on the standard output. Each string is preceded by the source file name and a colon.</p> <p>The first step is to use exstr -e to extract a list of strings and save it in a file. Next, examine this list and determine which strings can be translated and subsequently retrieved by the message retrieval function. Then, modify this file by deleting lines that can't be translated and, for lines that can be translated, by adding the message file names and the message numbers as the fourth (<i>msgfile</i>) and fifth (<i>msgnum</i>) entries on a line. The message files named must have been created by mkmsgs(1) and exist in <i>/usr/lib/locale/locale/LC_MESSAGES</i>. (The directory <i>locale</i> corresponds to the language in which the text strings are written; see setlocale(3C)). The message numbers used must correspond to the sequence numbers of strings in the message files.</p> <p>Now use this modified file as input to exstr -r to produce a new version of the original C-language source file in which the strings have been replaced by calls to the message retrieval function gettext(). The <i>msgfile</i> and <i>msgnum</i> fields are used to construct the first argument to gettext(). The second argument to gettext() is printed if the message retrieval fails at run-time. This argument is the null string, unless the -d option is used.</p> <p>This utility cannot replace strings in all instances. For example, a static initialized character string cannot be replaced by a function call. A second example is that a string could be in a form of an escape sequence which could not be translated. In order not to break existing code, the files created by invoking exstr -e must be examined and lines containing strings not replaceable by function calls must be deleted. In some cases the code may require modifications so that strings can be extracted and replaced by calls to the message retrieval function.</p>
OPTIONS	<p>-e Extract a list of strings from the named C-language source files, with positional information. This list is produced on standard output in the following format:</p> <pre> <i>file:line:position:msgfile:msgnum:string</i> </pre> <p> <i>file</i> the name of a C-language source file <i>line</i> line number in the file <i>position</i> character position in the line <i>msgfile</i> null <i>msgnum</i> null <i>string</i> the extracted text string </p>

Normally you would redirect this output into a file. Then you would edit this file to add the values you want to use for *msgfile* and *msgnum*:

msgfile the file that contains the text strings that will replace *string*. A file with this name must be created and installed in the appropriate place by the **mkmsgs(1)** utility.

msgnum the sequence number of the string in *msgfile*.

The next step is to use **exstr -r** to replace *strings* in *file*.

- r** Replace strings in a C-language source file with function calls to the message retrieval function **gettext()**.
- d** This option is used together with the **-r** option. If the message retrieval fails when **gettext()** is invoked at run-time, then the extracted string is printed. You would use the capability provided by **exstr** on an application program that needs to run in an international environment and have messages print in more than one language. **exstr** replaces text strings with function calls that point at strings in a message data base. The data base used depends on the run-time value of the **LC_MESSAGES** environment variable (see **environ(5)**).

EXAMPLES

The following examples show uses of **exstr**.

Assume that the file **example.c** contains two strings:

```
main()
{
    printf("This is an example\n");
    printf("Hello world!\n");
}
```

The **exstr** utility, invoked with the argument **example.c** extracts strings from the named file and prints them on the standard output.

```
example% exstr example.c
```

produces the following output:

```
example.c:This is an example\n
example.c:Hello world!\n
```

```
example% exstr -e example.c > example.stringsout
```

produces the following output in the file **example.stringsout**:

```
example.c:3:8::This is an example\n
example.c:4:8::Hello world!\n
```

You must edit **example.stringsout** to add the values you want to use for the *msgfile* and *msgnum* fields before these strings can be replaced by calls to the retrieval function. If **UX** is the name of the message file, and the numbers **1** and **2** represent the sequence number of the strings in the file, here is what **example.stringsout** looks like after you add this information:

```
example.c:3:8:UX:1:This is an example\n
example.c:4:8:UX:2:Hello world!\n
```

The **exstr** utility can now be invoked with the **-r** option to replace the strings in the source file by calls to the message retrieval function **gettext()**.

```
example% exstr -r example.c <example.stringsout >intexample.c
```

produces the following output:

```
extern char *gettext();
main()
{
    printf(gettext("UX:1", ""));
    printf(gettext("UX:2", ""));
}
```

```
example% exstr -rd example.c <example.stringsout >intexample.c
```

uses the extracted strings as a second argument to **gettext()**.

```
extern char *gettext();
main()
{
    printf(gettext("UX:1", "This is an example\n"));
    printf(gettext("UX:2", "Hello world!\n"));
}
```

FILES /usr/lib/locale/locale/LC_MESSAGES/* files created by **mkmsgs(1)**

SEE ALSO **gettext(1)**, **mkmsgs(1)**, **printf(1)**, **srchtxt(1)**, **gettext(3C)**, **printf(3S)**, **setlocale(3C)**, **environ(5)**

DIAGNOSTICS The error messages produced by **exstr** are intended to be self-explanatory. They indicate errors in the command line or format errors encountered within the input file.

NAME	face – executable for the Framed Access Command Environment Interface
SYNOPSIS	face [-i <i>init_file</i>] [-c <i>command_file</i>] [-a <i>alias_file</i>] [<i>filename</i> . . .]
DESCRIPTION	<i>filename</i> is the full pathname of the file describing the object to be opened initially, and must follow the naming convention Menu.xxx for a menu, Form.xxx for a form, and Text.xxx for a text file, where <i>xxx</i> is any string that conforms to the UNIX system file naming conventions. The Form and Menu Language Interpreter (FMLI) descriptor lifetime will be ignored for all frames opened by argument to face . These frames have a lifetime of immortal by default. If <i>filename</i> is not specified on the command line, the AT&T FACE Menu will be opened along with those objects specified by the LOGINWIN environment variables. These variables are found in the user's .environ file.
OPTIONS	-a <i>alias_file</i> Alias file. -c <i>command_file</i> Command file. -i <i>init_file</i> Initial file.
FILES	\$HOME/pref/.environ
SEE ALSO	env(1)
DIAGNOSTICS	The face command will exit with a non-zero exit code if the user is not properly set up as a FACE user.

NAME	factor – obtain the prime factors of a number
SYNOPSIS	factor [<i>integer</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	<p>When you use factor without an argument, it waits for you to give it an integer. After you give it a positive integer less than or equal to 10^{14}, it factors the integer, prints its prime factors the proper number of times, and then waits for another integer. factor exits if it encounters a zero or any non-numeric character.</p> <p>If you invoke factor with an argument, it factors the integer as described above, and then it exits.</p> <p>The maximum time to factor an integer is proportional to \sqrt{n}. factor will take this time when n is prime or the square of a prime.</p>
DIAGNOSTICS	factor prints the error message, Ouch , for input out of range or for garbage input.

NAME	fastboot, fasthalt – reboot/halt the system without checking the disks
SYNOPSIS	<i>/usr/ucb/fastboot</i> [<i>boot-options</i>] <i>/usr/ucb/fasthalt</i> [<i>halt-options</i>]
AVAILABILITY	SUNWscpu
DESCRIPTION	fastboot and fasthalt are shell scripts that invoke reboot and halt with the proper arguments. These commands are provided for compatibility only.
SEE ALSO	fsck(1M) , halt(1M) , init(1M) , reboot(1M) , init.d(4)

NAME	fdformat – format floppy diskette
SYNOPSIS	fdformat [-dDeEfHILmMUqvx] [-b label] [-B filename] [-t dostype] [<i>devname</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>fdformat is a utility for formatting diskettes. All new, blank diskettes must be formatted before they can be used. fdformat formats and verifies each track on the diskette, and indicates whether any bad sectors were encountered. All existing data on the diskette, if any, is destroyed by formatting.</p> <p>By default, fdformat uses the configured capacity of the drive to format the diskette. A 3.5" high-density drive uses diskettes with a formatted capacity of 1.44 megabytes. A 5.25" high-density drive uses diskettes with a formatted capacity of 1.2 megabytes. In either case, a density option does not have to be specified to fdformat. However, a density option must be specified when using a diskette with a lower capacity than the drive's default. Use the -H option to format high-density diskettes (1.44-megabyte capacity) in an extra-high-density (ED) drive. Use the -D option, the -I option, or the -L option to format double-density (or "low-density") diskettes (720KB capacity) in an HD or ED drive. To format medium-density diskettes (1.2-megabyte capacity), use the -M option with -t nec (this is the same as using the -m option with -t nec).</p> <p>Extended density uses double-sided, extended-density (or extra-high-density) (DS/ED) diskettes. Medium and high densities use the same media: double-sided, high-density (DS/HD) diskettes. Double ("low") density uses double-sided, double-density (DS/DD) diskettes. Substituting diskettes of one density for diskettes of either a higher or lower density generally will not work. Data integrity cannot be assured whenever a diskette is formatted to a capacity not matching its density.</p> <p>fdformat writes new identification and data fields for each sector on all tracks unless the -x option is specified. If the -v option is specified, each sector on all tracks is written and read back. Otherwise, a read of each sector is performed immediately after the track is formatted.</p> <p>After formatting and verifying, fdformat writes an operating-system label on block 0. Use the -t dos option (same as the -d option) to put an MS-DOS file system on the diskette after format is done or use the -t nec option with the -M option (same as the -m option) to put an NEC-DOS file system on the diskette. Otherwise, fdformat writes a SunOS label on the diskette in block 0. The label is required on Solaris systems if you intend to put a UNIX file system on the diskette.</p>
OPTIONS	<p>-D Format a 720KB (3.5") or 360KB (5.25") double-density diskette (same as the -I or -L options). This is the default for double-density type drives. It is needed if the drive is a high- or extended-density type.</p> <p>-e Eject the diskette when done. (This feature is not available on all systems).</p> <p>-E Format a 2.88-megabyte (3.5") extended-density diskette. This is the default for extended-density type drives.</p>

- f** Force. Do not ask for confirmation before starting format.
 - H** Format a 1.44-megabyte (3.5") or 1.2-megabyte (5.25") high-density diskette. This is the default for high-density type drives; it is needed if the drive is the extended-density type.
 - M** Write a 1.2-megabyte (3.5") medium-density format on a high-density diskette (use only with the **-t nec** option). This is the same as using **-m**. (This feature is not available on all systems.)
 - U** **umount** any file systems on the floppy and then format.
 - q** Quiet; do not print status messages.
 - v** Verify each block of the diskette after the format.
 - x** Skip the format, and only write a SunOS label or an MS-DOS file system.
 - b label** Label the diskette with volume *label*. A SunOS volume label is restricted to 8 characters. A DOS volume label is restricted to 11 upper-case characters.
 - B filename** Install special boot loader in *filename* on an MS-DOS diskette. This option is only meaningful when the **-d option** (or **-t dos**) is also specified.
 - t dos** Install an MS-DOS file system and boot sector on the disk after formatting. This is equivalent to the DOS format command or the **-d** option.
 - t nec** Install an NEC-DOS file system and boot sector on the disk after formatting. This should be used only with the **-M** option. (This feature is not available on all systems.)
 - devname* Replace *devname* with **rdiskette0** (systems without Volume Management) or **floppy0** (systems with Volume Management) to use the first drive or **rdiskette1** (systems without Volume Management) or **floppy1** (systems with Volume Management) to use the second drive. If *devname* is omitted, the default drive, if one exists, will be used.
- The following options are provided for compatibility with previous versions of **fdformat**; their use is discouraged.
- d** Format an MS-DOS floppy diskette (same as **-t dos**). This is equivalent to the MS-DOS **FORMAT** command.
 - I** Format a 720KB (3.5") or 360KB (5.25") double-density diskette (same as **-D** or **-L**). This is the default for double-density type drives; it is needed if the drive is the high- or extended-density type.
 - L** Format a 720KB (3.5") or 360KB (5.25") double-density diskette (same as **-I** or **-D**). This is the default for double-density type drives; it is needed if the drive is the high- or extended-density type.
 - m** Write a 1.2-megabyte (3.5") medium-density format on a high-density diskette (use only with the **-t nec** option). This is the same as using **-M**. (This feature is not available on all systems.)

FILES	<p>/vol/dev/diskette0 Directory providing block device access for the media in floppy drive 0.</p> <p>/vol/dev/rdiskette0 Directory providing character device access for the media in floppy drive 0.</p> <p>/vol/dev/aliases/floppy0 Symbolic link to the character device for the media in floppy drive 0.</p> <p>/dev/rdiskette Directory providing character device access for the media in the primary floppy drive, usually drive 0.</p>
SEE ALSO	cpio(1), eject(1), tar(1), volcancel(1), volcheck(1), volmissing(1), mount(1M), newfs(1M), rmmount(1M), vold(1M), rmmount.conf(4), vold.conf(4), pcfs(7), volfs(7)
x86 Only	fd(7)
NOTES	<p>A diskette containing a ufs file system created on a SPARC (big endian) system (by using fdformat and newfs(1M)) is not identical to a diskette containing a ufs file system created on an x86 (little endian) system. Do not interchange ufs diskettes between these platforms; use cpio(1) or tar(1) to transfer files on diskettes between them.</p> <p>A diskette formatted using the -t dos option (or -d) for MS-DOS will not have the necessary system files, and is therefore not bootable. Trying to boot from it on a PC will result in the following message:</p> <p style="text-align: center;">Non-System disk or disk error Replace and strike any key when ready</p>
BUGS	Currently, bad sector mapping is not supported on floppy diskettes. Therefore, a diskette is unusable if fdformat finds an error (bad sector).

NAME	fgrep – search a file for a character string
SYNOPSIS	fgrep [-bchilnsvx] [-e <i>special string</i>] [-f <i>filename</i>] [<i>strings</i>] [<i>filename ...</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>fgrep (fast grep) searches files for a character string and prints all lines that contain that string. fgrep is different from grep(1) and egrep(1) because it searches for a string, instead of searching for a pattern that matches an expression. It uses a fast and compact algorithm.</p> <p>The characters \$, *, [, ^, , (,), and \ are interpreted literally by fgrep, that is, fgrep does not recognize full regular expressions as does egrep. Since these characters have special meaning to the shell, it is safest to enclose the entire <i>string</i> in single quotes '...'. If no files are specified, fgrep assumes standard input. Normally, each line found is copied to the standard output. The file name is printed before each line found if there is more than one input file.</p>
OPTIONS	<p>-b Precede each line by the block number on which it was found. This can be useful in locating block numbers by context (first block is 0).</p> <p>-c Print only a count of the lines that contain the pattern.</p> <p>-h Suppress printing of filenames when searching multiple files.</p> <p>-i Ignore upper/lower case distinction during comparisons.</p> <p>-l Print the names of files with matching lines once, separated by new-lines. Does not repeat the names of files when the pattern is found more than once.</p> <p>-n Precede each line by its line number in the file (first line is 1).</p> <p>-s Work silently, that is, display nothing except error messages. This is useful for checking the error status.</p> <p>-v Print all lines except those that contain the pattern.</p> <p>-x Print only lines matched entirely.</p> <p>-e <i>special_string</i> Search for a <i>special string</i> (<i>string</i> begins with a -).</p> <p>-f <i>filename</i> Take the list of <i>strings</i> from <i>filename</i>.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of fgrep for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how fgrep behaves.</p>

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S./English).

SEE ALSO

ed(1), egrep(1), grep(1), sed(1), sh(1), environ(5)

DIAGNOSTICS

Exit status is 0 if any matches are found, 1 if none, 2 for syntax errors or inaccessible files (even if matches were found).

NOTES

Ideally there should be only one **grep** command, but there is not a single algorithm that spans a wide enough range of space-time tradeoffs. Lines are limited to BUFSIZ characters; longer lines are truncated. BUFSIZ is defined in `<stdio.h>`.

NAME	file – determine file type
SYNOPSIS	file [-h] [-m <i>mfilename</i>] [-f <i>ffilename</i>] <i>argument</i> ... file [-h] [-m <i>mfilename</i>] -f <i>ffilename</i> file -c [-m <i>mfilename</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>file performs a series of tests on each file supplied by <i>argument</i> and, optionally, on each file supplied in <i>ffilename</i> in an attempt to classify it. If <i>argument</i> appears to be a text file, file examines the first 512 bytes and tries to guess its programming language. If <i>argument</i> is an executable a.out, file prints the version stamp, provided it is greater than 0. If <i>argument</i> is a symbolic link, by default the link is followed and file tests the file that the symbolic link references.</p> <p>file uses /etc/magic to identify files that have a magic number. A magic number is a numeric or string constant that indicates the file type. Commentary at the beginning of /etc/magic explains its format.</p>
OPTIONS	<p>-c Check the magic file for format errors. For reasons of efficiency, this validation is normally not carried out.</p> <p>-h Do not follow symbolic links.</p> <p>-f <i>ffilename</i> <i>ffilename</i> contains the names of the files to be examined.</p> <p>-m <i>mfilename</i> Use <i>mfilename</i> as an alternate magic file, instead of /etc/magic .</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of file for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how file behaves.</p> <p>LC_CTYPE Determines how file handles characters. When LC_CTYPE is set to a valid value, file can display and handle text and filenames containing valid characters for that locale. file can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. file can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p> <p>LC_MESSAGES Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).</p>

FILES | */etc/magic*

SEE ALSO | *filehdr(4), magic(4), environ(5),*

DIAGNOSTICS | If the **-h** option is specified and *argument* is a symbolic link, **file** prints the error message:
symbolic link to *argument*

NAME	file – determine the type of a file by examining its contents
SYNOPSIS	<code>/usr/ucb/file [-f <i>ffile</i>] [-cL] [-m <i>mfile</i>] <i>filename</i> . .</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>file performs a series of tests on each <i>filename</i> in an attempt to determine what it contains. If the contents of a file appear to be ASCII text, file examines the first 512 bytes and tries to guess its language.</p> <p>file uses the file <code>/etc/magic</code> to identify files that have some sort of <i>magic number</i>, that is, any file containing a numeric or string constant that indicates its type.</p>
OPTIONS	<p><code>-c</code> Check for format errors in the magic number file. For reasons of efficiency, this validation is not normally carried out. No file type-checking is done under <code>-c</code>.</p> <p><code>-f <i>ffile</i></code> Get a list of filenames to identify from <i>ffile</i>.</p> <p><code>-L</code> If a file is a symbolic link, test the file the link references rather than the link itself.</p> <p><code>-m <i>mfile</i></code> Use <i>mfile</i> as the name of an alternate magic number file.</p>
EXAMPLES	<p>This example illustrates the use of file on all the files in a specific user's directory:</p> <pre> example% pwd /usr/blort/misc example% /usr/ucb/file * code: mc68020 demand paged executable code.c: c program text counts: ascii text doc: roff, nroff , or eqn input text empty.file: empty libz: archive random library memos: directory project: symbolic link to /usr/project script: executable shell script titles: ascii text s5.stuff: cpio archive example%</pre>

ENVIRONMENT

The environment variables **LC_CTYPE**, **LANG**, and **LC_default** control the character classification throughout **file**. On entry to **file**, these environment variables are checked in the following order: **LC_CTYPE**, **LANG**, and **LC_default**. When a valid value is found, remaining environment variables for character classification are ignored. For example, a new setting for **LANG** does not override the current valid character classification rules of **LC_CTYPE**. When none of the values is valid, the shell character classification defaults to the POSIX.1 "C" locale.

FILES

/etc/magic

SEE ALSO

magic(4)

BUGS

file often makes mistakes. In particular, it often suggests that command files are C programs.

file does not recognize Pascal or LISP.

NAME	find – find files
SYNOPSIS	find <i>path-name-list expression</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	find recursively descends the directory hierarchy for each path name in the <i>path-name-list</i> (that is, one or more path names) seeking files that match a Boolean <i>expression</i> written in the primaries given below. In the descriptions, the argument <i>n</i> is used as a decimal integer where <i>+n</i> means more than <i>n</i> , <i>-n</i> means less than <i>n</i> and <i>n</i> means exactly <i>n</i> .
USAGE Expressions	Valid expressions are: <ul style="list-style-type: none"> -name <i>pattern</i> True if <i>pattern</i> matches the current file name. Normal shell file name generation characters (see sh(1)) may be used. A backslash (\) is used as an escape character within the pattern. The pattern should be escaped or quoted when find is invoked from the shell. -perm [-]<i>onum</i> True if the file permission flags exactly match the octal number <i>onum</i> (see chmod(1)). If <i>onum</i> is prefixed by a minus sign (-), only the bits that are set in <i>onum</i> are compared with the file permission flags, and the expression evaluates true if they match. -size <i>n</i>[<i>c</i>] True if the file is <i>n</i> blocks long (512 bytes per block). If <i>n</i> is followed by a <i>c</i>, the size is in characters. -atime <i>n</i> True if the file was accessed <i>n</i> days ago. The access time of directories in <i>path-name-list</i> is changed by find itself. -ls Always true; prints current pathname together with its associated statistics. These include (respectively) inode number, size in kilobytes (1024 bytes), protection mode, number of hard links, user, group, size in bytes, and modification time. If the file is a special file the size field will instead contain the major and minor device numbers. If the file is a symbolic link the pathname of the linked-to file is printed preceded by '→'. The format is identical to that of ls -gilds (see ls(1)). Note: Formatting is done internally, without executing the ls program. -mtime <i>n</i> True if the file's data was modified <i>n</i> days ago. -ctime <i>n</i> True if the file's status was changed <i>n</i> days ago. -exec <i>command</i> True if the executed <i>command</i> returns a zero value as exit status. The end of <i>command</i> must be punctuated by an escaped semicolon. A command argument {} is replaced by the current path name. -ok <i>command</i> Like -exec except that the generated command line is printed with a question mark first, and is executed only if the user responds by typing y. -cpio <i>device</i> Always true; write the current file on <i>device</i> in cpio format (5120-byte

	records).
-ncpio <i>device</i>	Always true; write the current file on <i>device</i> in cpio -c format (5120 byte records).
-print	Always true; causes the current path name to be printed.
-newer <i>filename</i>	True if the current file has been modified more recently than the argument <i>filename</i> .
-depth	Always true; causes descent of the directory hierarchy to be done so that all entries in a directory are acted on before the directory itself. This can be useful when find is used with cpio(1) to transfer files that are contained in directories without write permission.
-mount	Always true; restricts the search to the file system containing the directory specified. Does not list mount points to other file systems.
-xdev	Same as the -mount option.
-local	True if the file system type is not a remote file system type as defined in the /etc/dfs/fstypes file. nfs is used as the default remote filesystem type if the /etc/dfs/fstypes file is not present.
(<i>expression</i>)	True if the parenthesized expression is true (parentheses are special to the shell and must be escaped).
-type <i>c</i>	True if the type of the file is <i>c</i> , where <i>c</i> is b , c , d , l , p , or f for block special file, character special file, directory, symbolic link, fifo (named pipe), or plain file, respectively.
-follow	Always true; causes symbolic links to be followed. When following symbolic links, find keeps track of the directories visited so that it can detect infinite loops; for example, such a loop would occur if a symbolic link pointed to an ancestor. This expression should not be used with the -type l expression.
-links <i>n</i>	True if the file has <i>n</i> links.
-user <i>uname</i>	True if the file belongs to the user <i>uname</i> . If <i>uname</i> is numeric and does not appear as a login name in the /etc/passwd file, it is taken as a user ID.
-nouser	True if the file belongs to a user not in the /etc/passwd file.
-group <i>gname</i>	True if the file belongs to the group <i>gname</i> . If <i>gname</i> is numeric and does not appear in the /etc/group file, it is taken as a group ID.
-nogroup	True if the file belongs to a group not in the /etc/group file.
-fstype <i>type</i>	True if the filesystem to which the file belongs is of type <i>type</i> .
-inum <i>n</i>	True if the file has inode number <i>n</i> .
-prune	Always yields true. Do not examine any directories or files in the directory structure below the <i>pattern</i> just matched. See the examples, below.

Primaries

The primaries may be combined using the following operators (in order of decreasing precedence):

- 1) The negation of a primary (! is the unary *not* operator).
- 2) Concatenation of primaries (the *and* operation is implied by the juxtaposition of two primaries).
- 3) Alternation of primaries (-o is the *or* operator).

Note: When you use **find** in conjunction with **cpio**, if you use the **-L** option with **cpio** then you must use the **-follow** expression with **find** and vice versa. Otherwise there will be undesirable results.

EXAMPLES

Remove all files in your home directory named **a.out** or ***.o** that have not been accessed for a week:

```
example% find $HOME \ ( -name a.out -o -name '*.o' \ ) -atime +7 -exec rm {} \ ;
```

Recursively print all file names in the current directory and below, but skipping SCCS directories:

```
example% find . -name SCCS -prune -o -print
```

Recursively print all file names in the current directory and below, skipping the contents of SCCS directories, but printing out the SCCS directory name:

```
example% find . -print -name SCCS -prune
```

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

FILES

/etc/passwd	password file
/etc/group	group file
/etc/dfs/fstypes	file that registers distributed file system packages

SEE ALSO **chmod(1)**, **ls(1)**, **sh(1)**, **test(1)**, **stat(2)**, **umask(2)**, **environ(5)**

WARNINGS

The following options are obsolete and will not be supported in future releases:

- cpio device** Always true; write the current file on *device* in **cpio** format (5120-byte records).
- ncpio device** Always true; write the current file on *device* in **cpio -c** format (5120 byte records).

NOTES

When using **find** to determine files modified within a range of time, one must use the **?time** argument BEFORE the **-print** argument otherwise **find** will give all files.

NAME	finger – display information about local and remote users
SYNOPSIS	<p>finger [-bfhilmppqsw] [<i>username</i>. . .]</p> <p>finger [-I] [<i>username@hostname1</i>[<i>@hostname2</i>...<i>@hostnamen</i>] . . .]</p> <p>finger [-I] [<i>@hostname1</i>[<i>@hostname2</i>...<i>@hostnamen</i>] . . .]</p>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>By default, the finger command displays in multi-column format the following information about each logged-in user:</p> <ul style="list-style-type: none"> • user name • user's full name • terminal name (prefixed with a '*' (asterisk) if write-permission is denied) • idle time • login time • host name, if logged in remotely <p>Idle time is in minutes if it is a single integer, in hours and minutes if a ':' (colon) is present, or in days and hours if a 'd' is present.</p> <p>When one or more <i>username</i> arguments are given, more detailed information is given for each <i>username</i> specified, whether they are logged in or not. <i>username</i> must be that of a local user, and may be a first or last name, or an account name. Information is presented in multi-line format as follows:</p> <ul style="list-style-type: none"> • the user name and the user's full name • the user's home directory and login shell • time the user logged in if currently logged in, or the time the user last logged in; and the terminal or host from which the user logged in • last time the user received mail, and the last time the user read mail • the first line of the \$HOME/.project file, if it exists • the contents of the \$HOME/.plan file, if it exists <p>If the arguments <i>username@hostname1</i>[<i>@hostname2</i>...<i>@hostnamen</i>] or <i>@hostname1</i>[<i>@hostname2</i>...<i>@hostnamen</i>] are used, the request is sent first to <i>hostnamen</i> and forwarded through each <i>hostnamen-1</i> to <i>hostname1</i>. The program uses the finger user information protocol (see RFC 1288) to query that remote host for information about the named user (if <i>username</i> is specified), or about each logged-in user. The information displayed is server dependent.</p>
OPTIONS	<p>The <i>username@hostname</i> form supports only the -I option.</p> <p>-b Suppress printing the user's home directory and shell in a long format printout.</p> <p>-f Suppress printing the header that is normally printed in a non-long format printout.</p> <p>-h Suppress printing of the .project file in a long format printout.</p> <p>-i Force "idle" output format, which is similar to short format except that only the</p>

login name, terminal, login time, and idle time are printed.

- l** Force long output format.
- m** Match arguments only on user name (not first or last name).
- p** Suppress printing of the **.plan** file in a long format printout.
- q** Force quick output format, which is similar to short format except that only the login name, terminal, and login time are printed.
- s** Force short output format.
- w** Suppress printing the full name in a short format printout.

FILES	\$HOME/.plan	user's plan
	\$HOME/.project	user's projects
	/etc/passwd	password file
	/var/adm/lastlog	time of last login
	/var/adm/utmp	accounting

SEE ALSO **passwd(1)**, **who(1)**, **whois(1)**

NOTES The **finger user information protocol** limits the options that may be used with the remote form of this command.

NAME	fmlcut – cut out selected fields of each line of a file
SYNOPSIS	fmlcut – <i>c</i> <i>list</i> [<i>filename</i> ...] fmlcut – <i>f</i> <i>list</i> [– <i>d</i> <i>char</i>] [– <i>s</i>] [<i>filename</i> ...]
DESCRIPTION	The fmlcut function cuts out columns from a table or fields from each line in <i>filename</i> ; in database parlance, it implements the projection of a relation. fmlcut can be used as a filter; if <i>filename</i> is not specified or is –, the standard input is read. <i>list</i> specifies the fields to be selected. Fields can be fixed length (character positions) or variable length (separated by a field delimiter character), depending on whether – <i>c</i> or – <i>f</i> is specified. Note: Either the – <i>c</i> or the – <i>f</i> option must be specified.
OPTIONS	<p><i>list</i> A comma-separated list of integer field numbers (in increasing order), with optional – to indicate ranges. For example: 1,4,7; 1–3,8; –5,10 (short for 1–5,10); or 3– (short for third through last field).</p> <p>–<i>clist</i> If –<i>c</i> is specified, <i>list</i> specifies character positions (for instance, –c1–72 would pass the first 72 characters of each line). Note: No space intervenes between –<i>c</i> and <i>list</i>.</p> <p>–<i>flist</i> If –<i>f</i> is specified, <i>list</i> is a list of fields assumed to be separated in the file by the default delimiter character, TAB, or by <i>char</i> if the –<i>d</i> option is specified. For example, –f1,7 copies the first and seventh field only. Lines with no delimiter characters are passed through intact (useful for table subheadings), unless –<i>s</i> is specified. Note: No space intervenes between –<i>f</i> and <i>list</i>. The following options can be used if you have specified –<i>f</i>.</p> <p> –<i>dchar</i> If –<i>d</i> is specified, <i>char</i> is the field delimiter. Space or other characters with special meaning to FMLI must be quoted. Note: No space intervenes between –<i>d</i> and <i>char</i>. The default field delimiter is TAB.</p> <p> –<i>s</i> Suppresses lines with no delimiter characters. If –<i>s</i> is not specified, lines with no delimiters will be passed through untouched.</p>
EXAMPLES	The following example gets the login IDS and names. example% fmlcut –d: –f1,5 /etc/passwd The next example gets the current login name. example% `who am i fmlcut –f1 –d" "`
SEE ALSO	fmlgrep(1F)

DIAGNOSTICS

fmlcut returns the following exit values:

- 0 when the selected field is successfully cut out
- 2 on syntax errors

The following error messages may be displayed on the FMLI message line:

ERROR: line too long

A line has more than 1023 characters or fields, or there is no new-line character.

ERROR: bad list for c/f option

Missing **-c** or **-f** option or incorrectly specified *list*. No error occurs if a line has fewer fields than the *list* calls for.

ERROR: no fields

The *list* is empty.

ERROR: no delimiter

Missing *char* on **-d** option.

NOTES

fmlcut cannot correctly process lines longer than 1023 characters, or lines with no new-line character.

NAME	fmlexpr – evaluate arguments as an expression
SYNOPSIS	fmlexpr <i>arguments</i>
DESCRIPTION	<p>The fmlexpr function evaluates its arguments as an expression. After evaluation, the result is written on the standard output. Terms of the expression must be separated by blanks. Characters special to FMLI must be escaped. Note: 30 is returned to indicate a zero value, rather than the null string. Strings containing blanks or other special characters should be quoted. Integer-valued arguments may be preceded by a unary minus sign. Internally, integers are treated as 32-bit, 2s complement numbers.</p> <p>The operators and keywords are listed below. Characters that need to be escaped are preceded by \. The list is in order of increasing precedence, with equal precedence operators grouped within {} symbols.</p>
USAGE	
Expressions	<p><i>expr</i> \ <i>expr</i> Returns the first <i>expr</i> if it is neither null nor 0, otherwise returns the second <i>expr</i>.</p> <p><i>expr</i> \& <i>expr</i> Returns the first <i>expr</i> if neither <i>expr</i> is null or 0, otherwise returns 0.</p> <p><i>expr</i> { =, \>, \>=, \<, \<=, != } <i>expr</i> Returns the result of an integer comparison if both arguments are integers, otherwise returns the result of a lexical comparison.</p> <p><i>expr</i> { +, - } <i>expr</i> Addition or subtraction of integer-valued arguments.</p> <p><i>expr</i> { *, /, % } <i>expr</i> Multiplication, division, or remainder of the integer-valued arguments.</p> <p><i>expr</i> : <i>expr</i> The matching operator : compares the first argument with the second argument which must be a regular expression. Regular expression syntax is the same as that of ed(1), except that all patterns are “anchored” (that is, begin with ^) and, therefore, ^ is not a special character, in that context. Normally, the matching operator returns the number of bytes matched (0 on failure). Alternatively, the \(...\) pattern symbols can be used to return a portion of the first argument.</p>
EXAMPLES	<ol style="list-style-type: none"> Add 1 to the variable a: example% 'fmlexpr \$a + 1 set -l a' For \$a equal to either "/usr/abc/file" or just "file": example% fmlexpr \$a : .*\/(.*\) \ \$a returns the last segment of a path name (that is, <i>file</i>). Watch out for / alone as an argument: fmlexpr will take it as the division operator (see NOTES below).

3. A better representation of example 2.

example% fmlexpr //Sa : .*\/(.*\)

The addition of the // characters eliminates any ambiguity about the division operator (because it makes it impossible for the left-hand expression to be interpreted as the division operator), and simplifies the whole expression.

4. Return the number of characters in \$VAR.

example% fmlexpr \$VAR : .*

SEE ALSO

ed(1), expr(1), set(1F), sh(1)

DIAGNOSTICS

As a side effect of expression evaluation, **fmlexpr** returns the following exit values:

- 0 if the expression is neither null nor 0 (that is, TRUE)
- 1 if the expression is null or 0 (that is, FALSE)
- 2 for invalid expressions (that is, FALSE).

syntax error

for operator/operand errors

non-numeric argument

if arithmetic is attempted on such a string

In the case of syntax errors and non-numeric arguments, an error message will be printed at the current cursor position. Use **refresh** to redraw the screen.

NOTES

After argument processing by FMLI, **fmlexpr** cannot tell the difference between an operator and an operand except by the value. If \$a is an =, the command:

example% fmlexpr \$a = =

looks like:

example% fmlexpr = = =

as the arguments are passed to **fmlexpr** (and they will all be taken as the = operator). The following works, and returns TRUE:

example% fmlexpr X\$a = X=

NAME	fmlgrep – search a file for a pattern
SYNOPSIS	fmlgrep [-b] [-c] [-i] [-l] [-n] [-s] [-v] <i>limited_regular_expression</i> [<i>filename...</i>]
DESCRIPTION	<p>fmlgrep searches <i>filename</i> for a pattern and prints all lines that contain that pattern. The fmlgrep function uses limited regular expressions (expressions that have string values that use a subset of the possible alphanumeric and special characters) like those used with ed(1) to match the patterns. It uses a compact non-deterministic algorithm.</p> <p>Be careful when using FMLI special characters (for instance, \$, `, `;`, `") in <i>limited_regular_expression</i>. It is safest to enclose the entire <i>limited_regular_expression</i> in single quotes '... ' .</p> <p>If <i>filename</i> is not specified, fmlgrep assumes standard input. Normally, each line matched is copied to standard output. The file name is printed before each line matched if there is more than one input-file.</p>
OPTIONS	<p>-b Precede each line by the block number on which it was found. This can be useful in locating block numbers by context (first block is 0).</p> <p>-c Print only a count of the lines that contain the pattern.</p> <p>-i Ignore upper/lower case distinction during comparisons.</p> <p>-l Print only the names of files with matching lines, separated by new-lines. Does not repeat the names of files when the pattern is found more than once.</p> <p>-n Precede each line by its line number in the file (first line is 1).</p> <p>-s Suppress error messages about nonexistent or unreadable files.</p> <p>-v Print all lines except those that contain the pattern.</p>
SEE ALSO	ed (1), egrep (1), fgrep (1), fmlcut (1F), grep (1)
DIAGNOSTICS	<p>fmlgrep returns the following exit values:</p> <p>0 if the pattern is found (that is, TRUE) 1 if the pattern is not found (that is, FALSE) 2 if an invalid expression was used or <i>filename</i> is inaccessible</p>
NOTES	<p>Lines are limited to BUFSIZ characters; longer lines are truncated. BUFSIZ is defined in /usr/include/stdio.h.</p> <p>If there is a line with embedded nulls, fmlgrep will only match up to the first null; if it matches, it will print the entire line.</p>

NAME	fmli – invoke FMLI
SYNOPSIS	fmli [-a <i>alias_file</i>] [-c <i>command_file</i>] [-i <i>initialization_file</i>] <i>filename</i> ...
AVAILABILITY	SUNWcsu
DESCRIPTION	The fmli command invokes the Form and Menu Language Interpreter and opens the frame(s) specified by the <i>filename</i> argument. The <i>filename</i> argument is the pathname of the initial frame definition file(s), and must follow the naming convention Menu.xxx , Form.xxx or Text.xxx for a menu, form or text frame respectively, where <i>xxx</i> is any string that conforms to UNIX system file naming conventions. The FMLI descriptor lifetime will be ignored for all frames opened by argument to fmli . These frames have a lifetime of immortal by default.
OPTIONS	<p>-a <i>alias_file</i> If -a is specified, <i>alias_file</i> is the name of a file which contains lines of the form <i>alias=pathname</i>. Thereafter, <i>\$alias</i> can be used in definition files to simplify references to objects or devices with lengthy pathnames, or to define a search path (similar to \$PATH in the UNIX system shell).</p> <p>-c <i>command_file</i> If -c is specified, <i>command_file</i> is the name of a file in which default FMLI commands can be disabled, and new application-specific commands can be defined. The contents of <i>command_file</i> are reflected in the FMLI Command Menu.</p> <p>-i <i>initialization_file</i> If -i is specified, <i>initialization_file</i> is the name of a file in which the following characteristics of the application as a whole can be specified:</p> <ul style="list-style-type: none"> – A transient introductory frame displaying product information – A banner, its position, and other elements of the banner line – Color attributes for all elements of the screen – Screen Labeled Keys (SLKs) and their layout on the screen.
EXAMPLES	<p>To invoke fmli:</p> <p style="padding-left: 40px;">example% fmli Menu.start</p> <p>where Menu.start is an example of <i>filename</i> named according to the file name conventions for menu definition files explained above.</p> <p>To invoke fmli and name an initialization file:</p> <p style="padding-left: 40px;">example% fmli -i init.myapp Menu.start</p> <p>where init.myapp is an example of <i>initialization_file</i>.</p>

**ENVIRONMENT
Variables**

LOADPFK Leaving this environment variable unset tells FMli, for certain terminals like the AT&T 5620 and 630, to download its equivalent character sequences for using function keys into the terminal's programmable function keys, wiping out any settings the user may already have set in the function keys. Setting **LOADPFK=NO** in the environment will prevent this downloading.

COLUMNS Can be used to override the width of the logical screen defined for the terminal set in **TERM**. For terminals with a 132-column mode, for example, invoking FMli with the line
COLUMNS=132 fmli frame-file
 will allow this wider screen width to be used.

LINES Can be used to override the length of the logical screen defined for the terminal set in **TERM**.

FILES

/usr/bin/fmli

SEE ALSO

vsig(1F)

DIAGNOSTICS

If *filename* is not supplied to the **fmli** command, **fmli** returns the message:

Initial object must be specified.

If *filename* does not exist or is not readable, **fmli** returns an error message and exits. The example command line above returns the following message and exits:

Can't open object "Menu.start"

If *filename* exists, but does not start with one of the three correct object names (**Menu.**, **Form.**, or **Text.**) or if it is named correctly but does not contain the proper data, **fmli** starts to build the screen by putting out the screen labels for function keys, after which it flashes the message:

I do not recognize that kind of object

and then exits.

NAME	fmt – simple text formatters
SYNOPSIS	fmt [-c] [-s] [-w width -width] [<i>inputfile...</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>fmt is a simple text formatter that fills and joins lines to produce output lines of (up to) the number of characters specified in the -w width option. The default <i>width</i> is 72. fmt concatenates the <i>inputfiles</i> listed as arguments. If none are given, fmt formats text from the standard input.</p> <p>Blank lines are preserved in the output, as is the spacing between words. fmt does not fill lines beginning with a '.' (dot), for compatibility with nroff(1). Nor does it fill lines starting with "From:".</p> <p>Indentation is preserved in the output, and input lines with differing indentation are not joined (unless -c is used).</p> <p>fmt can also be used as an in-line text filter for vi(1); the vi command:</p> <pre style="margin-left: 40px;">!}fmt</pre> <p>reformats the text between the cursor location and the end of the paragraph.</p>
OPTIONS	<p>-c Crown margin mode. Preserve the indentation of the first two lines within a paragraph, and align the left margin of each subsequent line with that of the second line. This is useful for tagged paragraphs.</p> <p>-s Split lines only. Do not join short lines to form longer ones. This prevents sample lines of code, and other such formatted text, from being unduly combined.</p> <p>-w width -width Fill output lines to up to <i>width</i> columns.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of fmt for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how fmt behaves.</p> <p>LC_CTYPE Determines how fmt handles characters. When LC_CTYPE is set to a valid value, fmt can display and handle text and filenames containing valid characters for that locale. fmt can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. fmt can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p>

SEE ALSO**nroff(1), vi(1)****NOTES**

The *-width* option is acceptable for BSD compatibility, but it may go away in future releases.

NAME	fmtmsg – display a message on stderr or system console
SYNOPSIS	fmtmsg [-c <i>class</i>] [-u <i>subclass</i>] [-l <i>label</i>] [-s <i>severity</i>] [-t <i>tag</i>] [-a <i>action</i>] <i>text</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>Based on a message's classification component, fmtmsg either writes a formatted message to stderr or writes a formatted message to the console.</p> <p>A formatted message consists of up to five standard components as defined below. The classification and subclass components are not displayed as part of the standard message, but rather define the source of the message and direct the display of the formatted message.</p>
OPTIONS	<p>-c <i>class</i> Describes the source of the message. Valid keywords are:</p> <ul style="list-style-type: none"> hard The source of the condition is hardware. soft The source of the condition is software. firm The source of the condition is firmware. <p>-u <i>subclass</i> A list of keywords (separated by commas) that further defines the message and directs the display of the message. Valid keywords are:</p> <ul style="list-style-type: none"> appl The condition originated in an application. This keyword should not be used in combination with either util or opsys. util The condition originated in a utility. This keyword should not be used in combination with either appl or opsys. opsys The message originated in the kernel. This keyword should not be used in combination with either appl or util. recov The application will recover from the condition. This keyword should not be used in combination with nrecov. nrecov The application will not recover from the condition. This keyword should not be used in combination with recov. print Print the message to the standard error stream stderr. console Write the message to the system console. print, console, or both may be used. <p>-l <i>label</i> Identifies the source of the message.</p> <p>-s <i>severity</i> Indicates the seriousness of the error. The keywords and definitions of the standard levels of <i>severity</i> are:</p> <ul style="list-style-type: none"> halt The application has encountered a severe fault and is halting. error The application has detected a fault. warn The application has detected a condition that is out of the ordinary and might be a problem. info The application is providing information about a condition that is not in error. <p>-t <i>tag</i> The string containing an identifier for the message.</p>

-a *action* A text string describing the first step in the error recovery process. This string must be written so that the entire *action* argument is interpreted as a single argument. **fmtmsg** precedes each action string with the **TO FIX:** prefix.

text A text string describing the condition. Must be written so that the entire *text* argument is interpreted as a single argument.

ENVIRONMENT

The environment variables **MSGVERB** and **SEV_LEVEL** control the behavior of **fmtmsg**. **MSGVERB** is set by the administrator in the **/etc/profile** for the system. Users can override the value of **MSGVERB** set by the system by resetting **MSGVERB** in their own **.profile** files or by changing the value in their current shell session. **SEV_LEVEL** can be used in shell scripts.

MSGVERB tells **fmtmsg** which message components to select when writing messages to **stderr**. The value of **MSGVERB** is a colon separated list of optional keywords.

MSGVERB can be set as follows:

```
MSGVERB=[keyword[:keyword[:...]]]
export MSGVERB
```

Valid *keywords* are: **label**, **severity**, **text**, **action**, and **tag**. If **MSGVERB** contains a keyword for a component and the component's value is not the component's null value, **fmtmsg** includes that component in the message when writing the message to **stderr**. If **MSGVERB** does not include a keyword for a message component, that component is not included in the display of the message. The keywords may appear in any order. If **MSGVERB** is not defined, if its value is the null string, if its value is not of the correct format, or if it contains keywords other than the valid ones listed above, **fmtmsg** selects all components.

MSGVERB affects only which message components are selected for display. All message components are included in console messages.

SEV_LEVEL defines severity levels and associates print strings with them for use by **fmtmsg**. The standard severity levels shown below cannot be modified. Additional severity levels can be defined, redefined, and removed.

```
0 (no severity is used)
1 HALT
2 ERROR
3 WARNING
4 INFO
```

SEV_LEVEL is set as follows:

```
SEV_LEVEL=[description[:description[:...]]]
export SEV_LEVEL
```

description is a comma-separated list containing three fields:

```
description=severity_keyword,level,printstring
```

severity_keyword is a character string used as the keyword with the `-s severity` option to **fmtmsg**.

level is a character string that evaluates to a positive integer (other than **0**, **1**, **2**, **3**, or **4**, which are reserved for the standard severity levels). If the keyword *severity_keyword* is used, *level* is the severity value passed on to **fmtmsg**(3C).

printstring is the character string used by **fmtmsg** in the standard message format whenever the severity value *level* is used.

If **SEV_LEVEL** is not defined, or if its value is null, no severity levels other than the defaults are available. If a *description* in the colon separated list is not a comma separated list containing three fields, or if the second field of a comma separated list does not evaluate to a positive integer, that *description* in the colon separated list is ignored.

EXAMPLES

Example 1: The following example of **fmtmsg** produces a complete message in the standard message format and displays it to the standard error stream:

```
example% fmtmsg -c soft -u recov,print,appl -l UX:cat -s error -t UX:cat:001
-a "refer to manual" "invalid syntax"
```

produces:

```
UX:cat: ERROR: invalid syntax
TO FIX: refer to manual UX:cat:138
```

Example 2: When the environment variable **MSGVERB** is set as follows:

```
MSGVERB=severity:text:action
```

and Example 1 is used, **fmtmsg** produces:

```
ERROR: invalid syntax
TO FIX: refer to manual
```

Example 3: When the environment variable **SEV_LEVEL** is set as follows:

```
SEV_LEVEL=note,5,NOTE
```

the following **fmtmsg** command:

```
example% fmtmsg -c soft -u print -l UX:cat -s note -a "refer to manual"
"invalid syntax"
```

produces:

```
NOTE: invalid syntax
To FIX: refer to manual
```

and displays the message on **stderr**.

SEE ALSO

addseverity(3C), **fmtmsg**(3C)

DIAGNOSTICS

The exit codes for **fmtmsg** are the following:

- 0** All the requested functions were executed successfully.
- 1** The command contains a syntax error, an invalid option, or an invalid argument to an option.
- 2** The function executed with partial success, however the message was not displayed on **stderr**.
- 4** The function executed with partial success, however the message was not displayed on the system console.
- 32** No requested functions were executed successfully.

NAME	fold – fold long lines
SYNOPSIS	fold [-w <i>width</i> -width] [<i>filename</i> . . .]
AVAILABILITY	SUNWcsu
DESCRIPTION	Fold the contents of the specified <i>filenames</i> , or the standard input if no files are specified, breaking the lines to have maximum width <i>width</i> . The default for <i>width</i> is 80. <i>width</i> should be a multiple of 8 if tabs are present, or the tabs should be expanded.
OPTIONS	-w <i>width</i> -width The contents of the specified <i>filenames</i> or the standard input are folded breaking the lines to have maximum width <i>width</i> , given either of these options.
SEE ALSO	pr (1)
NOTES	Folding may not work correctly if underlining is present. The -width option is provided as a transition tool only. It will be removed in future releases.

NAME	for, foreach – shell built-in functions to repeatedly execute action(s) for a selected number of times
SYNOPSIS	
sh	for <i>name</i> [in <i>word...</i>]; do <i>actions</i> ; done
csh	foreach <i>var</i> (<i>wordlist</i>) ... end
ksh	for <i>identifier</i> [in <i>word...</i>]; do <i>actions</i> ; done
DESCRIPTION	
sh	Each time a for command is executed, <i>name</i> is set to the next <i>word</i> taken from the in <i>word</i> list. If in <i>word...</i> is omitted, then the for command executes the do <i>actions</i> once for each positional parameter that is set. Execution ends when there are no more words in the list.
csh	The variable <i>var</i> is successively set to each member of <i>wordlist</i> . The sequence of commands between this command and the matching end is executed for each new value of <i>var</i> . Both foreach and end must appear alone on separate lines. The built-in command continue may be used to terminate the execution of the current iteration of the loop and the built-in command break may be used to terminate execution of the foreach command. When this command is read from the terminal, the loop is read once prompting with ? before any statements in the loop are executed.
ksh	Each time a for command is executed, <i>identifier</i> is set to the next <i>word</i> taken from the in <i>word</i> list. If in <i>word...</i> is omitted, then the for command executes the do <i>actions</i> once for each positional parameter that is set. Execution ends when there are no more words in the list.
SEE ALSO	break(1) , csh(1) , ksh(1) , sh(1)
NOTES	Both the Bourne shell, sh , and the Korn shell, ksh , can use the semicolon and the carriage return interchangeably in their syntax of the if , for , and while built-in commands.

NAME	from – display the sender and date of newly-arrived mail messages
SYNOPSIS	/usr/ucb/from [<i>-s sender</i>] [<i>username</i>]
AVAILABILITY	SUNWscpu
DESCRIPTION	from prints out the mail header lines in your mailbox file to show you who your mail is from. If <i>username</i> is specified, then <i>username</i> 's mailbox is examined instead of your own.
OPTIONS	-s sender Only display headers for mail sent by <i>sender</i> .
FILES	/var/spool/mail/*
SEE ALSO	biff(1B), mail(1B)

NAME	ftp – file transfer program
SYNOPSIS	ftp [-dgintv] [<i>hostname</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The ftp command is the user interface to the Internet standard File Transfer Protocol (FTP). ftp transfers files to and from a remote network site.</p> <p>The client host with which ftp is to communicate may be specified on the command line. If this is done, ftp immediately attempts to establish a connection to an FTP server on that host; otherwise, ftp enters its command interpreter and awaits instructions from the user. When ftp is awaiting commands from the user, it displays the prompt ftp>.</p>
OPTIONS	<p>The following options may be specified at the command line, or to the command interpreter:</p> <ul style="list-style-type: none"> -d Enable debugging. -g Disable filename “globbing.” -i Turn off interactive prompting during multiple file transfers. -n Do not attempt “auto-login” upon initial connection. If auto-login is not disabled, ftp checks the .netrc file in the user’s home directory for an entry describing an account on the remote machine. If no entry exists, ftp will prompt for the login name of the account on the remote machine (the default is the login name on the local machine), and, if necessary, prompts for a password and an account with which to login. -t Enable packet tracing (unimplemented). -v Show all responses from the remote server, as well as report on data transfer statistics. This is turned on by default if ftp is running interactively with its input coming from the user’s terminal. <p>The following commands can be specified to the command interpreter:</p> <p>! [<i>command</i>] Run <i>command</i> as a shell command on the local machine. If no <i>command</i> is given, invoke an interactive shell.</p> <p>\$ <i>macro-name</i> [<i>args</i>] Execute the macro <i>macro-name</i> that was defined with the macdef command. Arguments are passed to the macro unglobbed.</p> <p>account [<i>passwd</i>] Supply a supplemental password required by a remote system for access to resources once a login has been successfully completed. If no argument is included, the user will be prompted for an account password in a non-echoing input mode.</p>

- append** *local-file* [*remote-file*]
Append a local file to a file on the remote machine. If *remote-file* is not specified, the local file name is used, subject to alteration by any **ntrans** or **nmap** settings. File transfer uses the current settings for “representation type”, “file structure”, and “transfer mode”.
- ascii** Set the “representation type” to “network ASCII”. This is the default type.
- bell** Sound a bell after each file transfer command is completed.
- binary** Set the “representation type” to “image”.
- bye** Terminate the FTP session with the remote server and exit **ftp**. An EOF will also terminate the session and exit.
- case** Toggle remote computer file name case mapping during **mget** commands. When **case** is on (default is off), remote computer file names with all letters in upper case are written in the local directory with the letters mapped to lower case.
- cd** *remote-directory*
Change the working directory on the remote machine to *remote-directory*.
- cdup** Change the remote machine working directory to the parent of the current remote machine working directory.
- close** Terminate the FTP session with the remote server, and return to the command interpreter. Any defined macros are erased.
- cr** Toggle RETURN stripping during “network ASCII” type file retrieval. Records are denoted by a RETURN/LINEFEED sequence during “network ASCII” type file transfer. When **cr** is on (the default), RETURN characters are stripped from this sequence to conform with the UNIX system single LINEFEED record delimiter. Records on non-UNIX-system remote hosts may contain single LINEFEED characters; when an “network ASCII” type transfer is made, these LINEFEED characters may be distinguished from a record delimiter only when **cr** is off.
- delete** *remote-file*
Delete the file *remote-file* on the remote machine.
- debug**
Toggle debugging mode. When debugging is on, **ftp** prints each command sent to the remote machine, preceded by the string —>.
- dir** [*remote-directory*] [*local-file*]
Print a listing of the directory contents in the directory, *remote-directory*, and, optionally, placing the output in *local-file*. If no directory is specified, the current working directory on the remote machine is used. If no local file is specified, or *local-file* is —, output is sent to the terminal.
- disconnect**
A synonym for **close**.

- form** [*format-name*]
Set the carriage control format subtype of the “representation type” to *format-name*. The only valid *format-name* is **non-print**, which corresponds to the default “non-print” subtype.
- get remote-file** [*local-file*]
Retrieve the *remote-file* and store it on the local machine. If the local file name is not specified, it is given the same name it has on the remote machine, subject to alteration by the current **case**, **ntrans**, and **nmap** settings. The current settings for “representation type”, “file structure”, and “transfer mode” are used while transferring the file.
- glob** Toggle filename expansion, or “globbing”, for **mdelete**, **mget** and **mput**. If globbing is turned off, filenames are taken literally.
Globbing for **mput** is done as in **sh**(1). For **mdelete** and **mget**, each remote file name is expanded separately on the remote machine, and the lists are not merged.
Expansion of a directory name is likely to be radically different from expansion of the name of an ordinary file: the exact result depends on the remote operating system and FTP server, and can be previewed by doing **mls remote-files -**.
mget and **mput** are not meant to transfer entire directory subtrees of files. You can do this by transferring a **tar**(1) archive of the subtree (using a “representation type” of “image” as set by the **binary** command).
- hash** Toggle hash-sign (#) printing for each data block transferred. The size of a data block is 8192 bytes.
- help** [*command*]
Print an informative message about the meaning of *command*. If no argument is given, **ftp** prints a list of the known commands.
- lcd** [*directory*]
Change the working directory on the local machine. If no *directory* is specified, the user’s home directory is used.
- ls** [*remote-directory* | **-al**] [*local-file*]
Print an abbreviated listing of the contents of a directory on the remote machine. If *remote-directory* is left unspecified, the current working directory is used.
The **-a** option lists all entries, including those that begin with a dot (.), which are normally not listed. The **-l** option lists files in long format, giving mode, number of links, owner, group, size in bytes, and time of last modification for each file. If the file is a special file, the size field instead contains the major and minor device numbers rather than a size. If the file is a symbolic link, the filename is printed followed by “→” and the pathname of the referenced file.
If no local file is specified, or if *local-file* is –, the output is sent to the terminal.

macdef *macro-name*

Define a macro. Subsequent lines are stored as the macro *macro-name*; a null line (consecutive NEWLINE characters in a file or RETURN characters from the terminal) terminates macro input mode. There is a limit of 16 macros and 4096 total characters in all defined macros. Macros remain defined until a **close** command is executed.

The macro processor interprets **\$** and **** as special characters. A **\$** followed by a number (or numbers) is replaced by the corresponding argument on the macro invocation command line. A **\$** followed by an **i** signals that macro processor that the executing macro is to be looped. On the first pass **\$i** is replaced by the first argument on the macro invocation command line, on the second pass it is replaced by the second argument, and so on. A **** followed by any character is replaced by that character. Use the **** to prevent special treatment of the **\$**.

mdelete *remote-files*

Delete the *remote-files* on the remote machine.

mdir *remote-files local-file*

Like **dir**, except multiple remote files may be specified. If interactive prompting is on, **ftp** will prompt the user to verify that the last argument is indeed the target local file for receiving **mdir** output.

mget *remote-files*

Expand the *remote-files* on the remote machine and do a **get** for each file name thus produced. See **glob** for details on the filename expansion. Resulting file names will then be processed according to **case**, **ntrans**, and **nmap** settings. Files are transferred into the local working directory, which can be changed with **lcd directory**; new local directories can be created with **! mkdir directory**.

mkdir *directory-name*

Make a directory on the remote machine.

mls *remote-files local-file*

Like **ls(1)**, except multiple remote files may be specified. If interactive prompting is on, **ftp** will prompt the user to verify that the last argument is indeed the target local file for receiving **mls** output.

mode [*mode-name*]

Set the “transfer mode” to *mode-name*. The only valid *mode-name* is **stream**, which corresponds to the default “stream” mode. This implementation only supports **stream**, and requires that it be specified.

mput *local-files*

Expand wild cards in the list of local files given as arguments and do a **put** for each file in the resulting list. See **glob** for details of filename expansion. Resulting file names will then be processed according to **ntrans** and **nmap** settings.

nmap [*inpattern outpattern*]

Set or unset the filename mapping mechanism. If no arguments are specified, the filename mapping mechanism is unset. If arguments are specified, remote filenames are mapped during **mput** commands and **put** commands issued without a specified remote target filename. If arguments are specified, local filenames are mapped during **mget** commands and **get** commands issued without a specified local target filename.

This command is useful when connecting to a non-UNIX-system remote host with different file naming conventions or practices. The mapping follows the pattern set by *inpattern* and *outpattern*. *inpattern* is a template for incoming filenames (which may have already been processed according to the **ntrans** and **case** settings). Variable templating is accomplished by including the sequences **\$1**, **\$2**, ..., **\$9** in *inpattern*. Use `\` to prevent this special treatment of the `$` character. All other characters are treated literally, and are used to determine the **nmap** *inpattern* variable values.

For example, given *inpattern* **\$1.\$2** and the remote file name **mydata.data**, **\$1** would have the value **mydata**, and **\$2** would have the value **data**.

The *outpattern* determines the resulting mapped filename. The sequences **\$1**, **\$2**, ..., **\$9** are replaced by any value resulting from the *inpattern* template. The sequence **\$0** is replaced by the original filename. Additionally, the sequence [*seq1* , *seq2*] is replaced by *seq1* if *seq1* is not a null string; otherwise it is replaced by *seq2*.

For example, the command **nmap \$1.\$2.\$3 [\$1,\$2].[\$2,file]** would yield the output filename **myfile.data** for input filenames **myfile.data** and **myfile.data.old**, **myfile.file** for the input filename **myfile**, and **myfile.myfile** for the input filename **.myfile**. SPACE characters may be included in *outpattern*, as in the example **nmap \$1 | sed "s/ *\$/" > \$1**. Use the `\` character to prevent special treatment of the `$`, `[`, `]`, and `,` characters.

ntrans [*inchars* [*outchars*]]

Set or unset the filename character translation mechanism. If no arguments are specified, the filename character translation mechanism is unset. If arguments are specified, characters in remote filenames are translated during **mput** commands and **put** commands issued without a specified remote target filename, and characters in local filenames are translated during **mget** commands and **get** commands issued without a specified local target filename.

This command is useful when connecting to a non-UNIX-system remote host with different file naming conventions or practices. Characters in a filename matching a character in *inchars* are replaced with the corresponding character in *outchars*. If the character's position in *inchars* is longer than the length of *outchars*, the character is deleted from the file name.

Only 16 characters can be translated when using the **ntrans** command under **ftp**. Use **case** (described above) if needing to convert the entire alphabet.

open *host* [*port*]

Establish a connection to the specified *host* FTP server. An optional port number may be supplied, in which case, **ftp** will attempt to contact an FTP server at that port. If the *auto-login* option is on (default setting), **ftp** will also attempt to automatically log the user in to the FTP server.

prompt

Toggle interactive prompting. Interactive prompting occurs during multiple file transfers to allow the user to selectively retrieve or store files. By default, prompting is turned on. If prompting is turned off, any **mget** or **mput** will transfer all files, and any **mdelete** will delete all files.

proxy *ftp-command*

Execute an FTP command on a secondary control connection. This command allows simultaneous connection to two remote FTP servers for transferring files between the two servers. The first **proxy** command should be an **open**, to establish the secondary control connection. Enter the command **proxy ?** to see other FTP commands executable on the secondary connection.

The following commands behave differently when prefaced by **proxy**: **open** will not define new macros during the auto-login process, **close** will not erase existing macro definitions, **get** and **mget** transfer files from the host on the primary control connection to the host on the secondary control connection, and **put**, **mputd**, and **append** transfer files from the host on the secondary control connection to the host on the primary control connection.

Third party file transfers depend upon support of the **PASV** command by the server on the secondary control connection.

put *local-file* [*remote-file*]

Store a local file on the remote machine. If *remote-file* is left unspecified, the local file name is used after processing according to any **ntrans** or **nmap** settings in naming the remote file. File transfer uses the current settings for “representation type”, “file structure”, and “transfer mode”.

pwd Print the name of the current working directory on the remote machine.

quit A synonym for **bye**.

quote *arg1 arg2 ...*

Send the arguments specified, verbatim, to the remote FTP server. A single FTP reply code is expected in return. (The **remotehelp** command displays a list of valid arguments.)

quote should be used only by experienced users who are familiar with the FTP protocol.

recv *remote-file* [*local-file*]

A synonym for **get**.

remotehelp [*command-name*]

Request help from the remote FTP server. If a *command-name* is specified it is supplied to the server as well.

rename *from to*

Rename the file *from* on the remote machine to have the name *to*.

reset Clear reply queue. This command re-synchronizes command/reply sequencing with the remote FTP server. Resynchronization may be necessary following a violation of the FTP protocol by the remote server.

rmdir *directory-name*

Delete a directory on the remote machine.

runique

Toggle storing of files on the local system with unique filenames. If a file already exists with a name equal to the target local filename for a **get** or **mget** command, a **.1** is appended to the name. If the resulting name matches another existing file, a **.2** is appended to the original name. If this process continues up to **.99**, an error message is printed, and the transfer does not take place. The generated unique filename will be reported. **runique** will not affect local files generated from a shell command. The default value is off.

send *local-file* [*remote-file*]

A synonym for **put**.

sendport

Toggle the use of **PORT** commands. By default, **ftp** will attempt to use a **PORT** command when establishing a connection for each data transfer. The use of **PORT** commands can prevent delays when performing multiple file transfers. If the **PORT** command fails, **ftp** will use the default data port. When the use of **PORT** commands is disabled, no attempt will be made to use **PORT** commands for each data transfer. This is useful when connected to certain FTP implementations that ignore **PORT** commands but incorrectly indicate they have been accepted.

status Show the current status of **ftp**.

struct [*struct-name*]

Set the file structure to *struct-name*. The only valid *struct-name* is **file**, which corresponds to the default “file” structure. The implementation only supports **file**, and requires that it be specified.

sunique

Toggle storing of files on remote machine under unique file names. The remote FTP server must support the **STOU** command for successful completion. The remote server will report the unique name. Default value is off.

tenex Set the “representation type” to that needed to talk to TENEX machines.

trace Toggle packet tracing (unimplemented).

type [*type-name*]

Set the “representation type” to *type-name*. The valid *type-names* are **ascii** for “network ASCII”, **binary** or **image** for “image”, and **tenex** for “local byte size” with a byte size of 8 (used to talk to TENEX machines). If no type is specified, the current type is printed. The default type is “network ASCII”.

user *user-name* [*password*] [*account*]

Identify yourself to the remote FTP server. If the password is not specified and the server requires it, **ftp** will prompt the user for it (after disabling local echo). If an account field is not specified, and the FTP server requires it, the user will be prompted for it. If an account field is specified, an account command will be relayed to the remote server after the login sequence is completed if the remote server did not require it for logging in. Unless **ftp** is invoked with "auto-login" disabled, this process is done automatically on initial connection to the FTP server.

verbose

Toggle verbose mode. In verbose mode, all responses from the FTP server are displayed to the user. In addition, if verbose mode is on, when a file transfer completes, statistics regarding the efficiency of the transfer are reported. By default, verbose mode is on if **ftp**'s commands are coming from a terminal, and off otherwise.

? [*command*]

A synonym for **help**.

Command arguments which have embedded spaces may be quoted with quote (") marks. If any command argument which is not indicated as being optional is not specified, **ftp** will prompt for that argument.

ABORTING A FILE TRANSFER

To abort a file transfer, use the terminal interrupt key. Sending transfers will be immediately halted. Receiving transfers will be halted by sending an FTP protocol **ABOR** command to the remote server, and discarding any further data received. The speed at which this is accomplished depends upon the remote server's support for **ABOR** processing. If the remote server does not support the **ABOR** command, an **ftp>** prompt will not appear until the remote server has completed sending the requested file.

The terminal interrupt key sequence will be ignored when **ftp** has completed any local processing and is awaiting a reply from the remote server. A long delay in this mode may result from the **ABOR** processing described above, or from unexpected behavior by the remote server, including violations of the ftp protocol. If the delay results from unexpected remote server behavior, the local **ftp** program must be killed by hand.

FILE NAMING CONVENTIONS

Local files specified as arguments to **ftp** commands are processed according to the following rules.

- 1) If the file name `-` is specified, the standard input (for reading) or standard output (for writing) is used.
- 2) If the first character of the file name is `|`, the remainder of the argument is interpreted as a shell command. **ftp** then forks a shell, using **popen(3S)** with the argument supplied, and reads (writes) from the standard output (standard input) of that shell. If the shell command includes SPACE characters, the argument must be quoted; for example "`| ls -lt`". A particularly useful example of this mechanism is: "`dir | more`".

- 3) Failing the above checks, if globbing is enabled, local file names are expanded according to the rules used in the **sh**(1); see the **glob** command. If the **ftp** command expects a single local file (for example, **put**), only the first filename generated by the globbing operation is used.
- 4) For **mget** commands and **get** commands with unspecified local file names, the local filename is the remote filename, which may be altered by a **case**, **ntrans**, or **nmap** setting. The resulting filename may then be altered if **runique** is on.
- 5) For **mput** commands and **put** commands with unspecified remote file names, the remote filename is the local filename, which may be altered by a **ntrans** or **nmap** setting. The resulting filename may then be altered by the remote server if **sunique** is on.

FILE TRANSFER PARAMETERS

The FTP specification specifies many parameters which may affect a file transfer.

The “representation type” may be one of “network ASCII”, “EBCDIC”, “image”, or “local byte size” with a specified byte size (for PDP-10’s and PDP-20’s mostly). The “network ASCII” and “EBCDIC” types have a further subtype which specifies whether vertical format control (NEWLINE characters, form feeds, etc.) are to be passed through (“non-print”), provided in TELNET format (“TELNET format controls”), or provided in ASA (FORTRAN) (“carriage control (ASA)”) format. **ftp** supports the “network ASCII” (subtype “non-print” only) and “image” types, plus “local byte size” with a byte size of 8 for communicating with TENEX machines.

The “file structure” may be one of **file** (no record structure), **record**, or **page**. **ftp** supports only the default value, which is **file**.

The “transfer mode” may be one of **stream**, **block**, or **compressed**. **ftp** supports only the default value, which is **stream**.

FILES

~/.netrc

SEE ALSO

ls(1), **rcp**(1), **sh**(1), **tar**(1), **ftpd**(1M), **popen**(3S), **netrc**(4)

NOTES

Correct execution of many commands depends upon proper behavior by the remote server.

An error in the treatment of carriage returns in the 4.2 BSD code handling transfers with a “representation type” of “network ASCII” has been corrected. This correction may result in incorrect transfers of binary files to and from 4.2 BSD servers using a “representation type” of “network ASCII”. Avoid this problem by using the “image” type.

NAME	function – shell built-in command to define a function which is usable within this shell
SYNOPSIS ksh	function <i>identifier</i> { <i>list</i> ;} <i>identifier</i> () { <i>list</i> ;}
DESCRIPTION ksh	function defines a function which is referenced by <i>identifier</i> . The body of the function is the <i>list</i> of commands between { and }. Alternatively, omitting the function keyword and appending the <i>identifier</i> with a set of enclosed parentheses will accomplish the same function definition.
SEE ALSO	ksh(1)

NAME	gcore – get core images of running processes
SYNOPSIS	gcore [-o <i>filename</i>] [-p <i>procdir</i>] <i>process-id</i> . . .
DESCRIPTION	gcore creates a core image of each specified process. The name of the core image file for the process whose process ID is <i>process-id</i> will be core.process-id .
OPTIONS	-o Substitutes <i>filename</i> in place of core as the first part of the name of the core image files.
FILES	core.process-id core images
SEE ALSO	csh(1) , kill(1) , ptrace(2)

NAME	gencat – generate a formatted message catalog
SYNOPSIS	gencat <i>catfile msgfile...</i>
AVAILABILITY	SUNWloc
DESCRIPTION	<p>gencat merges the message text source file(s) <i>msgfile</i> into a formatted message database <i>catfile</i>. The database <i>catfile</i> is created if it does not already exist. If <i>catfile</i> does exist, its messages are included in the new <i>catfile</i>. If set and message numbers collide, the new message-text defined in <i>msgfile</i> replaces the old message text currently contained in <i>catfile</i>. The message text source file (or set of files) input to gencat can contain either set and message numbers or simply message numbers, in which case the set NL_SETD (see nl_types(5)) is assumed.</p> <p>The format of a message text source file is defined as follows. Note that the fields of a message text source line are separated by a single ASCII space or tab character. Any other ASCII spaces or tabs are considered as part of the subsequent field.</p> <p>\$set n comment Where <i>n</i> specifies the set identifier of the following messages until the next \$set, \$delset, or end-of-file appears. <i>n</i> must be a number in the range (1–{NL_SETMAX}). Set identifiers within a single source file need not be contiguous. Any string following the set identifier is treated as a comment. If no \$set directive is specified in a message text source file, all messages are located in the default message set NL_SETD.</p> <p>\$delset n comment Deletes message set <i>n</i> from an existing message catalog. Any string following the set number is treated as a comment. (Note: if <i>n</i> is not a valid set it is ignored.)</p> <p>\$ comment A line beginning with a dollar symbol \$ followed by an ASCII space or tab character is treated as a comment.</p> <p>m message-text The <i>m</i> denotes the message identifier, a number in the range (1–{NL_MSGMAX}). The message-text is stored in the message catalog with the set identifier specified by the last \$set directive, and with message identifier <i>m</i>. If the message-text is empty, and an ASCII space or tab field separator is present, an empty string is stored in the message catalog. If a message source line has a message number, but neither a field separator nor message-text, the existing message with that number (if any) is deleted from the catalog. Message identifiers need not be contiguous. The length of message-text must be in the range (0–{NL_TEXTMAX}).</p> <p>\$quote c This line specifies an optional quote character <i>c</i>, which can be used to surround message-text so that trailing spaces or null (empty) messages are visible in a message source line. By default, or if an empty \$quote directive is supplied, no quoting of message-text will be recognized.</p>

Empty lines in a message text source file are ignored.

Text strings can contain the special characters and escape sequences defined in the following table:

Description	Symbol	Sequence
newline	NL(LF)	\n
horizontal tab	HT	\t
vertical tab	VT	\v
backspace	BS	\b
carriage return	CR	\r
form feed	FF	\f
backslash	\	\\
bit pattern	ddd	\ddd

The escape sequence `\ddd` consists of backslash followed by 1, 2 or 3 octal digits, which are taken to specify the value of the desired character. If the character following a backslash is not one of those specified, the backslash is ignored.

Backslash followed by an ASCII newline character is also used to continue a string on the following line. Thus, the following two lines describe a single message string:

```
1 This line continues \  
to the next line
```

which is equivalent to:

```
1 This line continues to the next line
```

ENVIRONMENT

After message catalogs are generated and installed, the `LANG` environment variable, or the `LC_MESSAGES` environment if present, determines the prevailing message locale.

SEE ALSO

`mkmsgs(1)`, `catgets(3C)`, `catopen(3C)`, `gettext(3C)`, `environ(5)`, `nl_types(5)`

NAME	getfrm – returns the current frameID number
SYNOPSIS	getfrm
DESCRIPTION	getfrm returns the current frameID number. The frameID number is a number assigned to the frame by FMLI and displayed flush left in the frame's title bar. If a frame is closed its frameID number may be reused when a new frame is opened. getfrm takes no arguments.
EXAMPLES	If a menu whose frameID is 3 defines an item to have this action descriptor: action=open text stdtext `getfrm` the text frame defined in the definition file stdtext would be passed the argument 3 when it is opened.
NOTES	It is not a good idea to use getfrm in a backquoted expression coded on a line by itself. Stand-alone backquoted expressions are evaluated before any descriptors are parsed, thus the frame is not yet fully current, and may not have been assigned a frameID number.

NAME	getitems – returns a list of currently marked menu items
SYNOPSIS	getitems [<i>delimiter_string</i>]
DESCRIPTION	The getitems function returns the value of lininfo if defined, else it returns the value of the name descriptor, for all currently marked menu items. Each value in the list is delimited by <i>delimiter_string</i> . The default value of <i>delimiter_string</i> is newline.
EXAMPLES	<p>The done descriptor in the following menu definition file executes getitems when the user presses ENTER (note that the menu is multiselect):</p> <pre>Menu="Example" multiselect=TRUE done=`getitems ":" message` name="Item 1" action=`message "You selected item 1"` name="Item 2" lininfo="This is item 2" action=`message "You selected item 2"` name="Item 3" action=`message "You selected item 3"`</pre> <p>If a user marked all three items in this menu, pressing ENTER would cause the following string to be displayed on the message line:</p> <pre>Item 1:This is item 2:Item 3</pre> <p>Note: Because lininfo is defined for the second menu item, its value is displayed instead of the value of the name descriptor.</p>

NAME	getopt – parse command options
SYNOPSIS	set — `getopt <i>optstring</i> \$*`
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The getopts command supersedes getopt. For more information, see the NOTES below.</p> <p>getopt is used to break up options in command lines for easy parsing by shell procedures and to check for legal options. <i>optstring</i> is a string of recognized option letters; see getopt(3C). If a letter is followed by a colon, the option is expected to have an argument which may or may not be separated from it by white space. The special option — is used to delimit the end of the options. If it is used explicitly, getopt recognizes it; otherwise, getopt generates it; in either case, getopt places it at the end of the options. The positional parameters (\$1 \$2 ...) of the shell are reset so that each option is preceded by a — and is in its own positional parameter; each option argument is also parsed into its own positional parameter.</p>
EXAMPLES	<p>The following code fragment shows how one might process the arguments for a command that can take the options a or b, as well as the option o, which requires an argument:</p> <pre> set — `getopt abo: \$*` if [\$? != 0] then echo \$USAGE exit 2 fi for i in \$* do case \$i in -a -b) FLAG=\$i; shift;; -o) OARG=\$2; shift 2;; —) shift; break;; esac done </pre> <p>This code accepts any of the following as equivalent:</p> <pre> cmd -aoarg filename1 filename2 cmd -a -o arg filename1 filename2 cmd -oarg -a filename1 filename2 cmd -a -oarg — filename1 filename2 </pre>
SEE ALSO	shell_builtins(1) , sh(1) , getopt(3C)

DIAGNOSTICS

getopt prints an error message on the standard error when it encounters an option letter not included in *optstring*.

NOTES

getopt will not be supported in the next major release. For this release a conversion tool has been provided, **getoptcv**. For more information about **getopts** and **getoptcv**, see **getopts(1)**.

Reset **optind** to 1 when rescanning the options.

getopt does not support the part of Rule 8 of the command syntax standard (see **intro(1)**) that permits groups of option-arguments following an option to be separated by white space and quoted. For example,

```
cmd -a -b -o "xxx z yy" filename
```

is not handled correctly. To correct this deficiency, use the **getopts** command in place of **getopt**.

If an option that takes an option-argument is followed by a value that is the same as one of the options listed in *optstring* (referring to the earlier EXAMPLE section, but using the following command line: **cmd -o -a filename**), **getopt** always treats **-a** as an option-argument to **-o**; it never recognizes **-a** as an option. For this case, the **for** loop in the example shifts past the *filename* argument.

NAME	getoptcv – convert to getopt to parse command options
SYNOPSIS	<code>/usr/lib/getoptcv [-b] filename</code> <code>/usr/lib/getoptcv</code>
DESCRIPTION	<p><code>/usr/lib/getoptcv</code> reads the shell script in <i>filename</i>, converts it to use getopts instead of getopt, and writes the results on the standard output.</p> <p>getopts is a built-in Bourne shell command used to parse positional parameters and to check for valid options. See sh(1). It supports all applicable rules of the command syntax standard (see Rules 3-10, intro(1)). It should be used in place of the getopt command. (See the NOTES section below.) The syntax for the shell's built-in getopts command is:</p> <p>getopts <i>optstring name</i> [<i>argument . . .</i>]</p> <p><i>optstring</i> must contain the option letters the command using getopts will recognize; if a letter is followed by a colon, the option is expected to have an argument, or group of arguments, which must be separated from it by white space.</p> <p>Each time it is invoked, getopts places the next option in the shell variable <i>name</i> and the index of the next argument to be processed in the shell variable OPTIND. Whenever the shell or a shell script is invoked, OPTIND is initialized to 1.</p> <p>When an option requires an option-argument, getopts places it in the shell variable OPTARG.</p> <p>If an illegal option is encountered, ? will be placed in <i>name</i>.</p> <p>When the end of options is encountered, getopts exits with a non-zero exit status. The special option <code>—</code> may be used to delimit the end of the options.</p> <p>By default, getopts parses the positional parameters. If extra arguments (<i>argument . . .</i>) are given on the getopts command line, getopts parses them instead.</p> <p>So that all new commands will adhere to the command syntax standard described in intro(1), they should use getopts or getopt to parse positional parameters and check for options that are valid for that command (see the NOTES section below).</p>
OPTIONS	<p>-b Make the converted script portable to earlier releases of the UNIX system.</p> <p><code>/usr/lib/getoptcv</code> modifies the shell script in <i>filename</i> so that when the resulting shell script is executed, it determines at run time whether to invoke getopts or getopt.</p>
EXAMPLES	The following fragment of a shell program shows how one might process the arguments for a command that can take the options a or b , as well as the option o , which requires an option-argument:

```

while getoptcs abo: c
do
    case $c in
    a| b)    FLAG=$c;;
    o)      OARG=$OPTARG;;
    \?)    echo $USAGE
           exit 2;;
    esac
done
shift .expr $OPTARG - 1.

```

This code accepts any of the following as equivalent:

```

cmd -a -b -o "xxx z yy" filename
cmd -a -b -o "xxx z yy" -- filename
cmd -ab -o xxx,z,yy filename
cmd -ab -o "xxx z yy" filename
cmd -o xxx,z,yy -b -a filename

```

SEE ALSO [intro\(1\)](#), [sh\(1\)](#), [shell_builtins\(1\)](#), [getopt\(3C\)](#)

DIAGNOSTICS **getopts** prints an error message on the standard error when it encounters an option letter not included in *optstring*.

NOTES Although the following command syntax rule (see [intro\(1\)](#)) relaxations are permitted under the current implementation, they should not be used because they may not be supported in future releases of the system. As in the **EXAMPLES** section above, **a** and **b** are options, and the option **o** requires an option-argument. The following example violates Rule 5: options with option-arguments must not be grouped with other options:

```
example% cmd -abxxx filename
```

The following example violates Rule 6: there must be white space after an option that takes an option-argument:

```
example% cmd -ab -oxxx filename
```

Changing the value of the shell variable **OPTARG** or parsing different sets of arguments may lead to unexpected results.

NAME	getopts – shell built-in function to parse command-line options
SYNOPSIS	
sh	getopts <i>optstring name</i> [<i>argument ...</i>]
ksh	getopts <i>optstring name</i> [<i>arg... </i>]
DESCRIPTION	
sh	<p>getopts is a built-in Bourne shell command used to parse positional parameters and to check for valid options. See sh(1). It supports all applicable rules of the command syntax standard (see Rules 3-10, intro(1)). It should be used in place of the getopt command.</p> <p><i>optstring</i> must contain the option letters the command using getopts will recognize; if a letter is followed by a colon, the option is expected to have an argument, or group of arguments, which must be separated from it by white space.</p> <p>Each time it is invoked, getopts places the next option in the shell variable <i>name</i> and the index of the next argument to be processed in the shell variable OPTIND. Whenever the shell or a shell script is invoked, OPTIND is initialized to 1.</p> <p>When an option requires an option-argument, getopts places it in the shell variable OPTARG.</p> <p>If an illegal option is encountered, ? will be placed in <i>name</i>.</p> <p>When the end of options is encountered, getopts exits with a non-zero exit status. The special option — may be used to delimit the end of the options.</p> <p>By default, getopts parses the positional parameters. If extra arguments (<i>argument ...</i>) are given on the getopts command line, getopts parses them instead.</p> <p>/usr/lib/getoptcv reads the shell script in <i>filename</i>, converts it to use getopts instead of getopt, and writes the results on the standard output.</p> <p>So that all new commands will adhere to the command syntax standard described in intro(1), they should use getopts or getopt to parse positional parameters and check for options that are valid for that command.</p> <p>Examples:</p> <p>The following fragment of a shell program shows how one might process the arguments for a command that can take the options a or b, as well as the option o, which requires an option-argument:</p> <pre> while getopts abo: c do case \$c in a b) FLAG=\$c;; o) OARG=\$OPTARG;; \?) echo \$USAGE exit 2;; esac done </pre>

```
shift `expr $OPTIND -1`
```

This code accepts any of the following as equivalent:

```
cmd -a -b -o "xxx z yy" filename
cmd -a -b -o "xxx z yy" -- filename
cmd -ab -o xxx,z,yy filename
cmd -ab -o "xxx z yy" filename
cmd -o xxx,z,yy -b -a filename
```

getopts prints an error message on the standard error when it encounters an option letter not included in *optstring*.

Although the following command syntax rule (see **intro(1)**) relaxations are permitted under the current implementation, they should not be used because they may not be supported in future releases of the system. As in the EXAMPLES section above, **a** and **b** are options, and the option **o** requires an option-argument. The following example violates Rule 5: options with option-arguments must not be grouped with other options:

```
example% cmd -abxxx filename
```

The following example violates Rule 6: there must be white space after an option that takes an option-argument:

```
example% cmd -ab -oxxx filename
```

Changing the value of the shell variable **OPTIND** or parsing different sets of arguments may lead to unexpected results.

ksh

Checks *arg* for legal options. If *arg* is omitted, the positional parameters are used. An option argument begins with a + or a -. An option not beginning with + or - or the argument -- ends the options. *optstring* contains the letters that **getopts** recognizes. If a letter is followed by a :, that option is expected to have an argument. The options can be separated from the argument by blanks.

getopts places the next option letter it finds inside variable *name* each time it is invoked with a + prepended when *arg* begins with a +. The index of the next *arg* is stored in **OPTIND**. The option argument, if any, gets stored in **OPTARG**.

A leading : in *optstring* causes **getopts** to store the letter of an invalid option in **OPTARG**, and to set *name* to ? for an unknown option and to : when a required option is missing. Otherwise, **getopts** prints an error message. The exit status is non-zero when there are no more options.

For a further discussion of the Korn shell's **getopts** built-in command, see the previous discussion in the Bourne shell, **sh**, section of this manpage.

SEE ALSO

getopt(1), **getoptcv(1)**, **intro(1)**, **ksh(1)**, **sh(1)**

NAME	gettext – retrieve text string from message database
SYNOPSIS	gettext [<i>textdomain</i>] <i>msgid</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>gettext retrieves a translated text string corresponding to string <i>msgid</i> from a message object generated with msgfmt(1). The message object name is derived from the optional argument <i>textdomain</i> if present, otherwise from the TEXTDOMAIN environment. If no domain is specified, or if a corresponding string cannot be found, gettext prints <i>msgid</i>.</p> <p>Ordinarily gettext looks for its message object in /usr/lib/locale/lang/LC_MESSAGES where <i>lang</i> is the locale name. If present, the TEXTDOMAINDIR environment variable replaces the pathname component up to <i>lang</i>.</p> <p>This command interprets C escape sequences such as \t for tab. Use \\ to print a backslash. To produce a message on a line of its own, either put a \n at the end of <i>msgid</i>, or use this command in conjunction with printf(1).</p>
ENVIRONMENT	<p>LANG Specifies locale name.</p> <p>LC_MESSAGES Specifies messaging locale, and if present overrides LANG for messages.</p> <p>TEXTDOMAIN Specifies the text domain name, which is identical to the message object filename without .mo suffix.</p> <p>TEXTDOMAINDIR Specifies the pathname to the message database, and if present replaces /usr/lib/locale.</p>
SEE ALSO	msgfmt (1), printf (1), gettext (3I), setlocale (3C)
NOTES	This is the shell equivalent of the library routine gettext (3I).

NAME	gettext – retrieve a text string from a message database
SYNOPSIS	gettext <i>msgfile:msgnum</i> [<i>dflt_msg</i>]
AVAILABILITY	SUNWloc
DESCRIPTION	<p>gettext retrieves a text string from a message file in the directory /usr/lib/locale/locale/LC_MESSAGES. The directory name <i>locale</i> corresponds to the language in which the text strings are written; see setlocale(3C).</p> <p><i>msgfile</i> Name of the file in the directory /usr/lib/locale/locale/LC_MESSAGES to retrieve <i>msgnum</i> from. The name of <i>msgfile</i> can be up to 14 characters in length, but may not contain either <code>\0</code> (null) or the ASCII code for <code>/</code> (slash) or <code>:</code> (colon).</p> <p><i>msgnum</i> Sequence number of the string to retrieve from <i>msgfile</i>. The strings in <i>msgfile</i> are numbered sequentially from 1 to <i>n</i>, where <i>n</i> is the number of strings in the file.</p> <p><i>dflt_msg</i> Default string to be displayed if gettext fails to retrieve <i>msgnum</i> from <i>msgfile</i>. Nongraphic characters must be represented as alphabetic escape sequences.</p> <p>The text string to be retrieved is in the file <i>msgfile</i>, created by the mkmsgs(1) utility and installed under the directory /usr/lib/locale/locale/LC_MESSAGES. You control which directory is searched by setting the environment variable LC_MESSAGES. If LC_MESSAGES is not set, the environment variable LANG will be used. If LANG is not set, the files containing the strings are under the directory /usr/lib/locale/C/LC_MESSAGES.</p> <p>If gettext fails to retrieve a message in the requested language, it will try to retrieve the same message from /usr/lib/locale/C/LC_MESSAGES/msgfile. If this also fails, and if <i>dflt_msg</i> is present and non-null, then it will display the value of <i>dflt_msg</i>; if <i>dflt_msg</i> is not present or is null, then it will display the string Message not found!!.</p>
EXAMPLES	<p>If the environment variables LANG or LC_MESSAGES have not been set to other than their default values, the following example:</p> <pre style="margin-left: 40px;">example% gettext UX:10 "hello world\n"</pre> <p>will try to retrieve the 10th message from /usr/lib/locale/C/UX/msgfile. If the retrieval fails, the message "hello world," followed by a newline, will be displayed.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of gettext for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how gettext behaves.</p>

NAME	glob – shell built-in function to expand a word list
SYNOPSIS csh	glob <i>wordlist</i>
DESCRIPTION csh	glob performs filename expansion on <i>wordlist</i> . Like echo(1) , but no \ escapes are recognized. Words are delimited by NULL characters in the output.
SEE ALSO	csh(1) , echo(1)

NAME	gprof – display call-graph profile data
SYNOPSIS	gprof [-abcDsz] [-e <i>function-name</i>] [-E <i>function-name</i>] [-f <i>function-name</i>] [-F <i>function-name</i>] [<i>image-file</i> [<i>profile-file</i> . . .]] [-n <i>number of functions</i>]
DESCRIPTION	<p>gprof produces an execution profile of a program. The effect of called routines is incorporated in the profile of each caller. The profile data is taken from the call graph profile file which is created by programs compiled with the -xpg option of cc(1B), or -pg for other compilers. These options also link in versions of the library routines which are compiled for profiling. The symbol table in the executable image file <i>image-file</i> (a.out by default) is read and correlated with the call graph profile file <i>profile-file</i> (gmon.out by default). If more than one profile file is specified, the gprof output shows the sum of the profile information in the given profile files.</p> <p>First, execution times for each routine are propagated along the edges of the call graph. Cycles are discovered, and calls into a cycle are made to share the time of the cycle. The first listing shows the functions sorted according to the time they represent, including the time of their call graph descendants. Below each function entry is shown its (direct) call-graph children, and how their times are propagated to this function. A similar display above the function shows how this function's time and the time of its descendants is propagated to its (direct) call-graph parents.</p> <p>Cycles are also shown, with an entry for the cycle as a whole and a listing of the members of the cycle and their contributions to the time and call counts of the cycle.</p> <p>Next, a flat profile is given, similar to that provided by prof(1). This listing gives the total execution times and call counts for each of the functions in the program, sorted by decreasing time. Finally, an index is given, showing the correspondence between function names and call-graph profile index numbers.</p> <p>A single function may be split into subfunctions for profiling by means of the MARK macro (see prof(5)).</p> <p>Beware of quantization errors. The granularity of the sampling is shown, but remains statistical at best. It is assumed that the time for each execution of a function can be expressed by the total time for the function divided by the number of times the function is called. Thus the time propagated along the call-graph arcs to parents of that function is directly proportional to the number of times that arc is traversed.</p> <p>The profiled program must call exit(2) or return normally for the profiling information to be saved in the gmon.out file.</p>
OPTIONS	<p>-a Suppress printing statically declared functions. If this option is given, all relevant information about the static function (for instance, time samples, calls to other functions, calls from other functions) belongs to the function loaded just before the static function in the a.out file.</p> <p>-b Brief. Suppress descriptions of each field in the profile.</p> <p>-C Demangle C++ symbol names before printing them out.</p>

- c** Discover the static call-graph of the program by a heuristic which examines the text space of the object file. Static-only parents or children are indicated with call counts of 0.
- D** Produce a profile file **gmon.sum** that represents the difference of the profile information in all specified profile files. This summary profile file may be given to subsequent executions of **gprof** (also with **-D**) to summarize profile data across several runs of an **a.out** file. (See also the **-s** option.)

As an example, suppose function A calls function B **n** times in profile file **gmon.sum**, and **m** times in profile file **gmon.out**. With **-D**, a new **gmon.sum** file will be created showing the number of calls from A to B as **n-m**.
- E function-name**

Suppress printing the graph profile entry for routine *function-name* (and its descendants) as **-e**, below, and also exclude the time spent in *function-name* (and its descendants) from the total and percentage time computations. More than one **-E** option may be given. For example:

‘**-E mcount -E mcleanup**’

is the default.
- e function-name**

Suppress printing the graph profile entry for routine *function-name* and all its descendants (unless they have other ancestors that are not suppressed). More than one **-e** option may be given. Only one *function-name* may be given with each **-e** option.
- F function-name**

Print the graph profile entry only for routine *function-name* and its descendants (as **-f**, below) and also use only the times of the printed routines in total time and percentage computations. More than one **-F** option may be given. Only one *function-name* may be given with each **-F** option. The **-F** option overrides the **-E** option.
- f function-name**

Print the graph profile entry only for routine *function-name* and its descendants. More than one **-f** option may be given. Only one *function-name* may be given with each **-f** option.
- n** Limits the size of flat and graph profile listings to the top **n** offending functions.
- s** Produce a profile file **gmon.sum** which represents the sum of the profile information in all of the specified profile files. This summary profile file may be given to subsequent executions of **gprof** (also with **-s**) to accumulate profile data across several runs of an **a.out** file. (See also the **-D** option.)
- z** Display routines which have zero usage (as indicated by call counts and accumulated time). This is useful in conjunction with the **-c** option for discovering which routines were never called.

ENVIRONMENT	PROFDIR If this environment variable contains a value, place profiling output within that directory, in a file named <i>pid.programname</i> . <i>pid</i> is the process ID, and <i>programname</i> is the name of the program being profiled, as determined by removing any path prefix from the argv[0] with which the program was called. If the variable contains a NULL value, no profiling output is produced. Otherwise, profiling output is placed in the file gmon.out .
FILES	a.out executable file containing namelist gmon.out dynamic call-graph and profile gmon.sum summarized dynamic call-graph and profile SPROFDIR/pid.programname
SEE ALSO	cc(1B) , prof(1) , exit(2) , profil(2) , monitor(3C) , prof(5) Graham, S.L., Kessler, P.B., McKusick, M.K., 'gprof: A Call Graph Execution Profiler', <i>Proceedings of the SIGPLAN '82 Symposium on Compiler Construction</i> , SIGPLAN Notices, Vol. 17, No. 6, pp. 120-126, June 1982.
BUGS	Parents which are not themselves profiled will have the time of their profiled children propagated to them, but they will appear to be spontaneously invoked in the call-graph listing, and will not have their time propagated further. Similarly, signal catchers, even though profiled, will appear to be spontaneous (although for more obscure reasons). Any profiled children of signal catchers should have their times propagated properly, unless the signal catcher was invoked during the execution of the profiling routine, in which case all is lost.

NAME	graph – draw a graph
SYNOPSIS	graph [<i>-a spacing</i> [<i>start</i>]] [<i>-b</i>] [<i>-c string</i>] [<i>-g gridstyle</i>] [<i>-l label</i>] [<i>-m connectmode</i>] [<i>-s</i>] [<i>-x [1] lower [upper [spacing]]</i>] [<i>-y [1] lower [upper [spacing]]</i>] [<i>-h fraction</i>] [<i>-w fraction</i>] [<i>-r fraction</i>] [<i>-u fraction</i>] [<i>-t</i>] . . .
AVAILABILITY	SUNWesu
DESCRIPTION	graph with no options takes pairs of numbers from the standard input as abscissae and ordinates of a graph. Successive points are connected by straight lines. The standard output from graph contains plotting instructions suitable for input to plot (1B) or to the command lpr -g (see lpr (1B)). If the coordinates of a point are followed by a nonnumeric string, that string is printed as a label beginning on the point. Labels may be surrounded with quotes "...", in which case they may be empty or contain blanks and numbers; labels never contain NEWLINE characters. A legend indicating grid range is produced with a grid unless the <i>-s</i> option is present.
OPTIONS	Each option is recognized as a separate argument. If a specified lower limit exceeds the upper limit, the axis is reversed. <i>-a spacing</i> [<i>start</i>] Supply abscissae automatically (they are missing from the input); <i>spacing</i> is the spacing (default 1). <i>start</i> is the starting point for automatic abscissae (default 0 or lower limit given by <i>-x</i>). <i>-b</i> Break (disconnect) the graph after each label in the input. <i>-c string</i> <i>String</i> is the default label for each point. <i>-g gridstyle</i> <i>Gridstyle</i> is the grid style: 0 no grid, 1 frame with ticks, 2 full grid (default). <i>-l label</i> <i>label</i> is label for graph. <i>-m connectmode</i> Mode (style) of connecting lines: 0 disconnected, 1 connected (default). Some devices give distinguishable line styles for other small integers. <i>-s</i> Save screen, do not erase before plotting. <i>-x [1] lower [upper [spacing]]</i> If 1 is present, x axis is logarithmic. <i>lower</i> and <i>upper</i> are lower (and upper) x limits. <i>spacing</i> , if present, is grid spacing on x axis. Normally these quantities are determined automatically. <i>-y [1] lower [upper [spacing]]</i> If 1 is present, y axis is logarithmic. <i>lower</i> and <i>upper</i> are lower (and upper) y limits. <i>spacing</i> , if present, is grid spacing on y axis. Normally these quantities are determined automatically.

-h *fraction* *fraction* of space for height.
-w *fraction* *fraction* of space for width.
-r *fraction* *fraction* of space to move right before plotting.
-u *fraction* *fraction* of space to move up before plotting.
-t Transpose horizontal and vertical axes. Option **-x** now applies to the vertical axis.

SEE ALSO **lpr(1B), plot(1B), spline(1)**

BUGS **graph** stores all points internally and drops those for which there is no room. Segments that run out of bounds are dropped, not windowed. Logarithmic axes may not be reversed.

NAME	grep – search a file for a pattern
SYNOPSIS	grep [-bchilnsvw] <i>limited-regular-expression</i> [<i>filename ...</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>grep searches files for a pattern and prints all lines that contain that pattern. grep uses limited regular expressions (expressions that have string values that use a subset of the possible alphanumeric and special characters) like those used with ed(1) to match the patterns. It uses a compact non-deterministic algorithm.</p> <p>Be careful using the characters \$, *, [, ^, , (,), and \ in the <i>limited-regular-expression</i> because they are also meaningful to the shell. It is safest to enclose the entire <i>limited-regular-expression</i> in single quotes '...'. </p> <p>If no files are specified, grep assumes standard input. Normally, each line found is copied to standard output. The file name is printed before each line found if there is more than one input file.</p>
OPTIONS	<p>-b Precede each line by the block number on which it was found. This can be useful in locating block numbers by context (first block is 0).</p> <p>-c Print only a count of the lines that contain the pattern.</p> <p>-h Prevents the name of the file containing the matching line from being appended to that line. Used when searching multiple files.</p> <p>-i Ignore upper/lower case distinction during comparisons.</p> <p>-l Print only the names of files with matching lines, separated by NEWLINE characters. Does not repeat the names of files when the pattern is found more than once.</p> <p>-n Precede each line by its line number in the file (first line is 1).</p> <p>-s Suppress error messages about nonexistent or unreadable files</p> <p>-v Print all lines except those that contain the pattern.</p> <p>-w Search for the expression as a word as if surrounded by \< and \>.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of grep for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how grep behaves.</p>
SEE ALSO	ed (1), egrep (1), fgrep (1), sed (1), sh (1), environ (5)

DIAGNOSTICS Exit status is 0 if any matches are found, 1 if none, 2 for syntax errors or inaccessible files (even if matches were found).

NOTES Lines are limited to BUFSIZ characters; longer lines are truncated. BUFSIZ is defined in `/usr/include/stdio.h`. If there is a line with embedded nulls, **grep** will only match up to the first null; if it matches, it will print the entire line.

NAME	groups – print group membership of user
SYNOPSIS	groups [<i>user</i> . . .]
AVAILABILITY	SUNWcsu
DESCRIPTION	The command groups prints on standard output the groups to which you or the optionally specified user belong. Each user belongs to a group specified in /etc/passwd and possibly to other groups as specified in /etc/group . Note that /etc/passwd specifies the numerical ID (gid) of the group. The groups command converts gid to the group name in the output.
EXAMPLE	The output takes the following form: <pre>example% groups tester01 tester02 tester01 : staff tester02 : staff example%</pre>
FILES	/etc/passwd /etc/group
SEE ALSO	group(4) , passwd(4)

NAME	groups – display a user’s group memberships
SYNOPSIS	<code>/usr/ucb/groups [user ...]</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	With no arguments, groups displays the groups to which you belong; else it displays the groups to which the user belongs. Each user belongs to a group specified in the password file <code>/etc/passwd</code> and possibly to other groups as specified in the file <code>/etc/group</code> . If you do not own a file but belong to the group which it is owned by then you are granted group access to the file.
FILES	<code>/etc/passwd</code> <code>/etc/group</code>
SEE ALSO	<code>getgroups(2)</code>
NOTES	This command is obsolete.

NAME	grpck – check group database entries
SYNOPSIS	<code>/usr/etc/grpck [filename]</code>
DESCRIPTION	grpck checks that a file in group(4) does not contain any errors; it checks the /etc/group file by default.
FILES	/etc/group
SEE ALSO	groups(1) , group(4) , passwd(4)
DIAGNOSTICS	Too many/few fields An entry in the group file does not have the proper number of fields. No group name The group name field of an entry is empty. Bad character(s) in group name The group name in an entry contains characters other than lower-case letters and digits. Invalid GID The group ID field in an entry is not numeric or is greater than 65535. Null login name A login name in the list of login names in an entry is null. Logname not found in password file A login name in the list of login names in an entry is not in the password file. Line too long A line (including the newline character) in the group file exceeds the maximum length of 512 characters. Duplicate logname entry A login name appears more than once in the list of login names for a group file entry. Out of memory The program cannot allocate memory in order to continue. Maximum groups exceeded for logname A login name's group membership exceeds the maximum, NGROUPS_MAX .

NAME	hash, rehash, unhash, hashstat – shell built-in functions to evaluate the internal hash table of the contents of directories
SYNOPSIS	
sh	hash [-r] [<i>name</i> ...]
csh	rehash unhash hashstat
ksh	hash
DESCRIPTION	
sh	For each <i>name</i> , the location in the search path of the command specified by <i>name</i> is determined and remembered by the shell. The -r option to the hash built-in causes the shell to forget all remembered locations. If no arguments are given, hash provides information about remembered commands. <i>Hits</i> is the number of times a command has been invoked by the shell process. <i>Cost</i> is a measure of the work required to locate a command in the search path. If a command is found in a "relative" directory in the search path, after changing to that directory, the stored location of that command is recalculated. Commands for which this will be done are indicated by an asterisk (*) adjacent to the Hits information. <i>Cost</i> will be incremented when the recalculation is done.
csh	rehash recomputes the internal hash table of the contents of directories listed in the path environmental variable to account for new commands added. unhash disables the internal hash table. hashstat prints a statistics line indicating how effective the internal hash table has been at locating commands (and avoiding execs). An exec is attempted for each component of the <i>path</i> where the hash function indicates a possible hit and in each component that does not begin with a '/.
ksh	hash provides information about remembered commands. In the first column of the output, Hits , is the number of times a command has been invoked by the shell process. In the second column, Cost , is a measure of the work required to locate a command in the search path. The third column lists the command. If a command is found in a "relative" directory in the search path, after changing to that directory, the stored location of that command is recalculated. Commands for which this will be done are indicated by an asterisk (*) adjacent to the Hits information. Cost will be incremented when the recalculation is done.
SEE ALSO	csh(1) , ksh(1) , sh(1)

NAME	head – display first few lines of files
SYNOPSIS	head [<i>-n</i>] [<i>filename...</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	head copies the first <i>n</i> lines of each <i>filename</i> to the standard output. If no <i>filename</i> is given, head copies lines from the standard input. The default value of <i>n</i> is 10 lines. When more than one file is specified, the start of each file will look like: ==> <i>filename</i> <== Thus, a common way to display a set of short files, identifying each one, is: example% head -9999 filename1 filename2 ...
SEE ALSO	cat(1), more(1), pg(1), tail(1)

NAME	history, fc – shell built-in functions to re-use previous command-lines from the current shell
SYNOPSIS	
cs h	history [-hr] [<i>n</i>]
k sh	fc [-e <i>ename</i>] -nlr [<i>first</i> [<i>last</i>]] fc -e - [<i>old=new</i>] [<i>command</i>]
DESCRIPTION	
cs h	Display the history list; if <i>n</i> is given, display only the <i>n</i> most recent events. -r Reverse the order of printout to be most recent first rather than oldest first. -h Display the history list without leading numbers. This is used to produce files suitable for sourcing using the -h option to the cs h built-in command, source (1).
k sh	fc , in the first form, a range of commands from <i>first</i> to <i>last</i> is selected from the last HIST-SIZE commands that were typed at the terminal. The arguments <i>first</i> and <i>last</i> may be specified as a number or as a string. A string is used to locate the most recent command starting with the given string. A negative number is used as an offset to the current command number. If the -l flag is selected, the commands are listed on standard output. Otherwise, the editor program -e name is invoked on a file containing these keyboard commands. If <i>ename</i> is not supplied, then the value of the variable FCEDIT (default /bin/ed) is used as the editor. When editing is complete, the edited command(s) is executed. If <i>last</i> is not specified then it will be set to <i>first</i> . If <i>first</i> is not specified the default is the previous command for editing and -16 for listing. The flag -r reverses the order of the commands and the flag -n suppresses command numbers when listing. In the second form the <i>command</i> is re-executed after the substitution <i>old=new</i> is performed.
SEE ALSO	cs h(1), k sh(1), sh (1), source (1)

NAME	hostid – print the numeric identifier of the current host
SYNOPSIS	<code>/usr/ucb/hostid</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	The hostid command prints the identifier of the current host in hexadecimal. This numeric value is likely to differ when hostid is run on a different machine.
SEE ALSO	sysinfo(2) , gethostid(3B)

NAME	hostname – set or print name of current host system
SYNOPSIS	<code>/usr/ucb/hostname</code> [<i>name-of-host</i>]
AVAILABILITY	SUNWscpu
DESCRIPTION	The hostname command prints the name of the current host, as given before the login prompt. The super-user can set the hostname by giving an argument.
SEE ALSO	uname(1)

NAME	iconv – code set conversion utility
SYNOPSIS	iconv <i>-f fromcode -t tocode</i> [<i>filename</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>iconv converts the characters or sequences of characters in <i>filename</i> from one code set to another and writes the results to standard output. Should no conversion exist for a particular character then it is converted to the underscore '_' in the target codeset.</p> <p>The required arguments <i>fromcode</i> and <i>toctype</i> identify the input and output code sets, respectively. If no <i>filename</i> argument is specified on the command line, iconv reads the standard input.</p> <p>iconv will always convert to or from the ISO 8859-1 Latin alphabet No.1, from or to an ISO 646 ASCII variant codeset for a particular language. The ISO 8859-1 codeset will support the majority of 8 bit codesets. The conversions attempted by iconv accommodate the most commonly used languages.</p> <p>The following table lists the supported conversions.</p>

Code Set Conversions Supported				
Code	Symbol	Target Code	Symbol	comment
ISO 646	646	ISO 8859-1	8859	US ASCII
ISO 646de	646de	ISO 8859-1	8859	German
ISO 646da	646da	ISO 8859-1	8859	Danish
ISO 646en	646en	ISO 8859-1	8859	English ASCII
ISO 646es	646es	ISO 8859-1	8859	Spanish
ISO 646fr	646fr	ISO 8859-1	8859	French
ISO 646it	646it	ISO 8859-1	8859	Italian
ISO 646sv	646sv	ISO 8859-1	8859	Swedish
ISO 8859-1	8859	ISO 646	646	7 bit ASCII
ISO 8859-1	8859	ISO 646de	646de	German
ISO 8859-1	8859	ISO 646da	646da	Danish
ISO 8859-1	8859	ISO 646en	646en	English ASCII
ISO 8859-1	8859	ISO 646es	646es	Spanish
ISO 8859-1	8859	ISO 646fr	646fr	French
ISO 8859-1	8859	ISO 646it	646it	Italian
ISO 8859-1	8859	ISO 646sv	646sv	Swedish

The conversions are performed according to the tables found on **iconv(5)**.

OPTIONS

-f *fromcode* Identifies the input code set.
-t *to*code Identifies the output code set.

EXAMPLES

The following converts the contents of file **mail1** from code set **8859** to **646fr** and stores the results in file *mail.local*.

```
example% iconv -f 8859 -t 646fr mail1 > mail.local
```

FILES

/usr/lib/iconv/iconv_data lists the conversions supported
/usr/lib/iconv/*.t conversion tables

SEE ALSO

iconv(5)

DIAGNOSTICS

iconv returns 0 upon successful completion, 1 otherwise.

NAME	if, test – shell built-in functions to evaluate expression(s) or to make execution of actions dependent upon the evaluation of expression(s)
SYNOPSIS	
sh	if <i>condition</i> ; then <i>action</i> ; [elif <i>condition</i> ; then <i>action</i>] ... [else <i>action</i>]; fi test <i>expression</i> [<i>expression</i>]
csh	if (<i>expr</i>) <i>command</i> if (<i>expr</i>) then <i>commands</i> ... else if (<i>expr2</i>) then <i>commands</i> ... else <i>commands</i> ... endif
ksh	if <i>condition</i> ; then <i>action</i> [[; elif <i>condition</i> ; then <i>action</i>] ...]; else <i>action</i> ; fi test <i>expression</i> [<i>expression</i>]
DESCRIPTION	
sh	The <i>condition</i> following if is executed and, if it returns a zero exit status, the <i>action</i> following the first then is executed. Otherwise, the <i>condition</i> following elif is executed and, if its value is zero, the <i>action</i> following the next then is executed. Failing that, the else <i>action</i> is executed. If no else <i>action</i> or then <i>action</i> is executed, then the if command returns a zero exit status. test evaluates the expression <i>expression</i> and, if its value is true, sets a zero (true) exit status; otherwise, a non-zero (false) exit status is set; test also sets a non-zero exit status if there are no arguments. When permissions are tested, the effective user ID of the process is used. All operators, flags, and brackets (brackets used as shown in the second SYNOPSIS line) must be separate arguments to the test command; normally these items are separated by spaces. Primitives: The following primitives are used to construct <i>expression</i> : -r <i>filename</i> True if <i>filename</i> exists and is readable. -w <i>filename</i> True if <i>filename</i> exists and is writable. -x <i>filename</i> True if <i>filename</i> exists and is executable. -f <i>filename</i> True if <i>filename</i> exists and is a regular file. Alternatively, if /usr/bin/sh users specify /usr/ucb before /usr/bin in their PATH environment

	variable, then test will return true if <i>filename</i> exists and is (not-a-directory). This is also the default for /usr/bin/csh users.
-d filename	True if <i>filename</i> exists and is a directory.
-h filename	True if <i>filename</i> exists and is a symbolic link. With all other primitives (except -L filename), the symbolic links are followed by default.
-c filename	True if <i>filename</i> exists and is a character special file.
-b filename	True if <i>filename</i> exists and is a block special file.
-p filename	True if <i>filename</i> exists and is a named pipe (fifo).
-u filename	True if <i>filename</i> exists and its set-user-ID bit is set.
-g filename	True if <i>filename</i> exists and its set-group-ID bit is set.
-k filename	True if <i>filename</i> exists and its sticky bit is set.
-s filename	True if <i>filename</i> exists and has a size greater than zero.
-t [fildes]	True if the open file whose file descriptor number is <i>fildes</i> (1 by default) is associated with a terminal device.
-z s1	True if the length of string <i>s1</i> is zero.
-n s1	True if the length of the string <i>s1</i> is non-zero.
s1 = s2	True if strings <i>s1</i> and <i>s2</i> are identical.
s1 != s2	True if strings <i>s1</i> and <i>s2</i> are <i>not</i> identical.
s1	True if <i>s1</i> is <i>not</i> the null string.
n1 -eq n2	True if the integers <i>n1</i> and <i>n2</i> are algebraically equal. Any of the comparisons -ne , -gt , -ge , -lt , and -le may be used in place of -eq .
-L filename	True if <i>filename</i> exists and is a symbolic link. With all other primitives (except -h filename), the symbolic links are followed by default.

Operators:

These primaries may be combined with the following operators:

!	Unary negation operator.
-a	Binary <i>and</i> operator.
-o	Binary <i>or</i> operator (-a has higher precedence than -o).
(expression)	Parentheses for grouping. Notice also that parentheses are meaningful to the shell and, therefore, must be quoted.

The **not-a-directory** alternative to the **-f** option is a transition aid for BSD applications and may not be supported in future releases.

The **-L** option is a migration aid for users of other shells which have similar options and may not be supported in future releases.

If you test a file you own (the `-r` `-w` or `-x` tests), but the permission tested does not have the *owner* bit set, a non-zero (false) exit status will be returned even though the file may have the **group** or *other* bit set for that permission. The correct exit status will be set if you are super-user.

The `=` and `!=` operators have a higher precedence than the `-r` through `-n` operators, and `=` and `!=` always expect arguments; therefore, `=` and `!=` cannot be used with the `-r` through `-n` operators.

If more than one argument follows the `-r` through `-n` operators, only the first argument is examined; the others are ignored, unless a `-a` or a `-o` is the second argument.

csh With the one-line form of `if`, if the specified expression evaluates to true, the single *command* with arguments is executed. Variable substitution on *command* happens early, at the same time it does for the rest of the `if` command. *command* must be a simple command, not a pipeline, a command list, or a parenthesized command list. Note: I/O redirection occurs even if *expr* is false, when *command* is *not* executed (this is a bug).

With the multi-line form of `if`, if *expr* is true, *commands* up to the first **else** are executed. Otherwise, if *expr2* is true, the *commands* between the **else if** and the second **else** are executed. Otherwise, *commands* between the **else** and the **endif** are executed. Any number of **else if** pairs are allowed, but only one **else**. Only one **endif** is needed, but it is required. The words **else** and **endif** must be the first nonwhite characters on a line. The **if** must appear alone on its input line or after an **else**.

ksh The *condition* following `if` is executed and, if it returns an exit status of zero, the *action* following the first **then** is executed. Otherwise, the *condition* following **elif** is executed and, if its value is zero, the *action* following the next **then** is executed. Failing that, the **else action** is executed. If no **else action** or **then action** is executed, then the `if` command returns a zero exit status.

For a description of the `test` built-in, see the (**sh**) Bourne shell's `test` built-in above.

EXAMPLES

In the `if` command examples, three conditions are tested, and if all three evaluate as true or successful, then their validities are printed to the screen. The two forms of the `test` built-in follow the Bourne shell's `if` example.

```
sh      ZERO=0 ONE=1 TWO=2 ROOT=root
        if [ $ONE -gt $ZERO ] &&
          [ $TWO -eq 2 ] &&
          [ 'grep $ROOT /etc/passwd >&1 /dev/null' ]
        then
          echo "$ONE is greater than zero, $TWO equals 2, and root is a user-name
            in the password file"
        else
          echo "At least one of the three test conditions is false"
        fi
```

Examples of the **test** built-in:

```
test 'grep $ROOT /etc/passwd >&1 /dev/null'
echo $? /* test for success */
```

```
[ 'grep nosuchname /etc/passwd >&1 /dev/null' ]
echo $? /* test for failure */
```

```
csh @ ZERO = 0; @ ONE = 1; @ TWO = 2; set ROOT = root
grep $ROOT /etc/passwd >&1 /dev/null
if ( "$status" == "0" && $ONE > $ZERO && $TWO == 2 ) then
    /* $status must be tested for immediately following grep */
    echo "$ONE is greater than zero, $TWO equals 2, and root is a user-name
        in the password file"
endif
```

```
ksh ZERO=0 ONE=1 TWO=2 ROOT=root
if [ $ONE -gt $ZERO ]; [ $TWO -eq 2 ];
[ 'grep $ROOT /etc/passwd >&1 /dev/null' ];
then
    echo "$ONE is greater than zero, $TWO equals 2, and root is a user-name
        in the password file"
fi
```

The Korn shell will also accept the syntax of both the **if** command and the **test** command of the Bourne shell.

When using the brackets ([]) within **if** commands, you must separate both ends of both brackets from other characters with a space.

SEE ALSO [csh\(1\)](#), [ksh\(1\)](#), [sh\(1\)](#), [test\(1B\)](#)

NOTES Both the Bourne shell, **sh**, and the Korn shell, **ksh**, can use the semicolon and the carriage return interchangeably in their syntax of the **if**, **for**, and **while** built-in commands.

NAME	indicator – display application specific alarms and/or the "working" indicator
SYNOPSIS	indicator [-b [<i>n</i>]] [-c <i>column</i>] [-l <i>length</i>] [-o] [-w] [<i>string</i> ...]
DESCRIPTION	The indicator function displays application specific alarms or the "working" indicator, or both, on the FMLI banner line. By default, indicator ????. The argument <i>string</i> is a string to be displayed on the banner line, and should always be the last argument given. Note that <i>string</i> is not automatically cleared from the banner line.
OPTIONS	<p>-bn The -b option rings the terminal bell <i>n</i> times, where <i>n</i> is an integer from 1 to 10. The default value is 1. If the terminal has no bell, the screen is flashed instead, if possible.</p> <p>-c <i>column</i> The -c option defines the column of the banner line at which to start the indicator string. The argument <i>column</i> must be an integer from 0 to DISPLAYW-1. If the -c option is not used, <i>column</i> defaults to 0.</p> <p>-l <i>length</i> The -l option defines the maximum length of the string displayed. If <i>string</i> is longer than <i>length</i> characters, it will be truncated. The argument <i>length</i> must be an integer from 1 to DISPLAYW. If the -l option is not used, <i>length</i> defaults to DISPLAYW. NOTE: if <i>string</i> doesn't fit it will be truncated.</p> <p>-o The -o option causes indicator to duplicate its output to <i>stdout</i>.</p> <p>-w The -w option turns on the "working" indicator.</p>
EXAMPLES	<p>When the value entered in a form field is invalid, the following use of indicator will ring the bell three times and display the word WRONG starting at column 1 of the banner line.</p> <pre>invalidmsg=`indicator -b 3 -c 1 "WRONG"`</pre> <p>To clear the indicator after telling the user the entry is wrong:</p> <pre>invalidmsg=`indicator -b 9 -c 1 "WRONG"; sleep 3; indicator -c 1 " "`</pre> <p>In this example the value of invalidmsg (in this case the default value Input is not valid), still appears on the FMLI message line.</p>

NAME	indxbib – create an inverted index to a bibliographic database								
SYNOPSIS	indxbib <i>database-file</i> ...								
AVAILABILITY	SUNWdoc								
DESCRIPTION	<p>indxbib makes an inverted index to the named <i>database-file</i> (which must reside within the current directory), typically for use by lookbib(1) and refer(1). A <i>database</i> contains bibliographic references (or other kinds of information) separated by blank lines.</p> <p>A bibliographic reference is a set of lines, constituting fields of bibliographic information. Each field starts on a line beginning with a '%', followed by a key-letter, then a blank, and finally the contents of the field, which may continue until the next line starting with '%'. indxbib is a shell script that calls two programs: /usr/lib/refer/mkey and /usr/lib/refer/inv. mkey truncates words to 6 characters, and maps upper case to lower case. It also discards words shorter than 3 characters, words among the 100 most common English words, and numbers (dates) < 1000 or > 2099. These parameters can be changed.</p> <p>indxbib creates an entry file (with a .ia suffix), a posting file (.ib), and a tag file (.ic), in the working directory.</p>								
FILES	<p>/usr/lib/refer/mkey /usr/lib/refer/inv</p> <table border="0"> <tr> <td>x.ia</td> <td>entry file</td> </tr> <tr> <td>x.ib</td> <td>posting file</td> </tr> <tr> <td>x.ic</td> <td>tag file</td> </tr> <tr> <td>x.ig</td> <td>reference file</td> </tr> </table>	x.ia	entry file	x.ib	posting file	x.ic	tag file	x.ig	reference file
x.ia	entry file								
x.ib	posting file								
x.ic	tag file								
x.ig	reference file								
SEE ALSO	addbib (1), lookbib (1), refer (1), roffbib (1), sortbib (1)								
BUGS	<p>All dates should probably be indexed, since many disciplines refer to literature written in the 1800s or earlier.</p> <p>indxbib does not recognize pathnames.</p>								

NAME	install – install files
SYNOPSIS	<pre> /usr/ucb/install [-cs] [-g group] [-m mode] [-o owner] filename1 filename2 /usr/ucb/install [-cs] [-g group] [-m mode] [-o owner] filename ... directory /usr/ucb/install -d [-g group] [-m mode] [-o owner] directory </pre>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>Install is used within makefiles to copy new versions of files into a destination directory and to create the destination directory itself.</p> <p>The first two forms are similar to the cp(1) command with the addition that executable files can be stripped during the copy and the owner, group, and mode of the installed file(s) can be given.</p> <p>The third form can be used to create a destination directory with the required owner, group and permissions.</p> <p>Note: install uses no special privileges to copy files from one place to another. The implications of this are:</p> <ul style="list-style-type: none"> • You must have permission to read the files to be installed. • You must have permission to copy into the destination file or directory. • You must have permission to change the modes on the final copy of the file if you want to use the -m option to change modes. • You must be superuser if you want to specify the ownership of the installed file with -o. If you are not the super-user, or if -o is not in effect, the installed file will be owned by you, regardless of who owns the original.
OPTIONS	<p>-c Copy files. In fact install <i>always</i> copies files, but the -c option is retained for backwards compatibility with old shell scripts that might otherwise break.</p> <p>-d Create a directory. Missing parent directories are created as required as in mkdir -p. If the directory already exists, the owner, group and mode will be set to the values given on the command line.</p> <p>-s Strip executable files as they are copied.</p> <p>-g group Set the group ownership of the installed file or directory. (staff by default.)</p> <p>-m mode Set the mode for the installed file or directory. (0755 by default.)</p> <p>-o owner If run as root, set the ownership of the installed file to the user-ID of <i>owner</i>.</p>
SEE ALSO	chgrp(1) , chmod(1) , chown(1) , cp(1) , mkdir(1) , strip(1) , install(1M)

NAME	ipcrm – remove a message queue, semaphore set, or shared memory ID
SYNOPSIS	ipcrm [-m <i>shmid</i>] [-q <i>msqid</i>] [-s <i>semid</i>] [-M <i>shmkey</i>] [-Q <i>msgkey</i>] [-S <i>semkey</i>]
AVAILABILITY	SUNWipc
DESCRIPTION	ipcrm removes one or more messages, semaphores, or shared memory identifiers.
OPTIONS	<p>The identifiers are specified by the following <i>options</i> :</p> <p>-m <i>shmid</i> Remove the shared memory identifier <i>shmid</i> from the system. The shared memory segment and data structure associated with it are destroyed after the last detach.</p> <p>-q <i>msqid</i> Remove the message queue identifier <i>msqid</i> from the system and destroy the message queue and data structure associated with it.</p> <p>-s <i>semid</i> Remove the semaphore identifier <i>semid</i> from the system and destroy the set of semaphores and data structure associated with it.</p> <p>-M <i>shmkey</i> Removes the shared memory identifier, created with key <i>shmkey</i> , from the system. The shared memory segment and data structure associated with it are destroyed after the last detach.</p> <p>-Q <i>msgkey</i> Remove the message queue identifier, created with key <i>msgkey</i> , from the system and destroy the message queue and data structure associated with it.</p> <p>-S <i>semkey</i> Remove the semaphore identifier, created with key <i>semkey</i> , from the system and destroy the set of semaphores and data structure associated with it.</p> <p>The details of the removes are described in msgctl(2), shmctl(2), and semctl(2). Use the ipcs command to find the identifiers and keys.</p>
SEE ALSO	ipcs(1) , msgctl(2) , msgget(2) , msgop(2) , semctl(2) , semget(2) , semop(2) , shmctl(2) , shmget(2) , shmop(2)

NAME	ipcs – report inter-process communication facilities status
SYNOPSIS	ipcs [-abcmopqst] [-C <i>corefile</i>] [-N <i>namelist</i>]
AVAILABILITY	SUNWipc
DESCRIPTION	<p>ipcs prints information about active inter-process communication facilities. Without <i>options</i>, information is printed in short format for message queues, shared memory, and semaphores that are currently active in the system.</p> <p>The information that is displayed is controlled by the options supplied.</p>
OPTIONS	<p>-m Print information about active shared memory segments.</p> <p>-q Print information about active message queues.</p> <p>-s Print information about active semaphores.</p> <p>If -q, -m, or -s are specified, information about only those indicated is printed. If none of these three are specified, information about all three is printed subject to these options:</p> <p>-a Use all print options. (This is a shorthand notation for -b, -c, -o, -p, and -t.)</p> <p>-b Print information on biggest allowable size: maximum number of bytes in messages on queue for message queues, size of segments for shared memory, and number of semaphores in each set for semaphores. See below for meaning of columns in a listing.</p> <p>-c Print creator's login name and group name. See below.</p> <p>-o Print information on outstanding usage: number of messages on queue and total number of bytes in messages on queue for message queues and number of processes attached to shared memory segments.</p> <p>-p Print process number information: process ID of last process to send a message, process ID of last process to receive a message on message queues, process ID of creating process, and process ID of last process to attach or detach on shared memory segments. See below.</p> <p>-t Print time information: time of the last control operation that changed the access permissions for all facilities, time of last msgsnd and last msgrcv on message queues, time of last shmat and last shmdt on shared memory, time of last semop on semaphores. See below.</p> <p>-C <i>corefile</i> Use the file <i>corefile</i> in place of /dev/mem and /dev/kmem. Use a core dump obtained from savecore(1M) in place of /dev/mem and /dev/kmem. Without the -C option (default), the running system image is used.</p> <p>-N <i>namelist</i> Use the file <i>namelist</i> in place of /dev/ksyms.</p>

The column headings and the meaning of the columns in an **ipcs** listing are given below; the letters in parentheses indicate the options that cause the corresponding heading to appear; “all” means that the heading always appears. Note: These options only determine what information is provided for each facility; they do not determine which facilities are listed.

T	(all)	Type of the facility: <ul style="list-style-type: none"> q message queue m shared memory segment s semaphore
ID	(all)	The identifier for the facility entry.
KEY	(all)	The key used as an argument to msgget , semget , or shmget to create the facility entry. (Note: The key of a shared memory segment is changed to IPC_PRIVATE when the segment has been removed until all processes attached to the segment detach it.)
MODE	(all)	The facility access modes and flags: The mode consists of 11 characters that are interpreted as follows. The first two characters are: <ul style="list-style-type: none"> R A process is waiting on a <i>msgrcv</i>. S A process is waiting on a <i>msgsnd</i>. D The associated shared memory segment has been removed. It will disappear when the last process attached to the segment detaches it. C The associated shared memory segment is to be cleared when the first attach is executed. <ul style="list-style-type: none"> – The corresponding special flag is not set. <p>The next nine characters are interpreted as three sets of three bits each. The first set refers to the owner’s permissions; the next to permissions of others in the user-group of the facility entry; and the last to all others. Within each set, the first character indicates permission to read, the second character indicates permission to write or alter the facility entry, and the last character is currently unused.</p> <p>The permissions are indicated as follows:</p> <ul style="list-style-type: none"> r Read permission is granted. w Write permission is granted. a Alter permission is granted. – The indicated permission is not granted.
OWNER	(all)	The login name of the owner of the facility entry.
GROUP	(all)	The group name of the group of the owner of the facility entry.
CREATOR	(a,c)	The login name of the creator of the facility entry.
CGROUP	(a,c)	The group name of the group of the creator of the facility entry.
CBYTES	(a,o)	The number of bytes in messages currently outstanding on the associated message queue.
QNUM	(a,o)	The number of messages currently outstanding on the associated

		message queue.
QBYTES	(a,b)	The maximum number of bytes allowed in messages outstanding on the associated message queue.
LSPID	(a,p)	The process ID of the last process to send a message to the associated queue.
LRPID	(a,p)	The process ID of the last process to receive a message from the associated queue.
STIME	(a,t)	The time the last message was sent to the associated queue.
RTIME	(a,t)	The time the last message was received from the associated queue.
CTIME	(a,t)	The time when the associated entry was created or changed.
NATTCH	(a,o)	The number of processes attached to the associated shared memory segment.
SEGSZ	(a,b)	The size of the associated shared memory segment.
CPID	(a,p)	The process ID of the creator of the shared memory entry.
LPID	(a,p)	The process ID of the last process to attach or detach the shared memory segment.
ATIME	(a,t)	The time the last attach was completed to the associated shared memory segment.
DTIME	(a,t)	The time the last detach was completed on the associated shared memory segment.
NSEMS	(a,b)	The number of semaphores in the set associated with the semaphore entry.
OTIME	(a,t)	The time the last semaphore operation was completed on the set associated with the semaphore entry.

FILES	/etc/group	group names
	/etc/passwd	user names
	/dev/mem	memory
	/dev/ksyms	system namelist

SEE ALSO **msgop(2)**, **semop(2)**, **shmop(2)**

NOTES If the user specifies either the **-C** or **-N** flag, the real and effective UID/GID is set to the real UID/GID of the user invoking **ipcs**.
Things can change while **ipcs** is running; the information it gives is guaranteed to be accurate only when it was retrieved.

NAME	jobs, fg, bg, stop, notify – shell built-in functions to control process execution
SYNOPSIS	
sh	jobs [-p -l] [<i>%jobid ...</i>] jobs -x <i>command</i> [<i>arguments</i>] fg [<i>%jobid ...</i>] bg [<i>%jobid ...</i>] stop <i>%jobid ...</i> stop <i>pid ...</i>
csh	jobs. [-l] fg [<i>%job</i>] bg [<i>%job</i>] ... notify [<i>%job</i>] ... stop <i>%jobid ...</i> stop <i>pid ...</i>
ksh	jobs [-lnp] [<i>%job ...</i>] fg [<i>%job ...</i>] bg [<i>%job ...</i>] stop <i>%jobid ...</i> stop <i>pid ...</i>
DESCRIPTION	
sh	<p>When Job Control is enabled, jobs reports all jobs that are stopped or executing in the background. If <i>%jobid</i> is omitted, all jobs that are stopped or running in the background will be reported. The following options will modify/enhance the output of jobs:</p> <ul style="list-style-type: none"> -l Report the process group ID and working directory of the jobs. -p Report only the process group ID of the jobs. -x Replace any <i>jobid</i> found in <i>command</i> or <i>arguments</i> with the corresponding process group ID, and then execute <i>command</i> passing it <i>arguments</i>. <p>When the shell is invoked as jsh, Job Control is enabled in addition to all of the functionality described previously for sh. Typically Job Control is enabled for the interactive shell only. Non-interactive shells typically do not benefit from the added functionality of Job Control.</p> <p>With Job Control enabled every command or pipeline the user enters at the terminal is called a <i>job</i>. All jobs exist in one of the following states: foreground, background or stopped. These terms are defined as follows: 1) a job in the foreground has read and write access to the controlling terminal; 2) a job in the background is denied read access and has conditional write access to the controlling terminal (see stty(1)); 3) a stopped job</p>

is a job that has been placed in a suspended state, usually as a result of a **SIGTSTP** signal (see **signal(5)**).

Every job that the shell starts is assigned a positive integer, called a *job number* which is tracked by the shell and will be used as an identifier to indicate a specific job. Additionally the shell keeps track of the *current* and *previous* jobs. The *current job* is the most recent job to be started or restarted. The *previous job* is the first non-current job.

The acceptable syntax for a Job Identifier is of the form:

%jobid

where, *jobid* may be specified in any of the following formats:

% or **+** for the current job

- for the previous job

?<string> specify the job for which the command line uniquely contains *string*.

n for job number *n*, where *n* is a job number

pref where *pref* is a unique prefix of the command name (for example, if the command **ls -l name** were running in the background, it could be referred to as **%ls**); *pref* cannot contain blanks unless it is quoted.

When Job Control is enabled, **fg** resumes the execution of a stopped job in the foreground, also moves an executing background job into the foreground. If *%jobid* is omitted the current job is assumed.

When Job Control is enabled, **bg** resumes the execution of a stopped job in the background. If *%jobid* is omitted the current job is assumed.

stop stops the execution of a background job(s) by using its *jobid*, or of any process by using its *pid*. (see **ps(1)**)

csh **jobs**, without an argument, lists the active jobs under job control.

-l List process IDs, in addition to the normal information.

The shell associates a numbered *job* with each command sequence to keep track of those commands that are running in the background or have been stopped with TSTP signals (typically CTRL-Z). When a command or command sequence (semicolon separated list) is started in the background using the **&** metacharacter, the shell displays a line with the job number in brackets and a list of associated process numbers:

[1] 1234

To see the current list of jobs, use the **jobs** built-in command. The job most recently stopped (or put into the background if none are stopped) is referred to as the *current* job and is indicated with a '+'. The previous job is indicated with a '-'; when the current job is terminated or moved to the foreground, this job takes its place (becomes the new current job).

To manipulate jobs, refer to the **bg**, **fg**, **kill**, **stop**, and **%** built-in commands.

A reference to a job begins with a '%'. By itself, the percent-sign refers to the current job.

% %+ %% The current job.

%- The previous job.

%j Refer to job *j* as in: '**kill -9 %j**'. *j* can be a job number, or a string that uniquely specifies the command line by which it was started; '**fg %vi**' might bring a stopped **vi** job to the foreground, for instance.

??string Specify the job for which the command line uniquely contains *string*.

A job running in the background stops when it attempts to read from the terminal. Background jobs can normally produce output, but this can be suppressed using the '**stty tostop**' command.

fg brings the current or specified *job* into the foreground.

bg runs the current or specified jobs in the background.

stop stops the execution of a background job(s) by using its *jobid*, or of any process by using its *pid*. (see **ps**(1)).

notify will notify the user asynchronously when the status of the current job or specified jobs changes.

ksh **jobs** lists information about each given job; or all active jobs if *job* is omitted. The **-l** flag lists process ids in addition to the normal information. The **-n** flag only displays jobs that have stopped or exited since last notified. The **-p** flag causes only the process group to be listed. If the **monitor** option of the **set** command is turned on, an interactive shell associates a **job** with each pipeline. It keeps a table of current jobs, printed by the **jobs** command, and assigns them small integer numbers. When a job is started asynchronously with **&**, the shell prints a line which looks like:

[1] 1234

indicating that the **job**, which was started asynchronously, was job number 1 and had one (top-level) process, whose process id was 1234.

If you are running a job and wish to do something else you may hit the key **^Z** (CTRL-Z) which sends a **STOP** signal to the current job. The shell will then normally indicate that the job has been '**Stopped**', and print another prompt. You can then manipulate the state of this job, putting it in the background with the **bg** command, or run some other commands and then eventually bring the job back into the foreground with the foreground command **fg**. A **^Z** takes effect immediately and is like an interrupt in that pending output and unread input are discarded when it is typed.

A job being run in the background will stop if it tries to read from the terminal. Background jobs are normally allowed to produce output, but this can be disabled by giving the command "**stty tostop**". If you set this tty option, then background jobs will stop when they try to produce output like they do when they try to read input.

There are several ways to refer to **jobs** in the shell. A **job** can be referred to by the process id of any process of the **job** or by one of the following:

%number The job with the given number.
%string Any job whose command line begins with *string*.
%?string Any job whose command line begins with *string*.
%?string Any job whose command line contains *string*.
%% Current job.
%+ Equivalent to %%.
%- Previous job.

The shell learns immediately whenever a process changes state. It normally informs you whenever a job becomes blocked so that no further progress is possible, but only just before it prints a prompt. This is done so that it does not otherwise disturb your work. When the monitor mode is on, each background job that completes triggers any trap set for **CHLD**. When you try to leave the shell while jobs are running or stopped, you will be warned that 'You have stopped(running) jobs.' You may use the **jobs** command to see what they are. If you do this or immediately try to exit again, the shell will not warn you a second time, and the stopped jobs will be terminated.

The **fg** command is only on systems that support job control. Each *job* specified is brought to the foreground. Otherwise, the current job is brought into the foreground.

The **bg** command is only on systems that support job control. **bg** puts each specified *job* into the background. The current job is put in the background if *job* is not specified.

stop stops the execution of a background job(s) by using its *jobid*, or of any process by using its *pid*. (see **ps**(1))

SEE ALSO

csh(1), **kill**(1), **ksh**(1), **ps**(1), **sh**(1), **stop**(1), **shell_builtins**(1), **stty**(1) **signal**(5)

NAME	join – relational database operator
SYNOPSIS	join [-an] [-e s] [-jn m] [-o list] [-tc] <i>filename1 filename2</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>join forms, on the standard output, a join of the two relations specified by the lines of <i>filename1</i> and <i>filename2</i>. If <i>filename1</i> is -, the standard input is used.</p> <p><i>filename1</i> and <i>filename2</i> must be sorted in increasing ASCII collating sequence on the fields on which they are to be joined, normally the first in each line (see sort(1)).</p> <p>There is one line in the output for each pair of lines in <i>filename1</i> and <i>filename2</i> that have identical join fields. The output line normally consists of the common field, then the rest of the line from <i>filename1</i>, then the rest of the line from <i>filename2</i>.</p> <p>The default input field separators are blank, tab, or new-line. In this case, multiple separators count as one field separator, and leading separators are ignored. The default output field separator is a blank.</p>
OPTIONS	<p>Some of the options below use the argument <i>n</i>. This argument should be a 1 or a 2 referring to either <i>filename1</i> or <i>filename2</i>, respectively.</p> <p>-an In addition to the normal output, produce a line for each unpairable line in file <i>n</i>, where <i>n</i> is 1 or 2.</p> <p>-e s Replace empty output fields with string <i>s</i>.</p> <p>-jn m Join on the <i>m</i>th field of file <i>n</i>. If <i>n</i> is missing, use the <i>m</i>th field in each file. Fields are numbered starting with 1.</p> <p>-o list Each output line includes the fields specified in <i>list</i>, each element of which has the form <i>n.m</i>, where <i>n</i> is a file number and <i>m</i> is a field number. The common field is not printed unless specifically requested.</p> <p>-tc Use character <i>c</i> as a separator (tab character). Every appearance of <i>c</i> in a line is significant. The character <i>c</i> is used as the field separator for both input and output.</p>
EXAMPLES	<p>The following command line will join the password file and the group file, matching on the numeric group ID, and outputting the login name, the group name and the login directory. It is assumed that the files have been sorted in ASCII collating sequence on the group ID fields.</p> <pre>example% join -j1 4 -j2 3 -o 1.1 2.1 1.6 -t: /etc/passwd /etc/group</pre>

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

awk(1), **comm(1)**, **sort(1)**, **uniq(1)**, **environ(5)**

NOTES

With default field separation, the collating sequence is that of **sort -b**; with **-t**, the sequence is that of a plain sort.

The conventions of the **join**, **sort**, **comm**, **uniq**, and **awk** commands are wildly incongruous.

Filenames that are numeric may cause conflict when the **-o** option is used just before listing filenames.

NAME	kbd – manipulate the state of keyboard or display the type of keyboard
SYNOPSIS	kbd [-r] [-t] [-c on off] [-d keyboard device]
AVAILABILITY	SUNWcsu
DESCRIPTION	kbd manipulates the state of the keyboard, or displays the keyboard type. The default keyboard device being set is /dev/kbd . Only keyboards that support a clicker respond to the -c option. If you want to turn clicking on by default, this can be set in the /etc/rc* files.
OPTIONS	-r Reset the keyboard as if power-up. -t Return the type of the keyboard being used. -c on/off state Turn the clicking of the keyboard on or off. on Enable clicking. off Disable clicking. -d keyboard device Specify the keyboard device being set. The default is /dev/kbd .
EXAMPLES	The following example displays the keyboard type. example% kbd -t type 4 Sun keyboard To enable clicking by default, add the following line to your /etc/rc* file. kbd -c on
FILES	/etc/rc* /dev/kbd
SEE ALSO	loadkeys(1) , keytables(4) , kb(7)
BUGS	There is no way to determine the state of the keyboard click setting.

NAME	kdestroy – destroy Kerberos tickets
SYNOPSIS	<code>/usr/bin/kdestroy [-fnq]</code>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>kdestroy destroys the user's active Kerberos authorization tickets by writing zeros to the file that contains them. If the ticket file does not exist, kdestroy displays a message to that effect.</p> <p>After overwriting the file, kdestroy removes the file from the system. The utility displays a message indicating the success or failure of the operation. If kdestroy is unable to destroy the ticket file, it will warn you by making your terminal beep.</p> <p>In addition to removing the ticket file, kdestroy also invalidates all Kerberos credentials for this user being held in the kernel for use with NFS requests.</p> <p>If desired, you can place the kdestroy command in your .logout file so that your tickets are destroyed automatically when you logout. Note, however, that doing this will cause NFS operations done on your behalf to fail after you logout.</p>
OPTIONS	<p>-f Do not display the status message.</p> <p>-n Do not invalidate NFS credentials in the kernel. The credentials will continue to be valid until their normal expiration time, although new ones cannot be obtained until kinit(1) is run again for this user.</p> <p>-q Do not make your terminal beep if kdestroy fails to destroy the tickets.</p>
FILES	The file specified by the KRBTKFILE environment variable if set, otherwise /tmp/tktuid
SEE ALSO	kerberos(1) , kinit(1) , klist(1)
BUGS	Only the tickets in the user's current ticket file are destroyed. Separate ticket files are used to hold root instance and password changing tickets. These files should probably be destroyed too, or all of a user's tickets should be kept in a single ticket file.
AUTHORS	Steve Miller, MIT Project Athena/Digital Equipment Corporation Clifford Neuman, MIT Project Athena Bill Sommerfeld, MIT Project Athena

NAME	kerberos – introduction to the Kerberos system
DESCRIPTION	<p>The Kerberos system authenticates individual users in a network environment. After authenticating yourself to Kerberos, you can use the kerberos authentication option of network services such as NFS. In addition, in some environments you can use network utilities such as rlogin(1), rnp(1), and rsh(1) without having to present passwords to remote hosts and without having to bother with .rhosts files. See your system administrator for more information about Kerberos support at your site.</p> <p>Before you can use Kerberos, you must be registered as a user in the Kerberos database. You can use the kinit(1) command to find out your status. This command tries to log you into the Kerberos system. kinit will prompt you for a username and password. Enter your username and password. If the utility lets you login without giving you a message, you have already been registered.</p> <p>If you enter your username and kinit responds with this message:</p> <p style="text-align: center;">Principal unknown (kerberos)</p> <p>you haven't been registered as a Kerberos user. See your system administrator.</p> <p>A Kerberos name contains three parts. The first is the <i>principal name</i>, which is usually a user's or service's name. The second is the <i>instance</i>, which in the case of a user is usually NULL. Some users may have privileged instances, however, such as root or admin. In the case of a service, the instance is the name of the machine on which it runs; that is, there can be an NFS service running on the machine ABC, which is different from the NFS service running on the machine XYZ. The third part of a Kerberos name is the <i>realm</i>. The realm corresponds to the Kerberos service providing authentication for the principal. For example, at MIT there is a Kerberos running at the Laboratory for Computer Science and one running at Project Athena.</p> <p>When writing a Kerberos name, the principal name is separated from the instance (if not NULL) by a period, and the realm (if not the local realm) follows, preceded by an "@" sign. The following are examples of valid Kerberos names:</p> <p style="text-align: center;">billb jis.admin srz@lcs.mit.edu treese.root@athena.mit.edu</p> <p>When you authenticate yourself with Kerberos, typically through the kinit command, Kerberos gives you an initial Kerberos <i>ticket</i>. (A Kerberos ticket is an encrypted protocol message that provides authentication.) Kerberos uses this ticket for network utilities such as NFS, rlogin and rnp. The ticket transactions are done transparently, so you do not have to worry about their management.</p>

Note, however, that tickets expire. Privileged tickets, such as root instance tickets, expire in a few minutes, while tickets that carry more ordinary privileges may be good for several hours or a day, depending on the installation's policy. If your login session extends beyond the time limit, you will have to re-authenticate yourself to Kerberos to get new tickets. Use the **kinit** command to re-authenticate yourself.

If you use the **kinit** command to get your tickets, you can use the **kdestroy(1)** command to destroy your tickets before you end your login session. For more information about the **kinit** and **kdestroy** commands, see the **kinit(1)** and **kdestroy(1)** manual pages.

Currently, Kerberos supports NFS and other RPC network services using the **AUTH_KERB** authentication type. In some environments, the following network services are also supported: **rlogin**, **rsh**, and **rcp**. Other services are being worked on, such as the **pop** mail system, but are not yet available.

SEE ALSO **kdestroy(1)**, **kinit(1)**, **klist(1)**, **kerbd(1M)**, **kerberos(3N)**, **krb.conf(4)**

BUGS Kerberos will not do authentication forwarding. In other words, if you use **rlogin** to login to a remote host, you cannot use Kerberos services from that host until you authenticate yourself explicitly on that host. Although you may need to authenticate yourself on the remote host, be aware that when you do so, **rlogin** sends your password across the network in clear text.

AUTHORS Steve Miller, MIT Project Athena/Digital Equipment Corporation
Clifford Neuman, MIT Project Athena

The following people helped out on various aspects of the system:

Jeff Schiller designed and wrote the administration server and its user interface, **kadmin**. He also wrote the **dbm** version of the database management system.

Mark Colan developed the Kerberos versions of **rlogin**, **rsh**, and **rcp**, as well as contributing work on the servers.

John Ostlund developed the Kerberos versions of **passwd** and **userreg**.

Stan Zandarotti pioneered Kerberos in a foreign realm (LCS), and made many contributions based on that experience.

Many people contributed code and/or useful ideas. These include, Jim Aspnes, Bob Baldwin, John Barba, Richard Basch, Jim Bloom, Bill Bryant, Rob French, Dan Geer, David Jedlinsky, John Kohl, John Kubiawicz, Bob McKie, Brian Murphy, Ken Raeburn, Chris Reed, Jon Rochlis, Mike Shanzer, Bill Sommerfeld, Jennifer Steiner, Ted Ts'o, and Win Treese.

RESTRICTIONS COPYRIGHT 1985,1986 Massachusetts Institute of Technology

NAME	keylogin – decrypt and store secret key with key serv
SYNOPSIS	<code>/usr/bin/keylogin [-r]</code>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The keylogin command prompts for a password, and uses it to decrypt the user's secret key. The key may be found in the <code>/etc/publickey</code> file (see publickey(4)) or the NIS map "publickey.byname" or the NIS+ table "cred.org_dir" in the user's home domain. The sources and their lookup order are specified in the <code>/etc/nsswitch.conf</code> file (see nsswitch.conf(4)). Once decrypted, the user's secret key is stored by the local key server process, key serv(1M). This stored key is used when issuing requests to any secure RPC services, such as NFS or NIS+. The program keylogout(1) can be used to delete the key stored by key serv.</p> <p>keylogin will fail if it cannot get the caller's key, or the password given is incorrect. For a new user or host, a new key can be added using newkey(1M), nisaddcred(1M), or nisclient(1M).</p>
OPTIONS	<p><code>-r</code> Update the <code>/etc.rootkey</code> file. This file holds the unencrypted secret key of the super-user. Only the super-user may use this option. It is used so that processes running as super-user can issue authenticated requests without requiring that the administrator explicitly run keylogin as super-user at system startup time (see key serv(1M)). The <code>-r</code> option should be used by the administrator when the host's entry in the publickey database has changed, and the <code>/etc.rootkey</code> file has become out-of-date with respect to the actual key pair stored in the publickey database. The permissions on the <code>/etc.rootkey</code> file are such that it may be read and written by the super-user but by no other user on the system.</p>
FILES	<code>/etc.rootkey</code> super-user's secret key
SEE ALSO	chkey (1), keylogout (1), login (1), key serv (1M), newkey (1M), nisaddcred (1M), nisclient (1M), publickey (4), nsswitch.conf (4)

NAME	keylogout – delete stored secret key with keyserv
SYNOPSIS	/usr/bin/keylogout [-f]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>keylogout deletes the key stored by the key server process keyserv(1M). Further access to the key is revoked; however, current session keys may remain valid until they expire or are refreshed.</p> <p>Deleting the keys stored by keyserv will cause any background jobs or scheduled at(1) jobs that need secure RPC services to fail. Since only one copy of the key is kept on a machine, it is a bad idea to place a call to this command in your .logout file since it will affect other sessions on the same machine.</p>
OPTIONS	<p>-f Force keylogout to delete the secret key for the super-user. By default, keylogout by the super-user is disallowed because it would break all RPC services, such as NFS, that are started by the super-user.</p>
SEE ALSO	at (1), chkey (1), login (1), keylogin (1), keyserv (1M), newkey (1M), publickey (4)

NAME	kill – terminate a process by default
SYNOPSIS	kill [<i>-signal</i>] <i>pid</i> . . . kill <i>-signal</i> <i>-pgid</i> . . . kill <i>-l</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>kill sends a signal to the specified processes. The value of <i>signal</i> may be numeric or symbolic. (See signal(5)). The symbolic signal name is the name as it appears in <code>/usr/include/sys/signal.h</code>, with the SIG prefix stripped off. Signal 15 (SIGTERM) is sent by default; this will normally kill processes that do not catch or ignore the signal.</p> <p><i>pid</i> and <i>pgid</i> are unsigned numeric strings that identify which process(es) should receive the signal. If <i>pid</i> is used, the process with process ID <i>pid</i> is selected. If <i>pgid</i> is used, all processes with process group ID <i>pgid</i> are selected.</p> <p>The process number of each asynchronous process started with & is reported by the shell (unless more than one process is started in a pipeline, in which case the number of the last process in the pipeline is reported). Process numbers can also be found by using ps(1).</p> <p>The details of the kill are described in kill(2). For example, if process number 0 is specified, all processes in the process group are signaled.</p> <p>The signalled process must belong to the current user unless the user is the super-user.</p>
OPTIONS	<p>-l kill will print a list of symbolic signal names.</p> <p>The 3 shells, sh(1), csh(1), and ksh(1) each have built-in versions of the kill command for processes identified with a jobid number.</p>
SEE ALSO	csh (1), jobs (1), ksh (1), ps (1), sh (1), shell_builtins (1), kill (2), signal (3C), signal (5)
NOTES	
sh	<p>The Bourne shell, sh, has a built-in version of kill to provide the functionality of the kill command for processes identified with a <i>jobid</i>. The sh syntax is:</p> <p>kill [<i>-signal</i>] <i>%jobid</i></p>
csh	<p>The C-shell, csh, also has a built-in kill command, whose syntax is:</p> <p>kill [<i>-sig</i>] [<i>pid</i>] [<i>%job</i>] . . . kill <i>-l</i></p> <p>The csh kill built-in sends the TERM (terminate) signal, by default, or the signal specified, to the specified process ID, the <i>job</i> indicated, or the current <i>job</i>. Signals are either given by number or by name. There is no default. Typing kill does not send a signal to the current job. If the signal being sent is TERM (terminate) or HUP (hangup), then the job or process is sent a CONT (continue) signal as well.</p> <p>-l List the signal names that can be sent.</p>

ksh The **ksh kill**'s syntax is:

kill [*-sig*] %*job* . . .

kill -l

The **ksh kill** sends either the TERM (terminate) signal or the specified signal to the specified jobs or processes. Signals are either given by number or by names (as given in **signal(5)** stripped of the prefix "SIG" with the exception that SIGCHD is named CHLD). If the signal being sent is TERM (terminate) or HUP (hangup), then the job or process will be sent a CONT (continue) signal if it is stopped. The argument *job* can be the process id of a process that is not a member of one of the active jobs. In the second form, **kill -l**, the signal numbers and names are listed.

NAME	kinit – Kerberos login utility
SYNOPSIS	kinit [-ilrv] [<i>username</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The kinit command is used to login to the Kerberos authentication and authorization system. Note that only registered Kerberos users can use the Kerberos system. For information about registering as a Kerberos user, see the kerberos(1) manual page.</p> <p>When you use kinit without options, the utility prompts for your <i>username</i> and Kerberos password, and tries to authenticate your login with the local Kerberos server. The <i>username</i> can be specified on the command line if desired.</p> <p>If Kerberos authenticates the login attempt, kinit retrieves your initial ticket and puts it in the ticket file specified by your KRBTKFILE environment variable. If this variable is undefined, your ticket will be stored in the file /tmp/tktuid, where <i>uid</i> specifies your user identification number. Tickets expire after a specified lifetime, after which kinit must be run again to refresh the tickets. The default ticket lifetime is 8 hours.</p> <p>The kdestroy(1) command may be used to destroy any active tickets before you end your login session.</p>
OPTIONS	<p>-i kinit prompts you for a Kerberos instance.</p> <p>-l kinit prompts you for a ticket lifetime in minutes. Due to protocol restrictions in Kerberos Version 4, this value must be between 5 and 1275 minutes.</p> <p>-r kinit prompts you for a Kerberos realm. This option lets you authenticate yourself with a remote Kerberos server.</p> <p>-v Verbose mode. kinit prints the name of the ticket file used, and a status message indicating the success or failure of your login attempt.</p>
SEE ALSO	kdestroy (1), kerberos (1), klist (1)
BUGS	The -r option has not been fully implemented.
AUTHORS	Steve Miller, MIT Project Athena/Digital Equipment Corporation Clifford Neuman, MIT Project Athena

NAME	klist – list currently held Kerberos tickets	
SYNOPSIS	klist [-st] [-file <i>name</i>] [-srvtab]	
AVAILABILITY	SUNWcsu	
DESCRIPTION	<p>klist prints the name of the ticket file, the identity of the principal that the tickets are for (as listed in the ticket file), and the principal names of all Kerberos tickets currently held by the user, along with the issue and expire time for each authenticator. Principal names are listed in the form <i>name.instance@realm</i>, with the '.' omitted if the instance is NULL, and the '@' omitted if the realm is NULL.</p> <p>The value of the KRBTKFILE environment variable is used as the name of the ticket file. If this environment variable is not set, then the file /tmp/tktuid is used, where <i>uid</i> is the current user-id of the user.</p>	
OPTIONS	-s	Silent. Do not print the issue and expire times, the name of the ticket file, or the identity of the principal.
	-t	klist checks for the existence of a non-expired ticket-granting-ticket in the ticket file. If one is present, it exits with status 0, else it exits with status 1. No output is generated when this option is specified.
	-file <i>name</i>	File <i>name</i> is used as the ticket file.
	-srvtab	The file is treated as a service key file, and the names of the keys contained therein are printed. If no file is specified with a -file option, the default is /etc/srvtab .
FILES	/etc/krb.conf	to get the name of the local realm
	/tmp/tktuid	as the default ticket file
	/etc/srvtab	as the default service key file
SEE ALSO	kerberos(1) , kinit(1) , kdestroy(1)	
BUGS	When reading a file as a service key file, very little sanity or error checking is performed.	

NAME	ksh, rksh – KornShell, a standard/restricted command and programming language
SYNOPSIS	ksh [<i>±aefhikmnoprstuvx</i>] [<i>±o option</i>] ... [<i>-c string</i>] [<i>arg ...</i>] rksh [<i>±aefhikmnoprstuvx</i>] [<i>±o option</i>] ... [<i>-c string</i>] [<i>arg ...</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	ksh is a command and programming language that executes commands read from a terminal or a file. rksh is a restricted version of the command interpreter ksh ; it is used to set up login names and execution environments whose capabilities are more controlled than those of the standard shell. See Invocation below for the meaning of arguments to the shell.
Definitions	A <i>metacharacter</i> is one of the following characters: ; & () < > new-line space tab A <i>blank</i> is a tab or a space . An <i>identifier</i> is a sequence of letters, digits, or underscores starting with a letter or underscore. Identifiers are used as names for <i>functions</i> and <i>variables</i> . A <i>word</i> is a sequence of <i>characters</i> separated by one or more non-quoted <i>metacharacters</i> . A <i>command</i> is a sequence of characters in the syntax of the shell language. The shell reads each command and carries out the desired action either directly or by invoking separate utilities. A <i>special command</i> is a command that is carried out by the shell without creating a separate process. Except for documented side effects, most special commands can be implemented as separate utilities.
Commands	A <i>simple-command</i> is a sequence of <i>blank</i> separated words which may be preceded by a variable assignment list. (See Environment below.) The first word specifies the name of the command to be executed. Except as specified below, the remaining words are passed as arguments to the invoked command. The command name is passed as argument 0 (see exec(2)). The <i>value</i> of a simple-command is its exit status if it terminates normally, or (octal) 200+ <i>status</i> if it terminates abnormally (see signal(3C) for a list of status values). A <i>pipeline</i> is a sequence of one or more <i>commands</i> separated by . The standard output of each command but the last is connected by a pipe(2) to the standard input of the next command. Each command is run as a separate process; the shell waits for the last command to terminate. The exit status of a pipeline is the exit status of the last command. A <i>list</i> is a sequence of one or more <i>pipelines</i> separated by ;, &, &&, or , and optionally terminated by ;, &, or &. Of these five symbols, ;, &, and & have equal precedence, which is lower than that of && and . The symbols && and also have equal precedence. A semicolon (;) causes sequential execution of the preceding pipeline; an ampersand (&) causes asynchronous execution of the preceding pipeline (i.e., the shell does <i>not</i> wait for that pipeline to finish). The symbol & causes asynchronous execution of the preceding command or pipeline with a two-way pipe established to the parent shell.

The standard input and output of the spawned command can be written to and read from by the parent shell using the **-p** option of the special commands **read** and **print** described in **Special Commands**. The symbol **&&** (**| |**) causes the *list* following it to be executed only if the preceding pipeline returns zero (or a non-zero) value. An arbitrary number of new-lines may appear in a *list*, instead of a semicolon, to delimit a command.

A *command* is either a *simple-command* or one of the following. Unless otherwise stated, the value returned by a command is that of the last simple-command executed in the command.

for identifier [in word ...] ;do list ;done

Each time a **for** command is executed, *identifier* is set to the next *word* taken from the **in word** list. If **in word ...** is omitted, then the **for** command executes the **do list** once for each positional parameter that is set (see **Parameter Substitution** below). Execution ends when there are no more words in the list.

select identifier [in word ...] ;do list ;done

A **select** command prints to standard error (file descriptor 2), the set of *words*, each preceded by a number. If **in word ...** is omitted, then the positional parameters are used instead (see **Parameter Substitution** below). The **PS3** prompt is printed and a line is read from the standard input. If this line consists of the number of one of the listed *words*, then the value of the variable *identifier* is set to the *word* corresponding to this number. If this line is empty the selection list is printed again. Otherwise the value of the variable *identifier* is set to **NULL**. (See **Blank Interpretation** about **NULL**). The contents of the line read from standard input is saved in the shell variable **REPLY**. The *list* is executed for each selection until a **break** or *end-of-file* is encountered. If the **REPLY** variable is set to **NULL** by the execution of *list*, then the selection list is printed before displaying the **PS3** prompt for the next selection.

case word in [(|)pattern [| pattern] ...) list ;;] ... esac

A **case** command executes the *list* associated with the first *pattern* that matches *word*. The form of the patterns is the same as that used for file-name generation (see **File Name Generation** below).

if list ;then list [(| elif list ; then list] ...] ; else list ;fi

The *list* following **if** is executed and, if it returns an exit status of zero, the *list* following the first **then** is executed. Otherwise, the *list* following **elif** is executed and, if its value is zero, the *list* following the next **then** is executed. Failing that, the **else list** is executed. If no **else list** or **then list** is executed, then the **if** command returns a zero exit status.

while list ;do list ;done

until list ;do list ;done

A **while** command repeatedly executes the **while list** and, if the exit status of the last command in the list is zero, executes the **do list**; otherwise the loop terminates. If no commands in the **do list** are executed, then the **while** command returns a zero exit status; **until** may be used in place of **while** to negate the loop termination test.

(list) Execute *list* in a separate environment. Note, that if two adjacent open parentheses are needed for nesting, a space must be inserted to avoid arithmetic evaluation as described below.

{list} *list* is simply executed. Note that unlike the metacharacters (and), { and } are *reserved words* and must occur at the beginning of a line or after a ; in order to be recognized.

[[expression]]
Evaluates *expression* and returns a zero exit status when *expression* is true. See **Conditional Expressions** below, for a description of *expression*.

function identifier { list ;}
identifier() { list ;}
Define a function which is referenced by *identifier*. The body of the function is the *list* of commands between { and }. (See **Functions** below).

time pipeline
The *pipeline* is executed and the elapsed time as well as the user and system time are printed to standard error.

The following reserved words are only recognized as the first word of a command and when not quoted:

if then else elif fi case esac for while until do done { } function select time [[]]

Comments A word beginning with # causes that word and all the following characters up to a new-line to be ignored.

Aliasing The first word of each command is replaced by the text of an **alias** if an **alias** for this word has been defined. An **alias** name consists of any number of characters excluding metacharacters, quoting characters, file expansion characters, parameter and command substitution characters, and =. The replacement string can contain any valid shell script including the metacharacters listed above. The first word of each command in the replaced text, other than any that are in the process of being replaced, will be tested for aliases. If the last character of the alias value is a *blank* then the word following the alias will also be checked for alias substitution. Aliases can be used to redefine special builtin commands but cannot be used to redefine the reserved words listed above. Aliases can be created, listed, and exported with the **alias** command and can be removed with the **unalias** command. Exported aliases remain in effect for scripts invoked by name, but must be reinitialized for separate invocations of the shell (see **Invocation** below).

Aliasing is performed when scripts are read, not while they are executed. Therefore, for an alias to take effect the **alias** definition command has to be executed before the command which references the alias is read.

Aliases are frequently used as a short hand for full path names. An option to the aliasing facility allows the value of the alias to be automatically set to the full pathname of the corresponding command. These aliases are called *tracked* aliases. The value of a *tracked* alias is defined the first time the corresponding command is looked up and becomes

undefined each time the **PATH** variable is reset. These aliases remain *tracked* so that the next subsequent reference will redefine the value. Several tracked aliases are compiled into the shell. The **-h** option of the **set** command makes each referenced command name into a tracked alias.

The following *exported aliases* are compiled into (and built-in to) the shell but can be unset or redefined:

```
autoload='typeset -fu'
false='let 0'
functions='typeset -f'
hash='alias -t'
history='fc -l'
integer='typeset -i'
nohup='nohup '
r='fc -e -'
true=':'
type='whence -v'
```

Tilde Substitution

After alias substitution is performed, each word is checked to see if it begins with an unquoted `~`. If it does, then the word up to a `/` is checked to see if it matches a user name. If a match is found, the `~` and the matched login name are replaced by the login directory of the matched user. This is called a *tilde* substitution. If no match is found, the original text is left unchanged. A `~` by itself, or in front of a `/`, is replaced by **\$HOME**. A `~` followed by a `+` or `-` is replaced by **\$PWD** and **\$OLDPWD** respectively.

In addition, *tilde* substitution is attempted when the value of a *variable assignment* begins with a `~`.

Command Substitution

The standard output from a *command* enclosed in parenthesis preceded by a dollar sign (**\$(command)**) or a pair of grave accents (**`**) may be used as part or all of a word; trailing new-lines are removed. In the second (archaic) form, the string between the quotes is processed for special quoting characters before the command is executed. (See **Quoting** below.) The command substitution **\$(cat file)** can be replaced by the equivalent but faster (**<file**). Command substitution of most special commands that do not perform input/output redirection are carried out without creating a separate process.

An arithmetic expression enclosed in double parenthesis preceded by a dollar sign (**\$(arithmetic-expression)**) is replaced by the value of the arithmetic expression within the double parenthesis.

Process Substitution

This feature is available in SunOS and only on versions of the UNIX operating system that support the **/dev/fd** directory for naming open files. Each command argument of the form **<(list)** or **>(list)** will run process *list* asynchronously connected to some file in **/dev/fd**. The name of this file will become the argument to the command. If the form with **>** is selected then writing on this file will provide input for *list*. If **<** is used, then the file passed as an argument will contain the output of the *list* process. For example,

paste <(cut -f1 *file1*) <(cut -f3 *file2*) | tee >(*process1*) >(*process2*)

cuts fields 1 and 3 from the files *file1* and *file2* respectively, **past**es the results together, and sends it to the processes *process1* and *process2*, as well as putting it onto the standard output. Note that the file, which is passed as an argument to the command, is a UNIX **pipe(2)** so programs that expect to **lseek(2)** on the file will not work.

Parameter Substitution

A *parameter* is an *identifier*, one or more digits, or any of the characters *, @, #, ?, -, \$, and !. A *variable* (a *parameter* denoted by an *identifier*) has a *value* and zero or more *attributes*. *variables* can be assigned *values* and *attributes* by using the **typeset** special command. The attributes supported by the shell are described later with the **typeset** special command. Exported variables pass values and attributes to the environment.

The shell supports a one-dimensional array facility. An element of an array variable is referenced by a *subscript*. A *subscript* is denoted by a [, followed by an *arithmetic expression* (see **Arithmetic Evaluation** below) followed by a]. To assign values to an array, use **set -A name value ...**. The *value* of all subscripts must be in the range of 0 through 1023. Arrays need not be declared. Any reference to a variable with a valid subscript is legal and an array will be created if necessary. Referencing an array without a subscript is equivalent to referencing the element zero.

The *value* of a *variable* may be assigned by writing:

name=value [*name=value*] ...

If the integer attribute, **-i**, is set for *name*, the *value* is subject to arithmetic evaluation as described below.

Positional parameters, parameters denoted by a number, may be assigned values with the **set** special command. Parameter **\$0** is set from argument zero when the shell is invoked.

The character **\$** is used to introduce substitutable *parameters*.

\${parameter}

The shell reads all the characters from **\${** to the matching **}** as part of the same word even if it contains braces or metacharacters. The value, if any, of the parameter is substituted. The braces are required when *parameter* is followed by a letter, digit, or underscore that is not to be interpreted as part of its name or when a variable is subscripted. If *parameter* is one or more digits then it is a positional parameter. A positional parameter of more than one digit must be enclosed in braces. If *parameter* is * or @, then all the positional parameters, starting with **\$1**, are substituted (separated by a field separator character). If an array *identifier* with subscript * or @ is used, then the value for each of the elements is substituted (separated by a field separator character).

\${#parameter}

If *parameter* is * or @, the number of positional parameters is substituted. Otherwise, the length of the value of the *parameter* is substituted.

\${#identifier[*]}

The number of elements in the array *identifier* is substituted.

`${parameter:-word}`

If *parameter* is set and is non-null then substitute its value; otherwise, substitute *word*.

`${parameter:=word}`

If *parameter* is not set or is NULL then set it to *word*; the value of the parameter is then substituted. Positional parameters may not be assigned to in this way.

`${parameter:?word}`

If *parameter* is set and is non-null then substitute its value; otherwise, print *word* and exit from the shell. If *word* is omitted then a standard message is printed.

`${parameter:+word}`

If *parameter* is set and is non-null then substitute *word*; otherwise substitute nothing.

`${parameter#pattern}`**`${parameter##pattern}`**

If the shell *pattern* matches the beginning of the value of *parameter*, then the value of this substitution is the value of the *parameter* with the matched portion deleted; otherwise, the value of this *parameter* is substituted. In the first form the smallest matching pattern is deleted and in the second form the largest matching pattern is deleted. The result is unspecified when *parameter* is @, *, or an array variable with subscript @, or *.

`${parameter%pattern}`**`${parameter%%pattern}`**

If the shell *pattern* matches the end of the value of *parameter*, then the value of this substitution is the value of the *parameter* with the matched part deleted; otherwise, substitute the value of *parameter*. In the first form the smallest matching pattern is deleted and in the second form the largest matching pattern is deleted. The result is unspecified when *parameter* is @, *, or an array variable with subscript @, or *.

In the above, *word* is not evaluated unless it is to be used as the substituted string, so that, in the following example, **pwd** is executed only if **d** is not set or is NULL (see **Blank Interpretation** for description about NULL);

```
echo ${d:-$(pwd)}
```

If the colon (:) is omitted from the above expressions, then the shell only checks whether *parameter* is set or not.

The following parameters are automatically set by the shell:

#	The number of positional parameters in decimal.
-	Flags supplied to the shell on invocation or by the set command.
?	The decimal value returned by the last executed command.
\$	The process number of this shell.
_	Initially, the value of _ is an absolute pathname of the shell or script

being executed as passed in the *environment*. Subsequently it is assigned the last argument of the previous command. This parameter is not set for commands which are asynchronous. This parameter is also used to hold the name of the matching **MAIL** file when checking for mail.

- !** The process number of the last background command invoked.
- ERRNO** The value of **errno** as set by the most recently failed system call. This value is system dependent and is intended for debugging purposes.
- LINENO** The line number of the current line within the script or function being executed.
- OLDPWD** The previous working directory set by the **cd** command.
- OPTARG** The value of the last option argument processed by the **getopts** special command.
- OPTIND** The index of the last option argument processed by the **getopts** special command.
- PPID** The process number of the parent of the shell.
- PWD** The present working directory set by the **cd** command.
- RANDOM** Each time this variable is referenced, a random integer, uniformly distributed between 0 and 32767, is generated. The sequence of random numbers can be initialized by assigning a numeric value to **RANDOM**.
- REPLY** This variable is set by the **select** statement and by the **read** special command when no arguments are supplied.
- SECONDS** Each time this variable is referenced, the number of seconds since shell invocation is returned. If this variable is assigned a value, then the value returned upon reference will be the value that was assigned plus the number of seconds since the assignment.

The following variables are used by the shell:

- CDPATH** The search path for the **cd** command.
- COLUMNS**
If this variable is set, the value is used to define the width of the edit window for the shell edit modes and for printing **select** lists.
- EDITOR** If the value of this variable ends in **emacs**, **gmacs**, or **vi** and the **VISUAL** variable is not set, then the corresponding option (see the **set** special command below) will be turned on.
- ENV** If this variable is set, then parameter substitution is performed on the value to generate the pathname of the script that will be executed when the **shell** is invoked. (See **Invocation** below.) This file is typically used for **alias** and **function** definitions.
- FCEDIT** The default editor name for the **fc** command.
- FPATH** The search path for function definitions. By default the **FPATH**

directories are searched after the **PATH** variable. If an executable file is found, then it is read and executed in the current environment. **FPATH** is searched before **PATH** when a function with the **-u** attribute is referenced. The preset alias **autoload** preset alias causes a function with the **-u** attribute to be created.

- IFS** Internal field separators, normally **space**, **tab**, and **new-line** that are used to separate command words which result from command or parameter substitution and for separating words with the special command **read**. The first character of the **IFS** variable is used to separate arguments for the **\$*** substitution (See **Quoting** below).
- HISTFILE** If this variable is set when the shell is invoked, then the value is the pathname of the file that will be used to store the command history. (See **Command re-entry** below.)
- HISTSIZE** If this variable is set when the shell is invoked, then the number of previously entered commands that are accessible by this shell will be greater than or equal to this number. The default is **128**.
- HOME** The default argument (home directory) for the **cd** command.
- LC_CTYPE** Determines how the shell handles characters. When **LC_CTYPE** is set to a valid value, the shell can display and handle text and filenames containing valid characters for that locale. However, the shell is not multibyte (EUC) capable. In the "C" locale, only ASCII characters are valid. If **LC_CTYPE** (see **environ(5)**) is not set in the environment, the operational behavior of the shell is determined by the value of the **LANG** environment variable. If **LC_ALL** is set, its contents are used to override both the **LANG** and the other **LC_*** variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale prevails.
- LINES** If this variable is set, the value is used to determine the column length for printing **select** lists. Select lists will print vertically until about two-thirds of **LINES** lines are filled.
- MAIL** If this variable is set to the name of a mail file *and* the **MAILPATH** variable is not set, then the shell informs the user of arrival of mail in the specified file.
- MAILCHECK**
This variable specifies how often (in seconds) the shell will check for changes in the modification time of any of the files specified by the **MAILPATH** or **MAIL** variables. The default value is **600** seconds. When the time has elapsed the shell will check before issuing the next prompt.
- MAILPATH**
A colon (;) separated list of file names. If this variable is set, then the shell informs the user of any modifications to the specified files that have occurred within the last **MAILCHECK** seconds. Each file name

can be followed by a ? and a message that will be printed. The message will undergo parameter substitution with the variable \$₀ defined as the name of the file that has changed. The default message is **you have mail in \$₀**.

- PATH** The search path for commands (see **Execution** below). The user may not change **PATH** if executing under **rksh** (except in **.profile**).
- PS1** The value of this variable is expanded for parameter substitution to define the primary prompt string which by default is "\$ ". The character ! in the primary prompt string is replaced by the *command* number (see *Command Re-entry* below). Two successive occurrences of ! will produce a single ! when the prompt string is printed.
- PS2** Secondary prompt string, by default "> ".
- PS3** Selection prompt string used within a **select** loop, by default "#? ".
- PS4** The value of this variable is expanded for parameter substitution and precedes each line of an execution trace. If omitted, the execution trace prompt is "+ ".
- SHELL** The pathname of the *shell* is kept in the environment. At invocation, if the basename of this variable is **rsh**, **rksh**, or **krsh**, then the shell becomes restricted.
- TMOUT** If set to a value greater than zero, the shell will terminate if a command is not entered within the prescribed number of seconds after issuing the **PS1** prompt. (Note that the shell can be compiled with a maximum bound for this value which cannot be exceeded.)
- VISUAL** If the value of this variable ends in **emacs**, **gmacs**, or **vi** then the corresponding option (see Special Command **set** below) will be turned on.

The shell gives default values to **PATH**, **PS1**, **PS2**, **PS3**, **PS4**, **MAILCHECK**, **FCEDIT**, **TMOUT** and **IFS**, while **HOME**, **SHELL ENV** and **MAIL** are not set at all by the shell (although **HOME** is set by **login(1)**). On some systems **MAIL** and **SHELL** are also set by **login(1)**.

Blank Interpretation

After parameter and command substitution, the results of substitutions are scanned for the field separator characters (those found in **IFS**) and split into distinct arguments where such characters are found. Explicit **NULL** arguments ("") or ("") are retained. Implicit **NULL** arguments (those resulting from *parameters* that have no values) are removed.

File Name Generation

Following substitution, each command *word* is scanned for the characters *, ?, and [unless the **-f** option has been **set**. If one of these characters appears then the word is regarded as a *pattern*. The word is replaced with lexicographically sorted file names that match the pattern. If no file name is found that matches the pattern, then the word is left unchanged. When a *pattern* is used for file name generation, the character period (.) at the start of a file name or immediately following a /, as well as the character / itself, must be matched explicitly. In other instances of pattern matching the / and . are not treated specially.

- * Matches any string, including the NULL string.
- ? Matches any single character.
- [. . .] Matches any one of the enclosed characters. A pair of characters separated by – matches any character lexically between the pair, inclusive. If the first character following the opening "[" is a "!", then any character not enclosed is matched. A – can be included in the character set by putting it as the first or last character.

A *pattern-list* is a list of one or more patterns separated from each other with a |. Composite patterns can be formed with one or more of the following:

- ?(*pattern-list*)
Optionally matches any one of the given patterns.
- *(*pattern-list*)
Matches zero or more occurrences of the given patterns.
- +(*pattern-list*)
Matches one or more occurrences of the given patterns.
- @(*pattern-list*)
Matches exactly one of the given patterns.
- !(*pattern-list*)
Matches anything, except one of the given patterns.

Quoting

Each of the *metacharacters* listed above (See **Definitions** above) has a special meaning to the shell and causes termination of a word unless quoted. A character may be *quoted* (i.e., made to stand for itself) by preceding it with a \. The pair \new-line is removed. All characters enclosed between a pair of single quote marks ('), are quoted. A single quote cannot appear within single quotes. Inside double quote marks (parameter and command substitution occur and \ quotes the characters \, ', and \$. The meaning of \$* and @\$ is identical when not quoted or when used as a parameter assignment value or as a file name. However, when used as a command argument, \$* is equivalent to "\$1d \$2d . . .", where *d* is the first character of the IFS variable, whereas @\$ is equivalent to \$1 \$2 Inside grave quote marks (`) \ quotes the characters \, ', and \$. If the grave quotes occur within double quotes, then \ also quotes the character ".

The special meaning of reserved words or aliases can be removed by quoting any character of the reserved word. The recognition of function names or special command names listed below cannot be altered by quoting them.

Arithmetic Evaluation

An ability to perform integer arithmetic is provided with the special command **let**. Evaluations are performed using *long* arithmetic. Constants are of the form [*base#*] *n* where *base* is a decimal number between two and thirty-six representing the arithmetic base and *n* is a number in that base. If *base* is omitted then base 10 is used.

An arithmetic expression uses the same syntax, precedence, and associativity of expression as the C language. All the integral operators, other than ++, --, ?, and , are supported. Variables can be referenced by name within an arithmetic expression without using the parameter substitution syntax. When a variable is referenced, its value is evaluated as an arithmetic expression.

An internal integer representation of a *variable* can be specified with the `-i` option of the `typeset` special command. Arithmetic evaluation is performed on the value of each assignment to a variable with the `-i` attribute. If you do not specify an arithmetic base, the first assignment to the variable determines the arithmetic base. This base is used when parameter substitution occurs.

Since many of the arithmetic operators require quoting, an alternative form of the `let` command is provided. For any command which begins with a `((`, all the characters until a matching `)` are treated as a quoted expression. More precisely, `((...))` is equivalent to `let "..."`.

Prompting

When used interactively, the shell prompts with the parameter expanded value of `PS1` before reading a command. If at any time a new-line is typed and further input is needed to complete a command, then the secondary prompt (i.e., the value of `PS2`) is issued.

Conditional Expressions

A *conditional expression* is used with the `[[` compound command to test attributes of files and to compare strings. Word splitting and file name generation are not performed on the words between `[[` and `]]`. Each expression can be constructed from one or more of the following unary or binary expressions:

<code>-a file</code>	True, if <i>file</i> exists.
<code>-b file</code>	True, if <i>file</i> exists and is a block special file.
<code>-c file</code>	True, if <i>file</i> exists and is a character special file.
<code>-d file</code>	True, if <i>file</i> exists and is a directory.
<code>-f file</code>	True, if <i>file</i> exists and is an ordinary file.
<code>-g file</code>	True, if <i>file</i> exists and is has its setgid bit set.
<code>-k file</code>	True, if <i>file</i> exists and is has its sticky bit set.
<code>-n string</code>	True, if length of <i>string</i> is non-zero.
<code>-o option</code>	True, if option named <i>option</i> is on.
<code>-p file</code>	True, if <i>file</i> exists and is a fifo special file or a pipe.
<code>-r file</code>	True, if <i>file</i> exists and is readable by current process.
<code>-s file</code>	True, if <i>file</i> exists and has size greater than zero.
<code>-t fildes</code>	True, if file descriptor number <i>fildes</i> is open and associated with a terminal device.
<code>-u file</code>	True, if <i>file</i> exists and has its setuid bit set.
<code>-w file</code>	True, if <i>file</i> exists and is writable by current process.
<code>-x file</code>	True, if <i>file</i> exists and is executable by current process. If <i>file</i> exists and is a directory, then the current process has permission to search in the directory.
<code>-z string</code>	True, if length of <i>string</i> is zero.
<code>-L file</code>	True, if <i>file</i> exists and is a symbolic link.
<code>-O file</code>	True, if <i>file</i> exists and is owned by the effective user id of this process.
<code>-G file</code>	True, if <i>file</i> exists and its group matches the effective group id of this process.
<code>-S file</code>	True, if <i>file</i> exists and is a socket.
<code>file1 -nt file2</code>	True, if <i>file1</i> exists and is newer than <i>file2</i> .
<code>file1 -ot file2</code>	True, if <i>file1</i> exists and is older than <i>file2</i> .

file1 **-ef** *file2* True, if *file1* and *file2* exist and refer to the same file.
string **=** *pattern* True, if *string* matches *pattern*.
string **!=** *pattern* True, if *string* does not match *pattern*.
string1 **<** *string2* True, if *string1* comes before *string2* based on ASCII value of their characters.
string1 **>** *string2* True, if *string1* comes after *string2* based on ASCII value of their characters.
exp1 **-eq** *exp2* True, if *exp1* is equal to *exp2*.
exp1 **-ne** *exp2* True, if *exp1* is not equal to *exp2*.
exp1 **-lt** *exp2* True, if *exp1* is less than *exp2*.
exp1 **-gt** *exp2* True, if *exp1* is greater than *exp2*.
exp1 **-le** *exp2* True, if *exp1* is less than or equal to *exp2*.
exp1 **-ge** *exp2* True, if *exp1* is greater than or equal to *exp2*.

In each of the above expressions, if *file* is of the form **/dev/fd/*n***, where *n* is an integer, then the test is applied to the open file whose descriptor number is *n*.

A compound expression can be constructed from these primitives by using any of the following, listed in decreasing order of precedence.

(*expression*) True, if *expression* is true. Used to group expressions.
!*expression* True if *expression* is false.
expression1 **&&** *expression2*
 True, if *expression1* and *expression2* are both true.
expression1 **||** *expression2*
 True, if either *expression1* or *expression2* is true.

Input/Output

Before a command is executed, its input and output may be redirected using a special notation interpreted by the shell. The following may appear anywhere in a simple-command or may precede or follow a *command* and are *not* passed on to the invoked command. Command and parameter substitution occur before *word* or *digit* is used except as noted below. File name generation occurs only if the pattern matches a single file, and blank interpretation is not performed.

<*word* Use file *word* as standard input (file descriptor 0).
>*word* Use file *word* as standard output (file descriptor 1). If the file does not exist then it is created. If the file exists, and the **noclobber** option is on, this causes an error; otherwise, it is truncated to zero length.
>|*word* Same as **>**, except that it overrides the **noclobber** option.
>>*word* Use file *word* as standard output. If the file exists then output is appended to it (by first seeking to the end-of-file); otherwise, the file is created.
<>*word* Open file *word* for reading and writing as standard input.
<< [-]*word* The shell input is read up to a line that is the same as *word*, or to an end-of-file. No parameter substitution, command substitution or file name generation is performed on *word*. The resulting document, called a *here-document*, becomes the standard input. If any character of *word* is

quoted, then no interpretation is placed upon the characters of the document; otherwise, parameter and command substitution occur, **\newline** is ignored, and **** must be used to quote the characters ****, **\$**, **'**, and the first character of *word*. If **-** is appended to **<<**, then all leading tabs are stripped from *word* and from the document.

<&digit The standard input is duplicated from file descriptor *digit* (see **dup(2)**). Similarly for the standard output using **>&digit**.

<&- The standard input is closed. Similarly for the standard output using **>&-**.

<&p The input from the co-process is moved to standard input.

>&p The output to the co-process is moved to standard output.

If one of the above is preceded by a digit, then the file descriptor number referred to is that specified by the digit (instead of the default 0 or 1). For example:

```
... 2>&1
```

means file descriptor 2 is to be opened for writing as a duplicate of file descriptor 1.

The order in which redirections are specified is significant. The shell evaluates each redirection in terms of the (*file descriptor*, *file*) association at the time of evaluation. For example:

```
... 1>fname 2>&1
```

first associates file descriptor 1 with file *fname*. It then associates file descriptor 2 with the file associated with file descriptor 1 (i.e. *fname*). If the order of redirections were reversed, file descriptor 2 would be associated with the terminal (assuming file descriptor 1 had been) and then file descriptor 1 would be associated with file *fname*.

If a command is followed by **&** and job control is not active, then the default standard input for the command is the empty file **/dev/null**. Otherwise, the environment for the execution of a command contains the file descriptors of the invoking shell as modified by input/output specifications.

Environment

The *environment* (see **environ(5)**) is a list of name-value pairs that is passed to an executed program in the same way as a normal argument list. The names must be *identifiers* and the values are character strings. The shell interacts with the environment in several ways. On invocation, the shell scans the environment and creates a variable for each name found, giving it the corresponding value and marking it *export*. Executed commands inherit the environment. If the user modifies the values of these variables or creates new ones, using the **export** or **typeset -x** commands they become part of the environment. The environment seen by any executed command is thus composed of any name-value pairs originally inherited by the shell, whose values may be modified by the current shell, plus any additions which must be noted in **export** or **typeset -x** commands.

The environment for any *simple-command* or *function* may be augmented by prefixing it with one or more variable assignments. A variable assignment argument is a word of the form *identifier=value*. Thus:

```
TERM=450 cmd args
and
(export TERM; TERM=450; cmd args)
```

are equivalent (as far as the above execution of *cmd* is concerned except for special commands listed below that are preceded with a dagger).

If the **-k** flag is set, *all* variable assignment arguments are placed in the environment, even if they occur after the command name. The following first prints **a=b c** and then **c**:

```
echo a=b c
set -k
echo a=b c
```

This feature is intended for use with scripts written for early versions of the shell and its use in new scripts is strongly discouraged. It is likely to disappear someday.

Functions

The **function** reserved word, described in the **Commands** section above, is used to define shell functions. Shell functions are read in and stored internally. Alias names are resolved when the function is read. Functions are executed like commands with the arguments passed as positional parameters. (See **Execution** below.)

Functions execute in the same process as the caller and share all files and present working directory with the caller. Traps caught by the caller are reset to their default action inside the function. A trap condition that is not caught or ignored by the function causes the function to terminate and the condition to be passed on to the caller. A trap on **EXIT** set inside a function is executed after the function completes in the environment of the caller. Ordinarily, variables are shared between the calling program and the function. However, the **typeset** special command used within a function defines local variables whose scope includes the current function and all functions it calls.

The special command **return** is used to return from function calls. Errors within functions return control to the caller.

The names of all functions can be listed with **typeset +f**. **typeset -f** lists all function names as well as the text of all functions. **typeset -f function-names** lists the text of the named functions only. Functions can be undefined with the **-f** option of the **unset** special command.

Ordinarily, functions are **unset** when the shell executes a shell script. The **-xf** option of the **typeset** command allows a function to be exported to scripts that are executed without a separate invocation of the shell. Functions that need to be defined across separate invocations of the shell should be specified in the **ENV** file with the **-xf** option of **typeset**.

Jobs

If the **monitor** option of the **set** command is turned on, an interactive shell associates a **job** with each pipeline. It keeps a table of current jobs, printed by the **jobs** command, and assigns them small integer numbers. When a job is started asynchronously with **&**, the shell prints a line which looks like:

[1] 1234

indicating that the **job**, which was started asynchronously, was job number 1 and had one (top-level) process, whose process id was 1234.

If you are running a job and wish to do something else you may hit the key **~Z** (CTRL-Z) which sends a **STOP** signal to the current job. The shell will then normally indicate that the job has been '**Stopped**', and print another prompt. You can then manipulate the state of this job, putting it in the background with the **bg** command, or run some other commands and then eventually bring the job back into the foreground with the foreground command **fg**. A **~Z** takes effect immediately and is like an interrupt in that pending output and unread input are discarded when it is typed.

A job being run in the background will stop if it tries to read from the terminal. Background jobs are normally allowed to produce output, but this can be disabled by giving the command "**stty tostop**". If you set this tty option, then background jobs will stop when they try to produce output like they do when they try to read input.

There are several ways to refer to **jobs** in the shell. A **job** can be referred to by the process id of any process of the **job** or by one of the following:

%number	The job with the given number.
%string	Any job whose command line begins with <i>string</i> .
%?string	Any job whose command line contains <i>string</i> .
%%	Current job.
%+	Equivalent to %% .
%-	Previous job.

The shell learns immediately whenever a process changes state. It normally informs you whenever a job becomes blocked so that no further progress is possible, but only just before it prints a prompt. This is done so that it does not otherwise disturb your work.

When the monitor mode is on, each background job that completes triggers any trap set for **CHLD**.

When you try to leave the shell while jobs are running or stopped, you will be warned that 'You have stopped(running) jobs.' You may use the **jobs** command to see what they are. If you do this or immediately try to exit again, the shell will not warn you a second time, and the stopped jobs will be terminated.

Signals The INT and QUIT signals for an invoked command are ignored if the command is followed by **&** and the **monitor** option is not active. Otherwise, signals have the values inherited by the shell from its parent (but see also the **trap** special command below).

Execution Each time a command is executed, the above substitutions are carried out. If the command name matches one of the **Special Commands** listed below, it is executed within the current shell process. Next, the command name is checked to see if it matches one of the user defined functions. If it does, the positional parameters are saved and then reset to the arguments of the **function** call. When the **function** completes or issues a **return**, the positional parameter list is restored and any trap set on **EXIT** within the function is executed. The value of a **function** is the value of the last command executed. A function is

also executed in the current shell process. If a command name is not a **special command** or a user defined **function**, a process is created and an attempt is made to execute the command via **exec(2)**.

The shell variable **PATH** defines the search path for the directory containing the command. Alternative directory names are separated by a colon (:). The default path is **/bin:/usr/bin:** (specifying **/bin**, **/usr/bin**, and the current directory in that order). The current directory can be specified by two or more adjacent colons, or by a colon at the beginning or end of the path list. If the command name contains a / then the search path is not used. Otherwise, each directory in the path is searched for an executable file. If the file has execute permission but is not a directory or an **a.out** file, it is assumed to be a file containing shell commands. A sub-shell is spawned to read it. All non-exported aliases, functions, and variables are removed in this case. A parenthesized command is executed in a sub-shell without removing non-exported quantities.

Command Re-entry

The text of the last **HISTSIZE** (default 128) commands entered from a terminal device is saved in a **history** file. The file **\$HOME/.sh_history** is used if the **HISTFILE** variable is not set or if the file it names is not writable. A shell can access the commands of all *interactive* shells which use the same named **HISTFILE**. The special command **fc** is used to list or edit a portion of this file. The portion of the file to be edited or listed can be selected by number or by giving the first character or characters of the command. A single command or range of commands can be specified. If you do not specify an editor program as an argument to **fc** then the value of the variable **FCEDIT** is used. If **FCEDIT** is not defined then **/bin/ed** is used. The edited command(s) is printed and re-executed upon leaving the editor. The editor name **-** is used to skip the editing phase and to re-execute the command. In this case a substitution parameter of the form *old=new* can be used to modify the command before execution. For example, if **r** is aliased to **'fc -e -'** then typing **'r bad=good c'** will re-execute the most recent command which starts with the letter **c**, replacing the first occurrence of the string **bad** with the string **good**.

In-line Editing Option

Normally, each command line entered from a terminal device is simply typed followed by a new-line (RETURN or LINE FEED). If either the **emacs**, **gmacs**, or **vi** option is active, the user can edit the command line. To be in either of these edit modes **set** the corresponding option. An editing option is automatically selected each time the **VISUAL** or **EDITOR** variable is assigned a value ending in either of these option names.

The editing features require that the user's terminal accept RETURN as carriage return without line feed and that a space (' ') must overwrite the current character on the screen.

The editing modes implement a concept where the user is looking through a window at the current line. The window width is the value of **COLUMNS** if it is defined, otherwise 80. If the window width is too small to display the prompt and leave at least 8 columns to enter input, the prompt is truncated from the left. If the line is longer than the window width minus two, a mark is displayed at the end of the window to notify the user. As the cursor moves and reaches the window boundaries the window will be centered about the cursor. The mark is a > if the line extends on the right side of the window, < if the line extends on the left, and * if the line extends on both sides of the window.

The search commands in each edit mode provide access to the history file. Only strings are matched, not patterns, although a leading `^` in the string restricts the match to begin at the first character in the line.

emacs Editing Mode

This mode is entered by enabling either the **emacs** or **gmacs** option. The only difference between these two modes is the way they handle `^T`. To edit, the user moves the cursor to the point needing correction and then inserts or deletes characters or words as needed. All the editing commands are control characters or escape sequences. The notation for control characters is caret (`^`) followed by the character. For example, `^F` is the notation for control **F**. This is entered by depressing 'f' while holding down the CTRL (control) key. The SHIFT key is *not* depressed. (The notation `^?` indicates the DEL (delete) key.)

The notation for escape sequences is **M-** followed by a character. For example, **M-f** (pronounced Meta f) is entered by depressing ESC (ascii **033**) followed by 'f'. (**M-F** would be the notation for ESC followed by SHIFT (capital) 'F'.)

All edit commands operate from any place on the line (not just at the beginning). Neither the RETURN nor the LINE FEED key is entered after edit commands except when noted.

<code>^F</code>	Move cursor forward (right) one character.
<code>M-f</code>	Move cursor forward one word. (The emacs editor's idea of a word is a string of characters consisting of only letters, digits and underscores.)
<code>^B</code>	Move cursor backward (left) one character.
<code>M-b</code>	Move cursor backward one word.
<code>^A</code>	Move cursor to start of line.
<code>^E</code>	Move cursor to end of line.
<code>]char</code>	Move cursor forward to character <i>char</i> on current line.
<code>M-]char</code>	Move cursor backward to character <i>char</i> on current line.
<code>^X^X</code>	Interchange the cursor and mark.
<i>erase</i>	(User defined erase character as defined by the stty(1) command, usually <code>^H</code> or <code>#</code> .) Delete previous character.
<code>^D</code>	Delete current character.
<code>M-d</code>	Delete current word.
<code>M-^H</code>	(Meta-backspace) Delete previous word.
<code>M-h</code>	Delete previous word.
<code>M-^?</code>	(Meta-DEL) Delete previous word (if your interrupt character is <code>^?</code> (DEL, the default) then this command will not work).
<code>^T</code>	Transpose current character with next character in emacs mode. Transpose two previous characters in gmacs mode.
<code>^C</code>	Capitalize current character.
<code>M-c</code>	Capitalize current word.
<code>M-l</code>	Change the current word to lower case.
<code>^K</code>	Delete from the cursor to the end of the line. If preceded by a numerical parameter whose value is less than the current cursor position, then delete from given position up to the cursor. If preceded by a numerical parameter whose value is greater than the current cursor position, then delete from cursor up to given cursor position.

^W	Kill from the cursor to the mark.
M-p	Push the region from the cursor to the mark on the stack.
kill	(User defined kill character as defined by the stty(1) command, usually ^G or @ .) Kill the entire current line. If two kill characters are entered in succession, all kill characters from then on cause a line feed (useful when using paper terminals).
^Y	Restore last item removed from line. (Yank item back to the line.)
^L	Line feed and print current line.
^@	(NULL character) Set mark.
M-space	(Meta space) Set mark.
^J	(New line) Execute the current line.
^M	(Return) Execute the current line.
eof	End-of-file character, normally ^D , is processed as an End-of-file only if the current line is null.
^P	Fetch previous command. Each time ^P is entered the previous command back in time is accessed. Moves back one line when not on the first line of a multi-line command.
M-<	Fetch the least recent (oldest) history line.
M->	Fetch the most recent (youngest) history line.
^N	Fetch next command line. Each time ^N is entered the next command line forward in time is accessed.
^Rstring	Reverse search history for a previous command line containing <i>string</i> . If a parameter of zero is given, the search is forward. <i>string</i> is terminated by a RETURN or NEW LINE. If string is preceded by a ^ , the matched line must begin with <i>string</i> . If <i>string</i> is omitted, then the next command line containing the most recent <i>string</i> is accessed. In this case a parameter of zero reverses the direction of the search.
^O	Operate – Execute the current line and fetch the next line relative to current line from the history file.
M-digits	(Escape) Define numeric parameter, the digits are taken as a parameter to the next command. The commands that accept a parameter are ^F , ^B , <i>erase</i> , ^C , ^D , ^K , ^R , ^P , ^N , ^] , M-. , M-^] , M-<u>_</u> , M-b , M-c , M-d , M-f , M-h , M-l and M-^H .
M-letter	Soft-key – Your alias list is searched for an alias by the name <i>_letter</i> and if an alias of this name is defined, its value will be inserted on the input queue. The <i>letter</i> must not be one of the above meta-functions.
M-[letter	Soft-key – Your alias list is searched for an alias by the name <i>__letter</i> and if an alias of this name is defined, its value will be inserted on the input queue. The can be used to program functions keys on many terminals.
M-.	The last word of the previous command is inserted on the line. If preceded by a numeric parameter, the value of this parameter determines which word to insert rather than the last word.
M-<u>_</u>	Same as M-.
M-*	Attempt file name generation on the current word. An asterisk is appended if the word doesn't match any file or contain any special

	pattern characters.
M-ESC	File name completion. Replaces the current word with the longest common prefix of all filenames matching the current word with an asterisk appended. If the match is unique, a / is appended if the file is a directory and a space is appended if the file is not a directory.
M-=	List files matching current word pattern if an asterisk were appended.
^U	Multiply parameter of next command by 4.
\	Escape next character. Editing characters, the user's erase, kill and interrupt (normally ^?) characters may be entered in a command line or in a search string if preceded by a \. The \ removes the next character's editing features (if any).
^V	Display version of the shell.
M-#	Insert a # at the beginning of the line and execute it. This causes a comment to be inserted in the history file.

vi Editing Mode

There are two typing modes. Initially, when you enter a command you are in the *input* mode. To edit, the user enters *control* mode by typing ESC (**033**) and moves the cursor to the point needing correction and then inserts or deletes characters or words as needed. Most control commands accept an optional repeat *count* prior to the command.

When in **vi** mode on most systems, canonical processing is initially enabled and the command will be echoed again if the speed is 1200 baud or greater and it contains any control characters or less than one second has elapsed since the prompt was printed. The ESC character terminates canonical processing for the remainder of the command and the user can then modify the command line. This scheme has the advantages of canonical processing with the type-ahead echoing of raw mode.

If the option **viraw** is also set, the terminal will always have canonical processing disabled. This mode is implicit for systems that do not support two alternate end of line delimiters, and may be helpful for certain terminals.

Input Edit Commands

By default the editor is in input mode.

<i>erase</i>	(User defined erase character as defined by the stty(1) command, usually ^H or #.) Delete previous character.
^W	Delete the previous blank separated word.
^D	Terminate the shell.
^V	Escape next character. Editing characters and the user's erase or kill characters may be entered in a command line or in a search string if preceded by a ^V. The ^V removes the next character's editing features (if any).
\	Escape the next <i>erase</i> or <i>kill</i> character.

Motion Edit Commands

These commands will move the cursor.

[count]l	Cursor forward (right) one character.
[count]w	Cursor forward one alpha-numeric word.
[count]W	Cursor to the beginning of the next word that follows a blank.
[count]e	Cursor to end of word.
[count]E	Cursor to end of the current blank delimited word.
[count]h	Cursor backward (left) one character.
[count]b	Cursor backward one word.
[count]B	Cursor to preceding blank separated word.
[count]	Cursor to column <i>count</i> .
[count]fc	Find the next character <i>c</i> in the current line.
[count]Fc	Find the previous character <i>c</i> in the current line.
[count]tc	Equivalent to f followed by h .
[count]Tc	Equivalent to F followed by l .
[count];	Repeats <i>count</i> times, the last single character find command, f , F , t , or T .
[count],	Reverses the last single character find command <i>count</i> times.
0	Cursor to start of line.
^	Cursor to first non-blank character in line.
\$	Cursor to end of line.
%	Moves to balancing (,) , { , } , [, or] . If cursor is not on one of the above characters, the remainder of the line is searched for the first occurrence of one of the above characters first.

Search Edit Commands

These commands access your command history.

[count]k	Fetch previous command. Each time k is entered the previous command back in time is accessed.
[count]-	Equivalent to k .
[count]j	Fetch next command. Each time j is entered the next command forward in time is accessed.
[count]+	Equivalent to j .
[count]G	The command number <i>count</i> is fetched. The default is the least recent history command.
/string	Search backward through history for a previous command containing <i>string</i> . <i>string</i> is terminated by a RETURN or NEW LINE. If <i>string</i> is preceded by a ^ , the matched line must begin with <i>string</i> . If <i>string</i> is NULL the previous string will be used.
?string	Same as / except that search will be in the forward direction.
n	Search for next match of the last pattern to / or ? commands.

**Text Modification
Edit Commands**

N Search for next match of the last pattern to / or ?, but in reverse direction. Search history for the *string* entered by the previous / command.

These commands will modify the line.

a Enter input mode and enter text after the current character.

A Append text to the end of the line. Equivalent to **Sa**.

[count]cmotion

c[count]motion

Delete current character through the character that *motion* would move the cursor to and enter input mode. If *motion* is **c**, the entire line will be deleted and input mode entered.

C Delete the current character through the end of line and enter input mode. Equivalent to **c\$**.

[count]s Delete *count* characters and enter input mode.

S Equivalent to **cc**.

D Delete the current character through the end of line. Equivalent to **d\$**.

[count]dmotion

d[count]motion

Delete current character through the character that *motion* would move to. If *motion* is **d**, the entire line will be deleted.

i Enter input mode and insert text before the current character.

I Insert text before the beginning of the line. Equivalent to **0i**.

[count]P Place the previous text modification before the cursor.

[count]p Place the previous text modification after the cursor.

R Enter input mode and replace characters on the screen with characters you type overlay fashion.

[count]rc Replace the *count* character(s) starting at the current cursor position with *c*, and advance the cursor.

[count]x Delete current character.

[count]X Delete preceding character.

[count]. Repeat the previous text modification command.

[count]~ Invert the case of the *count* character(s) starting at the current cursor position and advance the cursor.

[count]_ Causes the *count* word of the previous command to be appended and input mode entered. The last word is used if *count* is omitted.

***** Causes an ***** to be appended to the current word and file name generation attempted. If no match is found, it rings the bell. Otherwise, the word is replaced by the matching pattern and input mode is entered.

**** Filename completion. Replaces the current word with the longest

Other Edit Commands

Miscellaneous commands.

[count]ymotion

y[count]motion

Yank current character through character that *motion* would move the cursor to and puts them into the delete buffer. The text and cursor are unchanged.

Y Yanks from current position to end of line. Equivalent to *y\$*.

u Undo the last text modifying command.

U Undo all the text modifying commands performed on the line.

[count]v Returns the command `fc -e ${VISUAL:-${EDITOR:-vi}}` *count* in the input buffer. If *count* is omitted, then the current line is used.

^L Line feed and print current line. Has effect only in control mode.

^J (New line) Execute the current line, regardless of mode.

^M (Return) Execute the current line, regardless of mode.

If the first character of the command is a #, then this command deletes this # and each # that follows a newline. Otherwise, sends the line after inserting a # in front of each line in the command. Useful for causing the current line to be inserted in the history as a comment and removing comments from previous comment commands in the history file.

= List the file names that match the current word if an asterisk were appended it.

@letter Your alias list is searched for an alias by the name *_letter* and if an alias of this name is defined, its value will be inserted on the input queue for processing.

Special Commands

The following *simple-commands* are executed in the shell process. Input/Output redirection is permitted. Unless otherwise indicated, the output is written on file descriptor 1 and the exit status, when there is no syntax error, is zero. Commands that are preceded by one or two † (daggers) are treated specially in the following ways:

1. Variable assignment lists preceding the command remain in effect when the command completes.
2. I/O redirections are processed after variable assignments.
3. Errors cause a script that contains them to abort.
4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.

† : [*arg* ...]

The command only expands parameters.

† . *file* [*arg* ...]

Read the complete *file* then execute the commands. The commands are executed in the current shell environment. The search path specified by **PATH** is used to find the directory containing *file*. If any arguments *arg* are given, they become the positional parameters. Otherwise the positional parameters are unchanged. The exit status is the exit status of the last command executed.

†† **alias** [**-tx**] [*name*[=*value*]] ...

alias with no arguments prints the list of aliases in the form *name=value* on standard output. An *alias* is defined for each name whose *value* is given. A trailing space in *value* causes the next word to be checked for alias substitution. The **-t** flag is used to set and list tracked aliases. The value of a tracked alias is the full pathname corresponding to the given *name*. The value becomes undefined when the value of **PATH** is reset but the aliases remained tracked. Without the **-t** flag, for each *name* in the argument list for which no *value* is given, the name and value of the alias is printed. The **-x** flag is used to set or print *exported aliases*. An *exported alias* is defined for scripts invoked by name. The exit status is non-zero if a *name* is given, but no value, and no alias has been defined for the *name*.

bg [%*job*...]

This command is only on systems that support job control. Puts each specified *job* into the background. The current job is put in the background if *job* is not specified. See **Jobs** for a description of the format of *job*.

† **break** [*n*]

Exit from the enclosed **for**, **while**, **until**, or **select** loop, if any. If *n* is specified then **break** *n* levels.

† **continue** [*n*]

Resume the next iteration of the enclosed **for**, **while**, **until**, or **select** loop. If *n* is specified then resume at the *n*-th enclosed loop.

cd [*arg*]

cd *old new*

This command can be in either of two forms. In the first form it changes the current directory to *arg*. If *arg* is **-** the directory is changed to the previous directory. The shell variable **HOME** is the default *arg*. The variable **PWD** is set to the current directory. The shell variable **CDPATH** defines the search path for the directory containing *arg*. Alternative directory names are separated by a colon (:). The default path is null (specifying the current directory). Note that the current directory is specified by a null path name, which can appear immediately after the equal sign or between the colon delimiters anywhere else in the path list. If *arg* begins with a / then the search path is not used. Otherwise, each directory in the path is searched for *arg*.

The second form of **cd** substitutes the string *new* for the string *old* in the current directory name, **PWD** and tries to change to this new directory.

The **cd** command may not be executed by **rksh**.

echo [*arg* ...]

See **echo(1)** for usage and description.

† **eval** [*arg* ...]

The arguments are read as input to the shell and the resulting command(s) executed.

† **exec** [*arg* ...]

If *arg* is given, the command specified by the arguments is executed in place of this shell without creating a new process. Input/output arguments may appear and affect the current process. If no arguments are given the effect of this command is to modify file descriptors as prescribed by the input/output redirection list. In this case, any file descriptor numbers greater than 2 that are opened with this mechanism are closed when invoking another program.

† **exit** [*n*]

Causes the shell to exit with the exit status specified by *n*. The value will be the least significant 8 bits of the specified status. If *n* is omitted then the exit status is that of the last command executed. When **exit** occurs when executing a trap, the last command refers to the command that executed before the trap was invoked. An end-of-file will also cause the shell to exit except for a shell which has the **ignoreeof** option (See **set** below) turned on.

†† **export** [*name*[=*value*]] ...

The given *names* are marked for automatic export to the **environment** of subsequently-executed commands.

fc [**-e** *ename*] **-nlr** [*first* [*last*]]

fc **-e** - [*old=new*] [*command*]

In the first form, a range of commands from *first* to *last* is selected from the last **HISTSIZE** commands that were typed at the terminal. The arguments *first* and *last* may be specified as a number or as a string. A string is used to locate the most recent command starting with the given string. A negative number is used as an offset to the current command number. If the **-l** flag is selected, the commands are listed on standard output. Otherwise, the editor program *ename* is invoked on a file containing these keyboard commands. If *ename* is not supplied, then the value of the variable **FCEDIT** (default **/bin/ed**) is used as the editor. When editing is complete, the edited command(s) is executed. If *last* is not specified then it will be set to *first*. If *first* is not specified the default is the previous command for editing and **-16** for listing. The flag **-r** reverses the order of the commands and the flag **-n** suppresses command numbers when listing. In the second form the *command* is re-executed after the substitution *old=new* is performed.

fg [%*job*...]

This command is only on systems that support job control. Each *job* specified is brought to the foreground. Otherwise, the current job is brought into the foreground. See **Jobs** for a description of the format of *job*.

getopts *optstring name* [*arg* ...]

Checks *arg* for legal options. If *arg* is omitted, the positional parameters are used. An option argument begins with a + or a -. An option not beginning with + or - or the argument -- ends the options. *optstring* contains the letters that **getopts** recognizes. If a letter is followed by a :, that option is expected to have an argument. The options can be separated from the argument by blanks.

getopts places the next option letter it finds inside variable *name* each time it is invoked with a + prepended when *arg* begins with a +. The index of the next *arg* is stored in **OPTIND**. The option argument, if any, gets stored in **OPTARG**.

A leading : in *optstring* causes **getopts** to store the letter of an invalid option in **OPTARG**, and to set *name* to ? for an unknown option and to : when a required option is missing. Otherwise, **getopts** prints an error message. The exit status is non-zero when there are no more options. See **getoptcv**(1) for usage and description.

jobs [**-lnp**] [%*job* ...]

Lists information about each given job; or all active jobs if *job* is omitted. The **-l** flag lists process ids in addition to the normal information. The **-n** flag only displays jobs that have stopped or exited since last notified. The **-p** flag causes only the process group to be listed. See **Jobs** for a description of the format of *job*.

kill [**-sig**] %*job* ...

kill -l Sends either the **TERM** (terminate) signal or the specified signal to the specified jobs or processes. Signals are either given by number or by names (as given in **signal**(5) stripped of the prefix "SIG" with the exception that **SIGCHD** is named **CHLD**). If the signal being sent is **TERM** (terminate) or **HUP** (hangup), then the job or process will be sent a **CONT** (continue) signal if it is stopped. The argument *job* can be the process id of a process that is not a member of one of the active jobs. See **Jobs** for a description of the format of *job*. In the second form, **kill -l**, the signal numbers and names are listed.

let *arg* ...

Each *arg* is a separate *arithmetic expression* to be evaluated. See **Arithmetic Evaluation** above, for a description of arithmetic expression evaluation.

The exit status is **0** if the value of the last expression is non-zero, and **1** otherwise.

† **newgrp** [*arg* ...]

Equivalent to **exec /bin/newgrp** *arg* ...

print [**-Rnrpsu**[*n*]] [*arg* ...]

The shell output mechanism. With no flags or with flag **-** or **--**, the arguments are printed on standard output as described by **echo**(1). The exit status is **0**, unless the output file is not open for writing.

-n suppresses **new-line** from being added to the output.

-R

-r (raw mode) ignore the escape conventions of **echo**. The **-R** option

- will print all subsequent arguments and options other than **-n**.
- p** causes the arguments to be written onto the pipe of the process spawned with **| &** instead of standard output.
 - s** causes the arguments to be written onto the history file instead of standard output.
 - u [n]** flag can be used to specify a one digit file descriptor unit number *n* on which the output will be placed. The default is **1**.

pwd Equivalent to **print -r - \$PWD print -r - \$PWD**

read [-prsu [n]] [name?prompt] [name ...]

The shell input mechanism. One line is read and is broken up into fields using the characters in **IFS** as separators. The escape character, (****), is used to remove any special meaning for the next character and for line continuation. In raw mode, **-r**, the **** character is not treated specially. The first field is assigned to the first *name*, the second field to the second *name*, etc., with leftover fields assigned to the last *name*. The **-p** option causes the input line to be taken from the input pipe of a process spawned by the shell using **| &**. If the **-s** flag is present, the input will be saved as a command in the history file. The flag **-u** can be used to specify a one digit file descriptor unit *n* to read from. The file descriptor can be opened with the **exec** special command. The default value of *n* is **0**. If *name* is omitted then **REPLY** is used as the default *name*. The exit status is **0** unless the input file is not open for reading or an end-of-file is encountered. An end-of-file with the **-p** option causes cleanup for this process so that another can be spawned. If the first argument contains a **?**, the remainder of this word is used as a *prompt* on standard error when the shell is interactive. The exit status is **0** unless an end-of-file is encountered.

†† readonly [name[=value]] ...

The given *names* are marked **readonly** and these names cannot be changed by subsequent assignment.

† return [n]

Causes a shell function or **'.** script to return to the invoking script with the return status specified by *n*. The value will be the least significant 8 bits of the specified status. If *n* is omitted then the return status is that of the last command executed. If **return** is invoked while not in a function or a **'.** script, then it is the same as an **exit**.

set [±aefhkmnopstuvx] [±o option] ... [±A name] [arg ...]

The flags for this command have meaning as follows:

- A** Array assignment. Unset the variable *name* and assign values sequentially from the list *arg*. If **+A** is used, the variable *name* is not unset first.
- a** All subsequent variables that are defined are automatically exported.
- e** If a command has a non-zero exit status, execute the **ERR** trap, if set, and exit. This mode is disabled while reading profiles.
- f** Disables file name generation.

- h** Each command becomes a tracked alias when first encountered.
- k** All variable assignment arguments are placed in the environment for a command, not just those that precede the command name.
- m** Background jobs will run in a separate process group and a line will print upon completion. The exit status of background jobs is reported in a completion message. On systems with job control, this flag is turned on automatically for interactive shells.
- n** Read commands and check them for syntax errors, but do not execute them. Ignored for interactive shells.
- o** The following argument can be one of the following option names:
 - allexport** Same as **-a**.
 - errexit** Same as **-e**.
 - bgnice** All background jobs are run at a lower priority. This is the default mode.
 - emacs** Puts you in an **emacs** style in-line editor for command entry.
 - gmacs** Puts you in a **gmacs** style in-line editor for command entry.
 - ignoreeof** The shell will not exit on end-of-file. The command **exit** must be used.
 - keyword** Same as **-k**.
 - markdirs** All directory names resulting from file name generation have a trailing **/** appended.
 - monitor** Same as **-m**.
 - noclobber** Prevents redirection **>** from truncating existing files. Require **>|** to truncate a file when turned on.
 - noexec** Same as **-n**.
 - noglob** Same as **-f**.
 - nolog** Do not save function definitions in history file.
 - nounset** Same as **-u**.
 - privileged** Same as **-p**.
 - verbose** Same as **-v**.
 - trackall** Same as **-h**.
- vi** Puts you in insert mode of a **vi** style in-line editor until you hit escape character **033**. This puts you in control mode. A return

sends the line.

viraw Each character is processed as it is typed in **vi** mode.

xtrace Same as **-x**.

If no option name is supplied then the current option settings are printed.

- p** Disables processing of the **\$HOME/.profile** file and uses the file **/etc/suid_profile** instead of the **ENV** file. This mode is on whenever the effective uid is not equal to the real uid, or when the effective gid is not equal to the real gid. Turning this off causes the effective uid and gid to be set to the real uid and gid.
- s** Sort the positional parameters lexicographically.
- t** Exit after reading and executing one command.
- u** Treat unset parameters as an error when substituting.
- v** Print shell input lines as they are read.
- x** Print commands and their arguments as they are executed.
- Turns off **-x** and **-v** flags and stops examining arguments for flags.
- Do not change any of the flags; useful in setting **\$1** to a value beginning with **-**. If no arguments follow this flag then the positional parameters are unset.

Using **+** rather than **-** causes these flags to be turned off. These flags can also be used upon invocation of the shell. The current set of flags may be found in **\$-**. Unless **-A** is specified, the remaining arguments are positional parameters and are assigned, in order, to **\$1 \$2 ...**. If no arguments are given then the names and values of all variables are printed on the standard output.

† **shift** [*n*]

The positional parameters from **\$n+1 \$n+1 ...** are renamed **\$1 ...**, default *n* is 1. The parameter *n* can be any arithmetic expression that evaluates to a non-negative number less than or equal to **\$#**.

stop %*jobid* ...

stop *pid* ...

stop stops the execution of a background job(s) by using its *jobid*, or of any process by using its *pid*. (see **ps(1)**).

suspend

Stops the execution of the current shell (but not if it is the login shell).

† **times** Print the accumulated user and system times for the shell and for processes run from the shell.

† **trap** [*arg*] [*sig*] ...

arg is a command to be read and executed when the shell receives signal(s) *sig*. (Note that *arg* is scanned once when the trap is set and once when the trap is taken.) Each *sig* can be given as a number or as the name of the signal. **trap**

commands are executed in order of signal number. Any attempt to set a trap on a signal that was ignored on entry to the current shell is ineffective. If *arg* is omitted or is `-`, then the trap(s) for each *sig* are reset to their original values. If *arg* is the NULL (the empty string, e.g., `""`) string then this signal is ignored by the shell and by the commands it invokes. If *sig* is `ERR` then *arg* will be executed whenever a command has a non-zero exit status. If *sig* is `DEBUG` then *arg* will be executed after each command. If *sig* is `0` or `EXIT` and the `trap` statement is executed inside the body of a function, then the command *arg* is executed after the function completes. If *sig* is `0` or `EXIT` for a `trap` set outside any function then the command *arg* is executed on exit from the shell. The `trap` command with no arguments prints a list of commands associated with each signal number.

†† **typeset** [`±HLRZfilrtux`[*n*]] [*name*[=*value*]] ...

Sets attributes and values for shell variables and functions. When **typeset** is invoked inside a function, a new instance of the variables *name* is created. The variables *value* and *type* are restored when the function completes. The following list of attributes may be specified:

- H** This flag provides UNIX to host-name file mapping on non-UNIX machines.
- L** Left justify and remove leading blanks from *value*. If *n* is non-zero it defines the width of the field; otherwise, it is determined by the width of the value of first assignment. When the variable is assigned to, it is filled on the right with blanks or truncated, if necessary, to fit into the field. Leading zeros are removed if the **-Z** flag is also set. The **-R** flag is turned off.
- R** Right justify and fill with leading blanks. If *n* is non-zero it defines the width of the field, otherwise it is determined by the width of the value of first assignment. The field is left filled with blanks or truncated from the end if the variable is reassigned. The **-L** flag is turned off.
- Z** Right justify and fill with leading zeros if the first non-blank character is a digit and the **-L** flag has not been set. If *n* is non-zero it defines the width of the field; otherwise, it is determined by the width of the value of first assignment.
- f** The names refer to function names rather than variable names. No assignments can be made and the only other valid flags are **-t**, **-u** and **-x**. The flag **-t** turns on execution tracing for this function. The flag **-u** causes this function to be marked undefined. The `FPATH` variable will be searched to find the function definition when the function is referenced. The flag **-x** allows the function definition to remain in effect across shell procedures invoked by name.
- i** Parameter is an integer. This makes arithmetic faster. If *n* is non-zero it defines the output arithmetic base; otherwise, the first assignment determines the output base.
- l** All upper-case characters are converted to lower-case. The upper-case flag, **-u** is turned off.
- r** The given *names* are marked **readonly** and these names cannot be

- changed by subsequent assignment.
- t** Tags the variables. Tags are user definable and have no special meaning to the shell.
 - u** All lower-case characters are converted to upper-case characters. The lower-case flag, **-l** is turned off.
 - x** The given *names* are marked for automatic export to the **environment** of subsequently-executed commands.

The **-i** attribute can not be specified along with **-R**, **-L**, **-Z**, or **-f**.

Using **+** rather than **-** causes these flags to be turned off. If no *name* arguments are given but flags are specified, a list of *names* (and optionally the *values*) of the *variables* which have these flags set is printed. (Using **+** rather than **-** keeps the values from being printed.) If no *names* and flags are given, the *names* and *attributes* of all *variables* are printed.

ulimit [**-HSacdfmnpstv**] [*limit*]

Set or display a resource limit. The available resources limits are listed below. Many systems do not contain one or more of these limits. The limit for a specified resource is set when *limit* is specified. The value of *limit* can be a number in the unit specified below with each resource, or the value **unlimited**. The **H** and **S** flags specify whether the hard limit or the soft limit for the given resource is set. A hard limit cannot be increased once it is set. A soft limit can be increased up to the value of the hard limit. If neither the **H** or **S** options is specified, the limit applies to both. The current resource limit is printed when *limit* is omitted. In this case the soft limit is printed unless **H** is specified. When more than one resource is specified, then the limit name and unit is printed before the value.

- a** Lists all of the current resource limits.
- c** The number of 512-byte blocks on the size of core dumps.
- d** The number of K-bytes on the size of the data area.
- f** The number of 512-byte blocks on files written by child processes (files of any size may be read).
- m** The number of K-bytes on the size of physical memory.
- n** The number of file descriptors plus 1.
- p** The number of 512-byte blocks for pipe buffering.
- s** The number of K-bytes on the size of the stack area.
- t** The number of seconds to be used by each process.
- v** The number of K-bytes for virtual memory.

If no option is given, **-f** is assumed.

umask [*mask*]

The user file-creation mask is set to *mask* (see **umask**(2)). *mask* can either be an octal number or a symbolic value as described in **chmod**(1). If a symbolic value is given, the new **umask** value is the complement of the result of applying *mask* to the complement of the previous umask value. If *mask* is omitted, the current value of the mask is printed.

unalias *name...*

The *aliases* given by the list of *names* are removed from the *alias* list.

unset [*-f*] *name...*

The variables given by the list of *names* are unassigned, i.e., their values and attributes are erased. **readonly** variables cannot be unset. If the *-f* flag is set, then the names refer to *function* names. Unsetting **ERRNO**, **LINENO**, **MAILCHECK**, **OPTARG**, **OPTIND**, **RANDOM**, **SECONDS**, **TMOU**, and **_** removes their special meaning even if they are subsequently assigned to.

† **wait** [*job*]

Wait for the specified *job* and report its termination status. If *job* is not given then all currently active child processes are waited for. The exit status from this command is that of the process waited for. See **Jobs** for a description of the format of *job*.

whence [*-pv*] *name...*

For each *name*, indicate how it would be interpreted if used as a command name.

The *-v* flag produces a more verbose report.

The *-p* flag does a path search for *name* even if name is an alias, a function, or a reserved word.

Invocation

If the shell is invoked by **exec(2)**, and the first character of argument zero (**\$0**) is *-*, then the shell is assumed to be a **login** shell and commands are read from **/etc/profile** and then from either **.profile** in the current directory or **\$HOME/.profile**, if either file exists. Next, commands are read from the file named by performing parameter substitution on the value of the environment variable **ENV** if the file exists. If the *-s* flag is not present and *arg* is, then a path search is performed on the first *arg* to determine the name of the script to execute. The script *arg* must have read permission and any **setuid** and **setgid** settings will be ignored. If the script is not found on the path, *arg* is processed as if it named a builtin command or function. Commands are then read as described below; the following flags are interpreted by the shell when it is invoked:

-c string If the *-c* flag is present then commands are read from *string*.

-s If the *-s* flag is present or if no arguments remain then commands are read from the standard input. Shell output, except for the output of the **Special Commands** listed above, is written to file descriptor 2.

-i If the *-i* flag is present or if the shell input and output are attached to a terminal (as told by **ioctl(2)**) then this shell is *interactive*. In this case **TERM** is ignored (so that **kill 0** does not kill an interactive shell) and **INTR** is caught and ignored (so that **wait** is interruptible). In all cases, **QUIT** is ignored by the shell.

-r If the *-r* flag is present the shell is a restricted shell.

The remaining flags and arguments are described under the **set** command above.

rksh Only

rksh is used to set up login names and execution environments whose capabilities are more controlled than those of the standard shell. The actions of **rksh** are identical to those of **ksh**, except that the following are disallowed:

- changing directory (see **cd**(1))
- setting the value of **SHELL**, **ENV**, or **PATH**
- specifying path or command names containing /
- redirecting output (>, >|, <>, and >>)
- changing group (see **newgrp**(1M)).

The restrictions above are enforced after **.profile** and the **ENV** files are interpreted.

When a command to be executed is found to be a shell procedure, **rksh** invokes **ksh** to execute it. Thus, it is possible to provide to the end-user shell procedures that have access to the full power of the standard shell, while imposing a limited menu of commands; this scheme assumes that the end-user does not have write and execute permissions in the same directory.

The net effect of these rules is that the writer of the **.profile** has complete control over user actions, by performing guaranteed setup actions and leaving the user in an appropriate directory (probably *not* the login directory).

The system administrator often sets up a directory of commands (i.e., **/usr/rbin**) that can be safely invoked by **rksh**.

EXIT STATUS

Errors detected by the shell, such as syntax errors, cause the shell to return a non-zero exit status. Otherwise, the shell returns the exit status of the last command executed (see also the **exit** command above). If the shell is being used non-interactively then execution of the shell file is abandoned. Run time errors detected by the shell are reported by printing the command or function name and the error condition. If the line number that the error occurred on is greater than one, then the line number is also printed in square brackets ([]) after the command or function name.

FILES

/etc/profile
/etc/suid_profile
\$HOME/.profile
/tmp/sh*
/dev/null

SEE ALSO

cat(1), **chmod**(1), **cut**(1), **echo**(1), **env**(1), **getoptcv**(1), **paste**(1), **ps**(1), **shell_builtins**(1), **stty**(1), **vi**(1), **newgrp**(1M), **dup**(2), **exec**(2), **fork**(2), **ioctl**(2), **lseek**(2), **pipe**(2), **umask**(2), **ulimit**(2), **wait**(2), **rand**(3C), **signal**(3C), **a.out**(4), **profile**(4), **environ**(5), **signal**(5)

Morris I. Bolsky and David G. Korn, *The KornShell Command and Programming Language*, Prentice Hall, 1989.

NOTES

If a command which is a *tracked alias* is executed, and then a command with the same name is installed in a directory in the search path before the directory where the original command was found, the shell will continue to **exec** the original command. Use the **-t** option of the **alias** command to correct this situation.

Some very old shell scripts contain a `^` as a synonym for the pipe character `|`.

Using the `fc` built-in command within a compound command will cause the whole command to disappear from the history file.

The built-in command `. file` reads the whole file before any commands are executed. Therefore, **alias** and **unalias** commands in the file will not apply to any functions defined in the file.

Traps are not processed while a job is waiting for a foreground process. Thus, a trap on **CHLD** won't be executed until the foreground job terminates.

NAME	ksrvtgt – fetch and store Kerberos ticket-granting ticket using a service key						
SYNOPSIS	<code>/usr/bin/ksrvtgt name instance [[realm] srvtab]</code>						
AVAILABILITY	SUNWcsu						
DESCRIPTION	<p>ksrvtgt retrieves a ticket-granting ticket with a lifetime of five minutes for the principal <i>name.instance@realm</i> (or <i>name.instance@localrealm</i> if <i>realm</i> is not supplied on the command line), decrypts the response using the service key found in the file <i>srvtab</i> (or in <code>/etc/srvtab</code> if <i>srvtab</i> is not specified on the command line), and stores the ticket in the standard ticket cache.</p> <p>This command is intended primarily for use in shell scripts and other batch-type facilities.</p>						
DIAGNOSTICS	<p>Generic kerberos failure (kfailure) can indicate a whole range of problems, the most common of which is the inability to read the service key file.</p>						
FILES	<table><tr><td><code>/etc/krb.conf</code></td><td>to get the name of the local realm.</td></tr><tr><td><code>/tmp/tktuid</code></td><td>The default ticket file.</td></tr><tr><td><code>/etc/srvtab</code></td><td>The default service key file.</td></tr></table>	<code>/etc/krb.conf</code>	to get the name of the local realm.	<code>/tmp/tktuid</code>	The default ticket file.	<code>/etc/srvtab</code>	The default service key file.
<code>/etc/krb.conf</code>	to get the name of the local realm.						
<code>/tmp/tktuid</code>	The default ticket file.						
<code>/etc/srvtab</code>	The default service key file.						
SEE ALSO	kdestroy(1) , kerberos(1) , kinit(1) , klist(1)						

NAME	last – display login and logout information about users and terminals
SYNOPSIS	last [-n <i>number</i> -number] [-f <i>filename</i>] [<i>name</i> <i>tty</i>] ...
AVAILABILITY	SUNWesu
DESCRIPTION	<p>The last command looks in the /var/adm/wtmpx, file which records all logins and logouts, for information about a user, a terminal or any group of users and terminals. Arguments specify names of users or terminals of interest. Names of terminals may be given fully or abbreviated. For example last 10 is the same as last term/10. If multiple arguments are given, the information which applies to any of the arguments is printed. For example last root console lists all of root's sessions as well as all sessions on the console terminal. last displays the sessions of the specified users and terminals, most recent first, indicating the times at which the session began, the duration of the session, and the terminal which the session took place on. If the session is still continuing or was cut short by a reboot, last so indicates.</p> <p>The pseudo-user reboot logs in at reboots of the system, thus</p> <p style="padding-left: 40px;">last reboot</p> <p>will give an indication of mean time between reboot.</p> <p>last with no arguments displays a record of all logins and logouts, in reverse order. If last is interrupted, it indicates how far the search has progressed in /var/adm/wtmpx. If interrupted with a quit signal (generated by a CTRL-\\) last indicates how far the search has progressed so far, and the search continues.</p>
OPTIONS	<p>-n <i>number</i> / -number Limit the number of entries displayed to that specified by <i>number</i>. These options are identical; the -number option is provided as a transition tool only and will be removed in future releases.</p> <p>-f <i>filename</i> Use <i>filename</i> as the name of the accounting file instead of /var/adm/wtmpx.</p>
FILES	/var/adm/wtmpx accounting file
ENVIRONMENT	Date and time format is based on locale specified by the LC_ALL , LC_TIME or LANG environments, in that order of priority.
SEE ALSO	utmp(4)

NAME	lastcomm – display the last commands executed, in reverse order
SYNOPSIS	lastcomm [<i>command-name</i>] ... [<i>user-name</i>] ... [<i>terminal-name</i>] ...
AVAILABILITY	SUNWesu
DESCRIPTION	<p>The lastcomm command gives information on previously executed commands. lastcomm with no arguments displays information about all the commands recorded during the current accounting file's lifetime. If called with arguments, lastcomm only displays accounting entries with a matching <i>command-name</i>, <i>user-name</i>, or <i>terminal-name</i>.</p> <p>If <i>terminal-name</i> is `-' there was no controlling TTY for the process. The process was probably executed during boot time. If <i>terminal-name</i> is `??', the controlling TTY could not be decoded into a printable name.</p>
EXAMPLES	<p>The command:</p> <pre>example% lastcomm a.out root term/01</pre> <p>produces a listing of all the executions of commands named a.out, by user root while using the terminal term/01.</p> <p>The command:</p> <pre>example% lastcomm root</pre> <p>produces a listing of all the commands executed by user root.</p> <p>For each process entry, lastcomm displays the following items of information:</p> <ul style="list-style-type: none"> • The command name under which the process was called. • One or more flags indicating special information about the process. The flags have the following meanings: <ul style="list-style-type: none"> F The process performed a fork but not an exec. S The process ran as a set-user-id program. • The name of the user who ran the process. • The terminal which the user was logged in on at the time (if applicable). • The amount of CPU time used by the process (in seconds). • The date and time the process exited.
FILES	<i>/var/adm/pacct</i> accounting file
SEE ALSO	last(1) , sigvec(3B) , acct(4) , core(4)

NAME	ld – link editor for object files
SYNOPSIS	ld [-a -r] [-b] [-G] [-i] [-m] [-s] [-t] [-V] [-B dynamic static] [-B symbolic] [-d y n] [-D token] [-e epsym] [-F name] [-h name] [-I name] [-L path] [-l x] [-M mapfile] [-o outfile] [-Q y n] [-R path] [-u symname] [-Y P,dirlist] [-z defs nodefs] [-z muldefs] [-z text] filename ...
DESCRIPTION	<p>The ld command combines relocatable object files, performs relocation, and resolves external symbols. ld operates in two modes, static or dynamic, as governed by the -d option. In static mode, -dn, relocatable object files given as arguments are combined to produce an executable object file; if the -r option is specified, relocatable object files are combined to produce one relocatable object file. In dynamic mode, -dy, the default, relocatable object files given as arguments are combined to produce an executable object file that will be linked at execution with any shared object files given as arguments; if the -G option is specified, relocatable object files are combined to produce a shared object. In all cases, the output of ld is left in a.out by default.</p> <p>If any argument is a library, it is searched exactly once at the point it is encountered in the argument list. The library may be either a relocatable archive or a shared object. For an archive library, only those routines defining an unresolved external reference are loaded. The archive library symbol table (see ar(4)) is searched sequentially with as many passes as are necessary to resolve external references that can be satisfied by library members. Thus, the ordering of members in the library is functionally unimportant, unless there exist multiple library members defining the same external symbol. A shared object consists of a single entity all of whose references must be resolved within the executable being built or within other shared objects with which it is linked.</p>
OPTIONS	<p>-a In static mode only, produce an executable object file; give errors for undefined references. This is the default behavior for static mode. -a may not be used with the -r option.</p> <p>-r Combine relocatable object files to produce one relocatable object file. ld will not complain about unresolved references. This option cannot be used in dynamic mode or with -a.</p> <p>-b In dynamic mode only, when creating an executable, do not do special processing for relocations that reference symbols in shared objects. Without the -b option, the link editor creates special position-independent relocations for references to functions defined in shared objects and arranges for data objects defined in shared objects to be copied into the memory image of the executable by the runtime linker. With the -b option, the output code may be more efficient, but it will be less sharable.</p> <p>-G In dynamic mode only, produce a shared object. Undefined symbols are allowed.</p>

- i** Ignore **LD_LIBRARY_PATH** setting. This option is useful when an **LD_LIBRARY_PATH** setting is in effect to influence the runtime library search, which would interfere with the link editing being performed.
- m** Produce a memory map or listing of the input/output sections, together with any non-fatal multiply defined symbols, on the standard output.
- s** Strip symbolic information from the output file. Any debugging information, that is *.debug*, *.line*, and *.stab* sections, and their associated relocation entries will be removed. Except for relocatable files or shared objects, the symbol table and string table sections will also be removed from the output object file.
- t** Turn off the warning about multiply defined symbols that are not the same size.
- V** Output a message giving information about the version of **ld** being used.
- B dynamic | static** Options governing library inclusion. **-Bdynamic** is valid in dynamic mode only. These options may be specified any number of times on the command line as toggles: if the **-Bstatic** option is given, no shared objects will be accepted until **-Bdynamic** is seen. See also the **-l** option.
- B symbolic** In dynamic mode only, when building a shared object, bind references to global symbols to their definitions within the object, if definitions are available. Normally, references to global symbols within shared objects are not bound until runtime, even if definitions are available, so that definitions of the same symbol in an executable or other shared objects can override the object's own definition. **ld** will issue warnings for undefined symbols unless **-z defs** overrides.
- d y | n** When **-dy**, the default, is specified, **ld** uses dynamic linking; when **-dn** is specified, **ld** uses static linking. Also see **-B dynamic | static**.
- D token,token,...** Print debugging information, as specified by each *token*, to the standard error. The special token *help* indicates the full list of tokens available.
- e epsym** Set the entry point address for the output file to be that of the symbol *epsym*.
- F name** Useful only when building a shared object. Specifies that the symbol table of the shared object is used as a "filter" on the symbol table of the shared object specified by *name*.
- h name** In dynamic mode only, when building a shared object, record *name* in the object's dynamic section. *name* will be recorded in executables that are linked with this object rather than the object's UNIX System file name. Accordingly, *name* will be used by the runtime linker as the name of the shared object to search for at runtime.

- I *name*** When building an executable, use *name* as the path name of the interpreter to be written into the program header. The default in static mode is no interpreter; in dynamic mode, the default is the name of the runtime linker, `/usr/lib/ld.so.1`. Either case may be overridden by **-Iname**. `exec` will load this interpreter when it loads the `a.out` and will pass control to the interpreter rather than to the `a.out` directly.
- L *path*** Add *path* to the library search directories. `ld` searches for libraries first in any directories specified by the **-L** options, and then in the standard directories. This option is useful only if it precedes the **-I** options to which it applies on the command line. The environment variable `LD_LIBRARY_PATH` may be used to supplement the library search path (see `LD_LIBRARY_PATH` below).
- I *x*** Search a library `libx.so` or `libx.a`, the conventional names for shared object and archive libraries, respectively. In dynamic mode, unless the **-Bstatic** option is in effect, `ld` searches each directory specified in the library search path for a file `libx.so` or `libx.a`. The directory search stops at the first directory containing either. `ld` chooses the file ending in `.so` if **-I_x** expands to two files whose names are of the form `libx.so` and `libx.a`. If no `libx.so` is found, then `ld` accepts `libx.a`. In static mode, or when the **-Bstatic** option is in effect, `ld` selects only the file ending in `.a`. A library is searched when its name is encountered, so the placement of **-I** is significant.
- M *mapfile*** Read *mapfile* as a text file of directives to `ld`. See *Linker and Libraries Guide* for description of mapfiles.
- o *outfile*** Produce an output object file named *outfile*. The name of the default object file is `a.out`.
- Q *y* | *n*** Under **-Q_y**, an **ident** string is added to the `.comment` section of the output file to identify the version of the link editor used to create the file. This will result in multiple `ld ident`s when there have been multiple linking steps, such as when using `ld -r`. This is identical with the default action of the `cc` command. **-Q_n** suppresses version identification.
- R *path*** A colon-separated list of directories used to specify library search directories to the runtime linker. If present and not null, it is recorded in the output object file and passed to the runtime linker. Multiple instances of this option are concatenated together with each *path* separated by a colon.
- u *symname*** Enter *symname* as an undefined symbol in the symbol table. This is useful for loading entirely from an archive library, since initially the symbol table is empty and an unresolved reference is needed to force the loading of the first routine. The placement of this option on the command line is significant; it must be placed before the library that will define the symbol.
- Y *P,dirlist*** Change the default directories used for finding libraries. *dirlist* is a colon-separated path list.

- z defs** Force a fatal error if any undefined symbols remain at the end of the link. This is the default when building an executable. It is also useful when building a shared object to assure that the object is self-contained, that is, that all its symbolic references are resolved internally.
- z nodefs** Allow undefined symbols. This is the default when building a shared object. When used with executables, the behavior of references to such "undefined symbols" is unspecified.
- z muldefs** Allows multiple symbol definitions. By default, multiple symbol definitions occurring between relocatable objects will result in a fatal error condition. This option suppresses the error condition, and allows the first symbol definition to be taken.
- z text** In dynamic mode only, force a fatal error if any relocations against non-writable, allocatable sections remain.

ENVIRONMENT**LD_LIBRARY_PATH**

A list of directories in which to search for libraries specified with the **-l** option. Multiple directories are separated by a colon. In the most general case, it will contain two directory lists separated by a semicolon:

dirlist1;dirlist2

If **ld** is called with any number of occurrences of **-L**, as in:

ld ... -Lpath1 ... -Lpathn ...

then the search path ordering is:

dirlist1 path1 ... pathn dirlist2 LIBPATH

When the list of directories does not contain a semicolon, it is interpreted as *dirlist2*.

LD_LIBRARY_PATH is also used to specify library search directories to the runtime linker. That is, if **LD_LIBRARY_PATH** exists in the environment, the runtime linker will search the directories named in it, before its default directory, for shared objects to be linked with the program at execution.

Note: When running a set-user-ID or set-group-ID program, the runtime linker will only search for libraries in **/usr/lib** and any full pathname specified within the executable as a result of a **runpath** being specified when the executable was constructed. Any library dependencies specified as relative pathnames will be silently ignored.

LD_OPTIONS

A default set of options to **ld**. **LD_OPTIONS** is interpreted by **ld** just as though its value had been placed on the command line, immediately following the name used to invoke **ld**, as in:

ld \$LD_OPTIONS ... other-arguments ...

LD_PRELOAD

A list of shared objects that are to be interpreted by the runtime linker. The specified shared objects are linked in after the program being executed and before any other shared objects that the program references.

Note: When running a set-user-ID or set-group-ID program, this option is silently ignored.

LD_RUN_PATH

An alternative mechanism for specifying a runpath to the link editor (see **-R** option). If both **LD_RUN_PATH** and the **-R** option are specified, **-R** supersedes.

LD_DEBUG

Provide a list of tokens that will cause the runtime linker to print debugging information to the standard error. The special token *help* indicates the full list of tokens available.

Note: Environment variable-names beginning with the characters 'LD_' are reserved for possible future enhancements to **ld**.

NOTES**Options No Longer Supported**

The following SunOS 4.x.y options do not have any replacement in this release:

-Bnosymbolic (this is now the default if **-Bsymbolic** is not used), **-d**, **-dc**, and **-dp**, (these are now the default, see **-b** above to override the default), **-M**, **-S**, **-t**, **-x**, **-X**, and **-ysym**.

The following SunOS 4.x.y options are not supported: **-align datum**, **-A name**, **-D**, **-p**, **-T[text] hex**, **-T datahex**. Much of the functionality of these options can be achieved using the **-Mmapfile** option.

Obsolete Options

The following SunOS 4.x.y options are obsolete in this release: **-n**, **-N** and **-z**.

FILES

libx.so	libraries
libx.a	libraries
a.out	output file
LIBPATH	usually /usr/ccs/lib:/usr/lib

SEE ALSO

as(1), **cc(1B)**, **ld(1B)**, **exec(2)**, **exit(2)**, **elf(3E)**, **end(3C)**, **exit(3C)**, **a.out(4)**, **ar(4)**

Linker and Libraries Guide

Solaris Binary Compatibility Guide

NAME	ld – link editor, dynamic link editor
SYNOPSIS	<i>/usr/ucb/ld</i> [<i>options</i>]
AVAILABILITY	SUNWscpu
DESCRIPTION	<i>/usr/ucb/ld</i> is the link editor for the BSD Compatibility Package. <i>/usr/ucb/ld</i> is identical to <i>/usr/bin/ld</i> (see <i>ld(1)</i>) except that BSD libraries and routines are included <i>before</i> the base libraries and routines.
OPTIONS	<i>/usr/ucb/ld</i> accepts the same options as <i>/usr/bin/ld</i> , with the following exceptions: -Ldir Add <i>dir</i> to the list of directories searched for libraries by <i>/usr/bin/ld</i> . Directories specified with this option are searched before <i>/usr/ucb/lib</i> and <i>/usr/lib</i> . -Y LU,dir Change the default directory used for finding libraries. Warning: This option may have unexpected results, and should not be used.
FILES	<i>/usr/lib</i> <i>/usr/lib/libx.a</i> <i>/usr/ucb/lib</i> <i>/usr/ucb/lib/libx.a</i>
SEE ALSO	<i>ar(1)</i> , <i>as(1)</i> , <i>cc(1B)</i> , <i>ld(1)</i> , <i>lorder(1)</i> , <i>strip(1)</i> , <i>tsort(1)</i>

NAME	ldd – list dynamic dependencies of executable files or shared objects
SYNOPSIS	ldd [-d -r] [-s] <i>filename</i> . . .
DESCRIPTION	<p>ldd lists the dynamic dependencies of executable files or shared objects. If <i>filename</i> is an executable file, ldd lists the pathnames of all shared objects that would be loaded as a result of executing <i>filename</i>.</p> <p>If <i>filename</i> is a shared object, ldd lists the pathnames of all shared objects that would be loaded as a result of loading <i>filename</i>. ldd requires shared objects to have execute permission.</p> <p>ldd processes its input one file at a time. For each input file ldd does one of the following:</p> <ul style="list-style-type: none"> Lists the object dependencies if they exist. Succeeds quietly if dependencies do not exist. Prints an error message if processing fails.
OPTIONS	<p>ldd can also check the compatibility of <i>filename</i> with the shared objects it uses. The following options indicate to ldd to print warnings for any unresolved symbol references that would occur if <i>filename</i> were executed.</p> <ul style="list-style-type: none"> -d Check references to data objects. -r Check references to both data objects and functions. <p>Only one of the above options may be given during any single invocation of ldd.</p> <ul style="list-style-type: none"> -s Displays the search path used to locate shared object dependencies.
FILES	<p>/usr/lib/lddstub Fake executable loaded to check the dependencies of shared objects.</p>
SEE ALSO	<p>ld(1), dlopen(3X) <i>Linker and Libraries Guide</i></p>
DIAGNOSTICS	<p>ldd prints the record of shared object path names to stdout. The optional list of symbol resolution problems are printed to stderr. If <i>filename</i> is not an executable file or a shared object, or if it cannot be opened for reading, a non-zero exit status is returned.</p>
NOTES	<p>ldd does not list shared objects explicitly attached using dlopen(3X).</p> <p>Using the -d or -r option with shared objects can give misleading results. ldd does a “worst case” analysis of the shared objects. However, in practice some or all of the symbols reported as unresolved can be resolved by the executable file referencing the shared object.</p> <p>ldd uses the same algorithm as the dynamic linker to locate shared objects.</p>

NAME	let – shell built-in function to evaluate one or more arithmetic expressions
SYNOPSIS ksh	let <i>arg</i> ...
DESCRIPTION ksh	Each <i>arg</i> is a separate "arithmetic expression" to be evaluated. The exit status is 0 if the value of the last expression is non-zero, and 1 otherwise.
SEE ALSO	ksh(1), set(1), typeset(1)

NAME	lex – lexical analysis program generator
SYNOPSIS	lex [-cntv] [-e -w] [-V -Q [y n]] [<i>filename</i>] ...
DESCRIPTION	<p>lex generates programs to be used in simple lexical analysis of text. Each <i>filename</i> (the standard input by default) contains regular expressions to search for, and actions written in C to be executed when expressions are found.</p> <p>A C source program, lex.yy.c is generated, to be compiled as follows:</p> <pre style="margin-left: 40px;">cc lex.yy.c -ll</pre> <p>This program, when run, copies unrecognized portions of the input to the output, and executes the associated C action for each regular expression that is recognized. The actual string matched is left in yytext, an external character array (a wchar_t array when the -w option is given).</p> <p>Matching is done in order of the strings in the file. The strings may contain square braces to indicate character classes, as in [abx-z] to indicate a, b, x, y, and z; and the operators *, + and ?, which mean, respectively, any nonnegative number, any positive number, or either zero or one occurrences of the previous character or character-class. The “dot” character (.) is the class of all characters except NEWLINE.</p> <p>Parentheses for grouping and vertical bar for alternation are also supported. The notation r{d, e} in a rule indicates instances of regular expression r between d and e. It has a higher precedence than , but lower than that of *, ?, +, or concatenation. The ^ (carat character) at the beginning of an expression permits a successful match only immediately after a NEWLINE, and the \$ character at the end of an expression requires a trailing NEWLINE.</p> <p>The / character in an expression indicates trailing context; only the part of the expression up to the slash is returned in yytext, although the remainder of the expression must follow in the input stream.</p> <p>An operator character may be used as an ordinary symbol if it is within ” symbols or is preceded by \`.</p> <p>Three subroutines defined as macros are expected: input() to read a character; unput(c) to replace a character read; and output(c) to place an output character. They are defined in terms of the standard streams, but you can override them. For C++ code, input() is renamed lex_input(), and output() is renamed lex_output() to avoid name conflicts with iostreams. The program generated is named yylex(), and the lex library libl.a contains a main() which calls it. The action REJECT on the right side of the rule rejects this match and executes the next suitable match; the function yymore() accumulates additional characters into the same yytext; and the function yyles(p) pushes back the portion of the string matched beginning at p, which should be between yytext and yytext+yyleng. The macros input and output use files yyin and yyout to read from and write to, defaulted to stdin and stdout, respectively.</p>

In a **lex** program, any line beginning with a blank is assumed to contain only C text and is copied; if it precedes **%%** it is copied into the external definition area of the **lex.yy.c** file. All rules should follow a **%%**, as in YACC. Lines preceding **%%** which begin with a non-blank character define the string on the left to be the remainder of the line; it can be used later by surrounding it with **{}**. Note: curly brackets do not imply parentheses; only string substitution is done.

The external names generated by **lex** all begin with the prefix **yy** or **YY**.

Certain table sizes for the resulting finite-state machine can be set in the definitions section:

```
%p n   number of positions is n (default 2000)
%n n   number of states is n (default 500)
%e n   number of parse tree nodes is n (default 1000)
%a n   number of transitions is n (default 3000)
```

The use of one or more of the above automatically implies the **-v** option, unless the **-n** option is used.

Programs generated by **lex(1)** need either the **-e** or **-w** option to handle input that contains EUC characters from supplementary codesets. If neither of these options is specified, **yytext** is of the type **char[]**, and the generated program can handle only ASCII characters.

When the **-e** option is used, **yytext** is of the type **unsigned char[]** and **yylen** gives the total number of *bytes* in the matched string. With this option, the macros **input()**, **unput(c)**, and **output(c)** should do a byte-based I/O in the same way as with the regular ASCII **lex(1)**. Two more variables are available with the **-e** option, **yywtext** and **yywleng**, which behave the same as **yytext** and **yylen** would under the **-w** option.

When the **-w** option is used, **yytext** is of the type **wchar_t[]** and **yylen** gives the total number of *characters* in the matched string. If you supply your own **input()**, **unput(c)**, or **output(c)** macros with this option, they must return or accept EUC characters in the form of wide character (**wchar_t**). This allows a different interface between your program and the **lex** internals, to expedite some programs.

When either the **-e** or **-w** option is used, the generated C program must be linked with the wide character library **libw.a** using the **-lw** linker flag.

Pattern Matching

When either the **-e** or **-w** option is used, patterns used in rules can include characters from both primary and supplementary codesets. The generated program performs pattern matching correctly on an input stream containing EUC characters from supplementary codesets.

You may use any valid EUC characters in a character range **[A-Z]** as long as **A** and **Z** belong to the same codeset.

." matches any character from any codeset (except **NEWLINE**).

International Caveats

Start condition names must consist solely of ASCII characters.

The "%T" directive can not be used when either the `-w` or `-e` option is used.

The default `main()` found in the lex library (`libl.a`) does not have a `setlocale(3C)` call. Thus, the resulting program would not recognize non-ASCII characters correctly. You have to supply your own `main()` in order to have your program handle EUC characters correctly. The simplest `main()` would be:

```
#include <locale.h>
main(){
    setlocale(LC_ALL, "");
    yylex();
}
```

OPTIONS

- `-c` Indicates C actions and is the default.
- `-e` Generate a program that can handle EUC characters (cannot be used with the `-w` option).
`yytext[]` is of type `unsignedchar[]`.
- `-n` Opposite of `-v`; `-n` is the default.
- `-t` Place the result on the standard output instead of in file `lex.yy.c`.
- `-v` Print a one-line summary of statistics of the generated analyzer.
- `-w` Generate a program that can handle EUC characters (cannot be used with the `-e` option).
Unlike the `-e` option, `yytext[]` is of type `wchar_t[]`.
- `-V` Print out version information on standard error.
- `-Q[y|n]`
Print out version information to output file `lex.yy.c` by using `-Qy`. The `-Qn` option does not print out version information and is the default.

EXAMPLES

The command line,

```
lex lexcommands
```

draws `lex` instructions from the file `lexcommands`, and places the output in `lex.yy.c`.

The following example `lex` program converts uppercase to lower, removes blanks at the end of lines, and replaces multiple blanks by single blanks.

```
%%
[A-Z] putchar (yytext[0]+'a'-'A');
[ ]+$ ;
[ ]+ putchar( ' ');
```

INTERNATIONAL EXAMPLES

The following is a similar program for the "japanese" locale environment:

```
%%
[\x30001221-\x30001273] putwchar (yytext[0]+0x0080);
```

```

[ \x300010a1]+$      ;
[ \x300010a1]+      putchar( ' ');
%%
#include <locale.h>
main(){
    setlocale(LC_ALL, "");
    yylex();
}

```

This program converts every hiragana character (of which the EUC wide character value is between 0x30001221 and 0x30001273) to the corresponding katakana character. It also recognizes double-space character (0x300010a1). 0x0080 is the offset between corresponding hiragana and katakana characters when represented in wide characters. Note that use of the hexadecimal escape sequence in this example is not really needed. The corresponding EUC characters could have been used instead.

This program must be compiled with the `-lw` option and linked with the wide character library `libw.a`. Compilation and execution must be done in an environment where either the `LANG` or `LC_CTYPE` environment variable is set to `japanese`. The command line for compiling this program would be:

```

% lex -w sample.l
% cc -o sample lex.yy.c -ll -lw

```

FILES

<code>lex.yy.c</code>	default output file when <code>-t</code> is not specified
<code>/usr/ccs/lib/libl.a</code>	lex library
<code>ncform</code>	
<code>nceucform</code>	C-program prototypes

SEE ALSO `sed(1)`, `yacc(1)`, `setlocale(3C)`
`lex` in *Programming Utilities Guide*

NOTES Ratfor is no longer supported as a host language.

The way to use hexadecimal escape sequences for multibyte characters differs from the versions of `lex` of previous release of SunOS Asian Language Environment, namely JLE, KLE, CLE and HLE. In these versions, a multibyte character was written as a sequence of hexadecimal escape sequences, one per byte, rather than as one hexadecimal escape sequence representing the character's wide character value.

NAME	limit, ulimit, unlimit – shell built-in functions to set/get limitations on the system resources available to the current shell and it's descendents.				
SYNOPSIS					
sh	ulimit [-[HS] [a cdfnstv]] ulimit [-[HS]] c d f n s t v] <i>limit</i>				
csh	limit [-h] [<i>resource</i> [<i>max-use</i>]] unlimit [-h] [<i>resource</i>]				
ksh	ulimit [-HSacdfmnpstv] [<i>limit</i>]				
DESCRIPTION					
sh	ulimit prints or sets hard or soft resource limits. These limits are described in getrlimit(2) . If <i>limit</i> is not present, ulimit prints the specified limits. Any number of limits may be printed at one time. The -a option prints all limits. If <i>limit</i> is present, ulimit sets the specified limit to <i>limit</i> . The string unlimited requests the largest valid limit. Limits may be set for only one resource at a time. Any user may set a soft limit to any value below the hard limit. Any user may lower a hard limit. Only a super-user may raise a hard limit; see su(1M) . The -H option specifies a hard limit. The -S option specifies a soft limit. If neither option is specified, ulimit will set both limits and print the soft limit. The following options specify the resource whose limits are to be printed or set. If no option is specified, the file size limit is printed or set. -c maximum core file size (in 512-byte blocks) -d maximum size of data segment or heap (in kbytes) -f maximum file size (in 512-byte blocks) -n maximum file descriptor plus 1 -s maximum size of stack segment (in kbytes) -t maximum CPU time (in seconds) -v maximum size of virtual memory (in kbytes)				
csh	limit limits the consumption by the current process or any process it spawns, each not to exceed <i>max-use</i> on the specified <i>resource</i> . If <i>max-use</i> is omitted, print the current limit; if <i>resource</i> is omitted, display all limits. -h Use hard limits instead of the current limits. Hard limits impose a ceiling on the values of the current limits. Only the privileged user may raise the hard limits. <i>resource</i> is one of: <table border="0" style="margin-left: 2em;"> <tr> <td style="padding-right: 1em;">cputime</td> <td>Maximum CPU seconds per process.</td> </tr> <tr> <td>filesize</td> <td>Largest single file allowed.</td> </tr> </table>	cputime	Maximum CPU seconds per process.	filesize	Largest single file allowed.
cputime	Maximum CPU seconds per process.				
filesize	Largest single file allowed.				

datasize	Maximum data size (including stack) for the process.
stacksize	Maximum stack size for the process.
coredumpsize	Maximum size of a core dump (file).
descriptors	Maximum number of file descriptors.
memorysize	Maximum size of virtual memory.

max-use is a number, with an optional scaling factor, as follows:

nh	Hours (for cputime).
nk	<i>n</i> kilobytes. This is the default for all but cputime .
nm	<i>n</i> megabytes or minutes (for cputime).
mm:ss	Minutes and seconds (for cputime).

unlimit removes a limitation on *resource*. If no *resource* is specified, then all resource limitations are removed. See the description of the **limit** command for the list of resource names.

-h Remove corresponding hard limits. Only the privileged user may do this.

ksh

ulimit sets or displays a resource limit. The available resources limits are listed below. Many systems do not contain one or more of these limits. The limit for a specified resource is set when *limit* is specified. The value of *limit* can be a number in the unit specified below with each resource, or the value **unlimited**. The **H** and **S** flags specify whether the hard limit or the soft limit for the given resource is set. A hard limit cannot be increased once it is set. A soft limit can be increased up to the value of the hard limit. If neither the **H** or **S** options is specified, the limit applies to both. The current resource limit is printed when *limit* is omitted. In this case the soft limit is printed unless **H** is specified. When more than one resource is specified, then the limit name and unit is printed before the value.

- a** Lists all of the current resource limits.
- c** The number of 512-byte blocks on the size of core dumps.
- d** The number of K-bytes on the size of the data area.
- f** The number of 512-byte blocks on files written by child processes (files of any size may be read).
- m** The number of K-bytes on the size of physical memory.
- n** The number of file descriptors plus 1.
- p** The number of 512-byte blocks for pipe buffering.
- s** The number of K-bytes on the size of the stack area.
- t** The number of seconds to be used by each process.
- v** The number of K-bytes for virtual memory.

If no option is given, **-f** is assumed.

SEE ALSO

csh(1), **ksh(1)**, **sh(1)**, **su(1)**, **getrlimit(2)**

NAME	line – read one line
SYNOPSIS	line
AVAILABILITY	SUNWcsu
DESCRIPTION	line copies one line (up to a new-line) from the standard input and writes it on the standard output. It returns an exit code of 1 on EOF and always prints at least a new-line. It is often used within shell files to read from the user's terminal.
SEE ALSO	sh(1) , read(2)

NAME	listusers – list user login information
SYNOPSIS	listusers [-g <i>groups</i>] [-l <i>logins</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	Executed without any options, this command lists all user logins sorted by login. The output shows the login ID and the account field value from the system's password database as specified by /etc/nsswitch.conf .
OPTIONS	-g <i>groups</i> Lists all user logins belonging to group , sorted by login. Multiple groups can be specified as a comma-separated list. -l <i>logins</i> Lists the user login or logins specified by logins , sorted by login. Multiple logins can be specified as a comma-separated list.
SEE ALSO	nsswitch.conf(4)
NOTES	A user login is one that has a UID of 100 or greater. The -l and -g options can be combined. User logins will only be listed once, even if they belong to more than one of the selected groups.

NAME	ln – make hard or symbolic links to files
SYNOPSIS	ln [-f] [-n] [-s] <i>filename1</i> [<i>filename2</i> ... <i>filenamen</i>] <i>target</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The ln command links <i>filenamen</i> to <i>target</i> by creating a directory entry that refers to <i>target</i>. By using ln with one or more file names, the user may create one or more links to <i>target</i>.</p> <p>The ln command may be used to create both hard links and symbolic links; by default it creates hard links. A hard link to a file is indistinguishable from the original directory entry. Any changes to a file are effective independent of the name used to reference the file. Hard links may not span file systems and may not refer to directories.</p> <p>Without the -s option, ln is used to create hard links. <i>filenamen</i> is linked to <i>target</i>. If <i>target</i> is a directory, another file named <i>filenamen</i> is created in <i>target</i> and linked to the original <i>filenamen</i>. If <i>target</i> is a file, its contents are overwritten.</p> <p>If ln determines that the mode of <i>target</i> forbids writing, it will print the mode (see chmod(2)), ask for a response, and read the standard input for one line. If the line begins with y, the link occurs, if permissible; otherwise, the command exits.</p>
OPTIONS	<p>-f ln will link files without questioning the user, even if the mode of <i>target</i> forbids writing. Note: This is the default if the standard input is not a terminal.</p> <p>-n If the linkname is an existing file, do not overwrite the contents of the file. The -f option overrides this option.</p> <p>-s ln will create a symbolic link. A symbolic link contains the name of the file to which it is linked. Symbolic links may span file systems and may refer to directories.</p> <p>If the -s option is used with two arguments, <i>target</i> may be an existing directory or a non-existent file. If <i>target</i> already exists and is not a directory, an error is returned. <i>filenamen</i> may be any path name and need not exist. If it exists, it may be a file or directory and may reside on a different file system from <i>target</i>. If <i>target</i> is an existing directory, a file is created in directory <i>target</i> whose name is <i>filenamen</i> or the last component of <i>filenamen</i>. This file is a symbolic link that references <i>filenamen</i>. If <i>target</i> does not exist, a file with name <i>target</i> is created and it is a symbolic link that references <i>filenamen</i>.</p> <p>If the -s option is used with more than two arguments, <i>target</i> must be an existing directory or an error will be returned. For each <i>filenamen</i>, a file is created in <i>target</i> whose name is <i>filenamen</i> or its last component; each new <i>filenamen</i> is a symbolic link to the original <i>filenamen</i>. The <i>files</i> and <i>target</i> may reside on different file systems.</p> <p>File permissions for <i>target</i> may be different than those displayed with a -l listing of the ls(1) command. To display the permissions of <i>target</i> use ls -lL. See stat(2) for more information.</p>

SEE ALSO

chmod(1), cp(1), ls(1), mv(1), rm(1), link(1M), readlink(2), stat(2), symlink(2)

NAME	ln – make hard or symbolic links to files
SYNOPSIS	<code>/usr/ucb/ln [-fs] filename [linkname]</code> <code>/usr/ucb/ln [-fs] pathname... directory</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p><code>/usr/ucb/ln</code> creates an additional directory entry, called a link, to a file or directory. Any number of links can be assigned to a file. The number of links does not affect other file attributes such as size, protections, data, etc.</p> <p><i>filename</i> is the name of the original file or directory. <i>linkname</i> is the new name to associate with the file or filename. If <i>linkname</i> is omitted, the last component of <i>filename</i> is used as the name of the link.</p> <p>If the last argument is the name of a directory, symbolic links are made in that directory for each <i>pathname</i> argument; <code>/usr/ucb/ln</code> uses the last component of each <i>pathname</i> as the name of each link in the named <i>directory</i>.</p> <p>A hard link (the default) is a standard directory entry just like the one made when the file was created. Hard links can only be made to existing files. Hard links cannot be made across file systems (disk partitions, mounted file systems). To remove a file, all hard links to it must be removed, including the name by which it was first created; removing the last hard link releases the inode associated with the file.</p> <p>A symbolic link, made with the <code>-s</code> option, is a special directory entry that points to another named file. Symbolic links can span file systems and point to directories. In fact, you can create a symbolic link that points to a file that is currently absent from the file system; removing the file that it points to does not affect or alter the symbolic link itself.</p> <p>A symbolic link to a directory behaves differently than you might expect in certain cases. While an <code>ls(1)</code> on such a link displays the files in the pointed-to directory, an <code>'ls -l'</code> displays information about the link itself:</p> <pre> example% /usr/ucb/ln -s dir link example% ls link file1 file2 file3 file4 example% ls -l link lrwxrwxrwx 1 user 7 Jan 11 23:27 link -> dir </pre> <p>When you <code>cd(1)</code> to a directory through a symbolic link, you wind up in the pointed-to location within the file system. This means that the parent of the new working directory is not the parent of the symbolic link, but rather, the parent of the pointed-to directory. For instance, in the following case the final working directory is <code>/usr</code> and not <code>/home/user/linktest</code>.</p> <pre> example% pwd /home/user/linktest example% /usr/ucb/ln -s /var/tmp symlink example% cd symlink </pre>

```
example% cd ..
example% pwd
/usr
```

C shell user's can avoid any resulting navigation problems by using the **pushd** and **popd** built-in commands instead of **cd**.

OPTIONS

- f Force a hard link to a directory. This option is only available to the super-user, and should be used with extreme caution.
- s Create a symbolic link or links.

EXAMPLES

The commands below illustrate the effects of the different forms of the **/usr/ucb/ln** command:

```
example% /usr/ucb/ln file link
example% ls -F file link
file link
example% /usr/ucb/ln -s file symlink
example% ls -F file symlink
file symlink@
example% ls -li file link symlink
10606 -rw-r--r-- 2 user      0 Jan 12 00:06 file
10606 -rw-r--r-- 2 user      0 Jan 12 00:06 link
10607 lrwxrwxrwx 1 user      4 Jan 12 00:06 symlink -> file
example% /usr/ucb/ln -s nonesuch devoid
example% ls -F devoid
devoid@
example% cat devoid
devoid: No such file or directory
example% /usr/ucb/ln -s /proto/bin/* /tmp/bin
example% ls -F /proto/bin /tmp/bin
/proto/bin:
x*  y*  z*

/tmp/bin:
x@  y@  z@
```

SEE ALSO

cp(1), **ls(1)**, **mv(1)**, **rm(1)**, **link(2)**, **readlink(2)**, **stat(2)**, **symlink(2)**

NOTES

When the last argument is a directory, simple basenames should not be used for *pathname* arguments. If a basename is used, the resulting symbolic link points to itself:

```
example% /usr/ucb/ln -s file /tmp
example% ls -l /tmp/file
lrwxrwxrwx 1 user      4 Jan 12 00:16 /tmp/file -> file
example% cat /tmp/file
/tmp/file: Too many levels of symbolic links
```

To avoid this problem, use full pathnames, or prepend a reference to the **PWD** variable to files in the working directory:

```
example% rm /tmp/file
```

```
example% /usr/ucb/ln -s $PWD/file /tmp
```

```
lrwxrwxrwx 1 user 4   Jan 12 00:16 /tmp/file → /home/user/subdir/file
```

NAME	loadfont – display or change font information in the RAM of the video card on an x86 system in text mode
SYNOPSIS	loadfont [<i>-f BDF_file</i> <i>-c codeset</i>] [<i>-m mode</i>] [<i>-d</i>]
AVAILABILITY	x86 SUNWcsu
DESCRIPTION	<p>The loadfont utility allows a user to load and activate a different font into the RAM of the video card used by the console of the Solaris for x86 operating system in text mode. It can also be used to display information about the fonts currently in use. In addition, the <i>-m</i> option can be used to change the size of the characters on the screen; it can also be used to change the number of lines per screen. loadfont will always read from standard output; this will allow a system administrator to use it from a remote terminal.</p> <p>When used without arguments, loadfont displays the different ways the command can be used, as shown in the synopsis.</p>
Options	<p><i>-f BDF_file</i> This command reads the contents of <i>BDF_file</i> and subsequently loads the font specified in the file into the RAM of the video card. The file must be in the Binary Distribution Format version 2.1 as developed by Adobe Systems, Inc. (See loadfont(4).)</p> <p><i>-c codeset</i> <i>codeset</i> is the name of a codeset available for the current font size. This font will be loaded into the RAM of the video card and activated. Use ? to find out the valid <i>codesets</i> available. This option is a shorthand form of <i>-f</i>.</p> <p><i>-m mode</i> This option will attempt to change the mode of the console as specified. This will result in having a different font size and/or different number of lines and columns on the screen. Use ? to find out the valid <i>modes</i> available.</p> <p><i>-d</i> This <i>reads</i> the font information from the video RAM and <i>writes</i> it to standard output in a format compatible with the Binary Distribution Format version 2.1 as developed by Adobe Systems, Inc. (See loadfont(4).)</p>
Fonts	<p>A font is the representation of characters by images. The need to use different fonts can be imposed by:</p> <ol style="list-style-type: none"> 1. The codeset used to represent the characters internally. 2. The resolution used to display the characters. <p>Each font contains exactly 256 images. All supported fonts are fixed size (constant width and constant height), i.e., each character takes the same amount of space on the screen. When the monitor is not being used in graphics mode, the loadfont utility allows a user to modify the font used by the video card, so different images are displayed on the screen</p>

of the console for the various characters. The same video card may support different text modes. Video cards typically differ by the number of pixels they use to represent a single character. On any given video card, the same number of pixels is used for each character. For the standard VGA video cards, 8 by 16 (8 horizontally and 16 vertically) resolution is supported:

When **loadfont** is invoked to modify the existing font, it will attempt to do so for the font size currently in use. Use the **-m** option to switch to another font size.

**loadfont and
pcmapkeys**

There is an almost one-to-one relationship between the use of the **loadfont** utility and the **pcmapkeys** utility. Whereas **loadfont** is used to list or modify the images that correspond with the various characters, the **pcmapkeys** utility is used to determine how characters are generated from the keyboard and which code (a single byte code) will be used to represent the character internally. The default representation is the ISO 8859-1 codeset.

When a different codeset is used, both a different **pcmapkeys** input file and a different font set are required. If the default font does not satisfy your needs (because a different font size or a customized font is required, e.g., a Greek font), a **loadfont** description file to be used with the **-f** option is needed. A sample file that describes the IBM extended ASCII font for an 8 by 16 resolution is supplied (**437.bdf**). A second sample file, **646g.bdf**, contains a font file for German ASCII. See **pcmapkeys(1)** and **loadfont(4)** for additional details.

WARNING

When an attempt is made to switch to a mode that the video card does not support, you will get a blank screen. There is nothing wrong with the system; as super-user, simply type in the command to set the mode back, e.g.:

```
loadfont -m V80x25
```

FILES

/usr/share/lib/fonts/8859-1.bdf	the Binary Distribution Format (BDF) file for the default fonts
/usr/share/lib/fonts/437.bdf	sample Binary Distribution Format (BDF) file for IBM 437 font on a VGA
/usr/share/lib/fonts/646g.bdf	sample BDF file for German ASCII

SEE ALSO

pcmapkeys(1), **loadfont(4)**

NOTE

The default fonts on the system are those of the ISO 8859-1 codeset. The optional IBM DOS 437 codeset is supported *only* at internationalization level 1. That is, if you choose to download fonts of the optional IBM DOS 437 codeset, there will be no support for non-standard U.S. date, time, currency, numbers, unit, and collation. There will be no support for non-English message and text presentation, and no multi-byte character support. Therefore, non-Windows users should only use IBM DOS 437 codeset in the default C locale.

NAME	loadkeys, dumpkeys – load and dump keyboard translation tables
SYNOPSIS	loadkeys [-e] [<i>filename</i>] dumpkeys
AVAILABILITY	SPARC SUNWloc
DESCRIPTION	<p>loadkeys reads the file specified by <i>filename</i>, and modifies the keyboard streams module's translation tables. If no file is specified, and the keyboard is a Type-4 keyboard, a default file for the layout indicated by the DIP switches on the keyboard. The file is in the format specified by keytables(4).</p> <p>If the layout code in the DIP switches on the keyboard has the hexadecimal value 0xdd, the file loaded by loadkeys by default is /usr/share/lib/keytables/layout_dd. These files specify only the entries that change between the different Type-4 keyboard layouts.</p> <p>dumpkeys writes, to the standard output, the current contents of the keyboard streams module's translation tables, in the format specified by keytables(4).</p>
OPTIONS	-e loadkeys loads the requested keytable layout information into the EEPROM keytables. This allows the console monitor to select the correct key-values when a non-US keyboard is connected to the workstation.
FILES	/usr/share/lib/keytables/layout_dd default keytable files
SEE ALSO	kbd(1) , keytables(4) , kb(7)
BUGS	The -e option requires EEPROM rev 2.6 or later.

NAME	logger – add entries to the system log
SYNOPSIS	<code>/usr/ucb/logger [-i] [-f <i>filename</i>] [-p <i>priority</i>] [-t <i>tag</i>] [<i>message</i>] ...</code>
DESCRIPTION	logger provides a method for adding one-line entries to the system log file from the command line. One or more <i>message</i> arguments can be given on the command line, in which case each is logged immediately. If this is unspecified, either the file indicated with <code>-f</code> or the standard input is added to the log. Otherwise, a <i>filename</i> can be specified, in which case each line in the file is logged. If neither is specified, logger reads and logs messages on a line-by-line basis from the standard input.
OPTIONS	<p><code>-i</code> Log the process ID of the logger process with each line.</p> <p><code>-f <i>filename</i></code> Use the contents of <i>filename</i> as the message to log.</p> <p><code>-p <i>priority</i></code> Enter the message with the specified <i>priority</i>. The message priority can be specified numerically, or as a <i>facility.level</i> pair. For example, '<code>-p local3.info</code>' assigns the message priority to the info level in the local3 facility. The default priority is user.notice.</p> <p><code>-t <i>tag</i></code> Mark each line added to the log with the specified <i>tag</i>.</p>
EXAMPLES	<p>The following example:</p> <p style="padding-left: 40px;">example% logger System rebooted</p> <p>logs the message 'System rebooted' to the default priority level notice to be treated by syslogd as are other messages to the facility user.</p> <p>The next example:</p> <p style="padding-left: 40px;">logger -p local0.notice -t HOSTIDM -f /dev/idmc</p> <p>reads from the file <code>/dev/idmc</code> and logs each line in that file as a message with the tag 'HOSTIDM' at priority level notice to be treated by syslogd as are other messages to the facility local0.</p>
SEE ALSO	syslogd(1M) , syslog(3)

NAME	logger – add entries to the system log
SYNOPSIS	/usr/ucb/logger [-f <i>filename</i>] [-i] [-p <i>priority</i>] [-t <i>tag</i>] [<i>message</i>] ...
AVAILABILITY	SUNWscpu
DESCRIPTION	logger provides a method for adding one-line entries to the system log file from the command line. One or more <i>message</i> arguments can be given on the command line, in which case each is logged immediately. If <i>message</i> is unspecified, either the file indicated with -f or the standard input is added to the log. Otherwise, a <i>filename</i> can be specified, in which case each line in the file is logged. If neither is specified, logger reads and logs messages on a line-by-line basis from the standard input.
OPTIONS	<p>-i Log the process ID of the logger process with each line.</p> <p>-f <i>filename</i> Use the contents of <i>filename</i> as the message to log.</p> <p>-p <i>priority</i> Enter the message with the specified <i>priority</i>. The message priority can be specified numerically, or as a <i>facility.level</i> pair. For example, '-p local3.info' assigns the message priority to the info level in the local3 facility. The default priority is user.notice.</p> <p>-t <i>tag</i> Mark each line added to the log with the specified <i>tag</i>.</p>
EXAMPLES	<p>The command:</p> <p style="padding-left: 40px;">example% logger System rebooted</p> <p>will log the message 'System rebooted' to the facility at priority notice to be treated by syslogd as other messages to the facility notice are.</p> <p>The next command:</p> <p style="padding-left: 40px;">example% logger -p local0.notice -t HOSTIDM -f /dev/idmc</p> <p>will read from the file /dev/idmc and will log each line in that file as a message with the tag 'HOSTIDM' at priority notice to be treated by syslogd as other messages to the facility local0 are.</p>
SEE ALSO	syslogd(1M) , syslog(3)

NAME	login – sign on to the system
SYNOPSIS	login [-d <i>device</i>] [-h <i>hostname</i> [<i>terminal</i>] -r <i>hostname</i>] [<i>name</i> [<i>environ</i> ...]]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The login command is used at the beginning of each terminal session and allows you to identify yourself to the system. It is invoked by the system when a connection is first established and after the previous user has terminated the login shell by issuing the exit command.</p> <p>If login is invoked as a command it must replace the initial command interpreter. This is accomplished by typing:</p> <p style="text-align: center;">exec login</p> <p>from the initial shell.</p> <p>login asks for your user name (if it is not supplied as an argument), and if appropriate, your password. Echoing is turned off (where possible) during the typing of your password, so it will not appear on the written record of the session.</p> <p>If there are no lower-case characters in the first line of input processed, login assumes the connecting TTY is an upper-case-only terminal and sets the port's termio(7) options to reflect this.</p> <p>If you make any mistake in the login procedure, the message:</p> <p style="text-align: center;">Login incorrect</p> <p>is printed and a new login prompt will appear. If you make five incorrect login attempts, all five may be logged in /var/adm/loginlog (if it exists) and the TTY line will be dropped. If you do not complete the login successfully within a certain period of time, it is likely that you will be silently disconnected.</p> <p>After a successful login, accounting files are updated; device owner, group, and permissions are set according to the contents of the /etc/logindevperm file (see login-devperm(4)); and the time you last logged in is printed.</p> <p>The user-ID, group-ID, supplementary group list, and working directory are initialized, and the command interpreter (usually ksh) is started.</p> <p>The basic <i>environment</i> is initialized to:</p> <p style="text-align: center;">HOME=<i>your-login-directory</i> LOGNAME=<i>your-login-name</i> PATH=/usr/bin:</p>

SHELL=*last-field-of-passwd-entry*
MAIL=*/var/mail/your-login-name*
TZ=*timezone-specification*

For Bourne shell and Korn shell logins, the shell executes **/etc/profile** and **\$HOME/.profile** (if it exists). For C shell logins, the shell executes **/etc.login**, **\$HOME/.cshrc**, and **\$HOME/.login**. The default **/etc/profile** and **/etc.login** files check quotas (see **quota(1M)**), print **/etc/motd** and check for mail. None of the messages are printed if the file **\$HOME/.hushlogin** exists. The name of the command interpreter is set to **-** (dash), followed by the last component of the interpreter's path name (for example, **-sh**).

If the *login-shell* field in the password file (see **passwd(4)**) is empty, then the default command interpreter, **/usr/bin/sh**, is used. If this field is ***** (asterisk), then the named directory becomes the root directory. At that point **login** is re-executed at the new level which must have its own root structure.

The environment may be expanded or modified by supplying additional arguments to **login**, either at execution time or when **login** requests your login name. The arguments may take either the form **xxx** or **xxx=yyy**. Arguments without an equal sign are placed in the environment as:

Ln=xxx

where *n* is a number starting at 0 and is incremented each time a new variable name is required. Variables containing an **=** are placed in the environment without modification. If they already appear in the environment, then they replace the older value. There are two exceptions. The variables **PATH** and **SHELL** cannot be changed. This prevents people, logging into restricted shell environments, from spawning secondary shells which are not restricted. **login** understands simple single-character quoting conventions. Typing a **' \ '** (backslash) in front of a character quotes it and allows the inclusion of such characters as spaces and tabs.

To enable remote logins by root, edit the **/etc/default/login** file by inserting a **' # '** (pound-sign) before the **CONSOLE=/dev/console** entry. See **FILES** below.

OPTIONS

-d device **login** accepts a device option, *device*. *device* is taken to be the path name of the TTY port **login** is to operate on. The use of the device option can be expected to improve **login** performance, since **login** will not need to call **ttyname(3C)**. The **-d** option is available only to users whose **UID** and effective **UID** are root. Any other attempt to use **-d** will cause **login** to quietly exit.

-h hostname [terminal]

used by **in.telnetd(1M)** to pass information about the remote host and terminal type.

-r hostname

used by **in.rlogind(1M)** to pass information about the remote host.

FILES

\$HOME/.cshrc initial commands for each csh
\$HOME/.hushlogin suppresses login messages
\$HOME/.login user's login commands for csh

\$HOME/.profile user's login commands for sh and ksh
\$HOME/.rhosts private list of trusted hostname/username combinations
/etc/.login system-wide csh login commands
/etc/logindevperm login-based device permissions
/etc/motd message-of-the-day
/etc/passwd password file
/etc/profile system-wide sh and ksh login commands
/etc/shadow list of users' encrypted passwords
/usr/bin/sh user's default command interpreter
/var/adm/lastlog time of last login
/var/adm/loginlog record of failed login attempts
/var/adm/utmp accounting
/var/adm/wtmp accounting
/var/mail/*your-name* mailbox for user *your-name*
/etc/default/login Default value can be set for the following flags in **/etc/default/login**. For example: TIMEZONE=EST5EDT or PASSREQ=YES.

TIMEZONE: Sets the TZ environment variable of the shell.
HZ: Sets the HZ environment variable of the shell.
ULIMIT: Sets the file size limit for the login. Units are disk blocks. Default is zero (no limit).
CONSOLE: If set, root can login on that device only. This will not prevent execution of remote commands with **rsh(1)**. Comment out this line to allow login by root.
PASSREQ: Determines if login requires a password.
ALTSHELL: Determines if login should set the **SHELL** environment variable.
PATH: Sets the initial shell **PATH** variable.
SUPATH: Sets the initial shell **PATH** variable for root.
TIMEOUT: Sets the number of seconds (between 0 and 900) to wait before abandoning a login session.
UMASK: Sets the initial shell file creation mode mask. See **umask(1)**.
SYSLOG: determines whether the **syslog(3)** LOG_AUTH facility should be used to log all root logins at level LOG_NOTICE and multiple failed login attempts at LOG_CRIT.
SLEEPTIME If present sets the number of seconds to wait before login failure is printed to the screen and another login and another login attempt is allowed. Default is 4 seconds; Minimum is 0 seconds. Maximum is 5 seconds.

SEE ALSO

csh(1), **ksh(1)**, **mail(1)**, **mailx(1)**, **passwd(1)**, **rlogin(1)**, **rsh(1)**, **sh(1)**, **shell_builtins(1)**, **telnet(1)**, **admintool(1M)**, **in.rlogind(1M)**, **in.telnetd(1M)**, **logins(1M)**, **newgrp(1M)**, **quota(1M)**, **su(1M)**, **syslogd(1M)**, **useradd(1M)**, **userdel(1M)**, **syslog(3)**, **hosts.equiv(4)**,

logindevperm(4), loginlog(4), passwd(4), profile(4), shadow(4), environ(5)

DIAGNOSTICS

Login incorrect The user name or the password cannot be matched.

Not on system console

Root login denied. Check the **CONSOLE** setting in **/etc/default/login**.

No directory! Logging in with home=/

The user's home directory named in the **passwd(4)** database cannot be found or has the wrong permissions. Contact your system administrator.

No shell

Cannot execute the shell named in the **passwd(4)** database. Contact your system administrator.

WARNINGS

If you use the **CONSOLE** setting to disable root logins, you should arrange that remote command execution by root is also disabled. See **rsh(1)**, **rcmd(3N)**, and **hosts.equiv(4)** for further details.

NAME	logname – get the name of the user running the process
SYNOPSIS	logname
AVAILABILITY	SUNWesu
DESCRIPTION	logname returns the name of the user running the process.
FILES	/etc/profile
SEE ALSO	env(1) , login(1) , cuserid(3S) , environ(5)
EXIT CODES	logname returns the following exit codes: 0 If logname is successful. 1 If logname cannot identify the user.

NAME	logout – shell built-in function to exit from a login session
SYNOPSIS csh	logout
DESCRIPTION csh	Terminate a login shell.
SEE ALSO	csh(1) , login(1)

NAME	look – find words in the system dictionary or lines in a sorted list
SYNOPSIS	<code>/usr/bin/look [-d] [-f] [-tc] string [filename]</code>
AVAILABILITY	SUNWdoc
DESCRIPTION	<p>The look command consults a sorted <i>filename</i> and prints all lines that begin with <i>string</i>. If no <i>filename</i> is specified, look uses <code>/usr/share/lib/dict/words</code> with collating sequence -df.</p> <p>look limits the length of a word to search for to 256 characters.</p>
OPTIONS	<p>-d Dictionary order. Only letters, digits, TAB and SPACE characters are used in comparisons.</p> <p>-f Fold case. Upper case letters are not distinguished from lower case in comparisons.</p> <p>-tc Set termination character. All characters to the right of <i>c</i> in <i>string</i> are ignored.</p>
FILES	<code>/usr/share/lib/dict/words</code> spelling list
SEE ALSO	<code>grep(1)</code> , <code>sort(1)</code>

NAME	lookbib – find references in a bibliographic database								
SYNOPSIS	lookbib <i>database</i>								
AVAILABILITY	SUNWdoc								
DESCRIPTION	<p>A bibliographic reference is a set of lines, constituting fields of bibliographic information. Each field starts on a line beginning with a '%', followed by a key-letter, then a blank, and finally the contents of the field, which may continue until the next line starting with '%'. lookbib uses an inverted index made by indxib to find sets of bibliographic references. It reads keywords typed after the '>' prompt on the terminal, and retrieves records containing all these keywords. If nothing matches, nothing is returned except another '>' prompt.</p> <p>It is possible to search multiple databases, as long as they have a common index made by indxib(1). In that case, only the first argument given to indxib is specified to lookbib. If lookbib does not find the index files (the .i[abc] files), it looks for a reference file with the same name as the argument, without the suffixes. It creates a file with a .ig suffix, suitable for use with fgrep (see grep(1)). lookbib then uses this fgrep file to find references. This method is simpler to use, but the .ig file is slower to use than the .i[abc] files, and does not allow the use of multiple reference files.</p>								
FILES	<table><tr><td>x.ia</td><td></td></tr><tr><td>x.ib</td><td></td></tr><tr><td>x.ic</td><td>index files</td></tr><tr><td>x.ig</td><td>reference file</td></tr></table>	x.ia		x.ib		x.ic	index files	x.ig	reference file
x.ia									
x.ib									
x.ic	index files								
x.ig	reference file								
SEE ALSO	addbib (1), grep (1), indxib (1), refer (1), roffbib (1), sortbib (1)								
BUGS	Probably all dates should be indexed, since many disciplines refer to literature written in the 1800s or earlier.								

NAME	lorder – find ordering relation for an object or library archive
SYNOPSIS	lorder <i>filename</i> . . .
DESCRIPTION	<p>The input is one or more object or library archive <i>filenames</i> (see ar(1)). The standard output is a list of pairs of object file or archive member names; the first file of the pair refers to external identifiers defined in the second. The output may be processed by tsort(1) to find an ordering of a library suitable for one-pass access by ld. Note: The link editor ld is capable of multiple passes over an archive in the portable archive format (see ar(4)) and does not require that lorder be used when building an archive. The usage of the lorder command may, however, allow for a more efficient access of the archive during the link edit process.</p> <p>The following example builds a new library from existing .o files.</p> <pre>ar -cr library `lorder *.o tsort`</pre>
FILES	<p><i>TMPDIR</i>/*symref temporary files</p> <p><i>TMPDIR</i>/*symdef temporary files</p> <p><i>TMPDIR</i> usually /var/tmp but can be redefined by setting the environment variable TMPDIR (see tempnam() in tempnam(3S))</p>
SEE ALSO	ar (1), ld (1), tsort (1), tempnam (3S), ar (4)
NOTES	<p>lorder will accept as input any object or archive file, regardless of its suffix, provided there is more than one input file. If there is but a single input file, its suffix must be .o.</p> <p>The length of the filename for TMPDIR is limited to whatever sed allows.</p>

NAME	lp, cancel – send/cancel requests to an LP print service
SYNOPSIS	<pre> lp [-c] [-m] [-p] [-s] [-w] [-d <i>dest</i>] [-f <i>form-name</i> [-d <i>any</i>]] [-H <i>special-handling</i>] [-n <i>number</i>] [-o <i>option</i>] [-P <i>page-list</i>] [-q <i>priority-level</i>] [-S <i>character-set</i> [-d <i>any</i>]] [-S <i>print-wheel</i> [-d <i>any</i>]] [-t <i>title</i>] [-T <i>content-type</i> [-r]] [-y <i>mode-list</i>] [<i>file ...</i>] lp -i request-ID... [-c] [-m] [-p] [-s] [-w] [-d <i>dest</i>] [-f <i>form-name</i> [-d <i>any</i>]] [-H <i>special-handling</i>] [-n <i>number</i>] [-o <i>option</i>] [-P <i>page-list</i>] [-q <i>priority-level</i>] [-S <i>character-set</i> [-d <i>any</i>]] [-S <i>print-wheel</i> [-d <i>any</i>]] [-t <i>title</i>] [-T <i>content-type</i> [-r]] [-y <i>mode-list</i>] cancel [<i>request-ID...</i>] [<i>printer...</i>] cancel -u login-ID-list [<i>printer...</i>] </pre>
AVAILABILITY	SUNWlpu
DESCRIPTION	<p>The first form of the lp command arranges for the named <i>file(s)</i> and associated information (collectively called a <i>request</i>) to be printed. If no file names are specified on the command line, the standard input is assumed. The standard input may be specified along with a named <i>file(s)</i> on the command line by listing the file name(s) and specifying ‘-’ (dash) for the standard input. The <i>files</i> will be printed in the order in which they appear on the shell command line.</p> <p>The LP print service associates a unique <i>request-ID</i> (with the -i option) with each request and displays it on the standard output. This <i>request-ID</i> can be used later with the -i option when canceling or changing a request, or when determining its status. (See the section on cancel for details about canceling a request, and lpstat(1) for information about checking the status of a print request.)</p> <p>The second form of lp is used to change the options for a request. The print request identified by the <i>request-ID</i> is changed according to the printing options specified with this shell command. The printing options available are the same as those with the first form of the lp shell command. If the request has finished printing, the change is rejected. If the request is already printing, it will be stopped and restarted from the beginning (unless the -P option has been given).</p> <p>The cancel command allows users to cancel print requests previously sent with the lp command. The first form of cancel permits cancellation of requests based on their <i>request-ID</i>. The second form of cancel permits cancellation of requests based on the <i>login-ID</i> of their owner.</p>
Sending a Print Request	The first form of the lp command is used to send a print request to a particular printer or group of printers.

OPTIONS

Options to **lp** always precede any file names, but may be specified in any order. The following options are available for **lp** :

- c** Make a copy of the *file* before printing. Normally, *file* will not be copied, but will be linked whenever possible. If the **-c** option is not given, then the user should be careful not to remove any *file* before the request has been printed in its entirety. It should also be noted that if the **-c** option is not specified, any changes made to the named *file* after the request is made but before it is printed will be reflected in the printed output.

- d dest** Choose *dest* as the printer or class of printers that is to do the printing. If *dest* is a printer, then the request will be printed only on that specific printer. If *dest* is a class of printers, then the request will be printed on the first available printer that is a member of the class. If *dest* is **any**, then the request will be printed on any printer which can handle it. Under certain conditions, (unavailability of printers, file space limitations, and so on) requests for specific destinations may not be accepted (see **lpstat(1)**). By default, *dest* is taken from the environment variable **LPDEST** (if it is set). Otherwise, a default destination (if one exists) for the computer system is used. Destination names vary between systems (see **lpstat(1)**).

- f form-name [-d any]**
 Print the request on the form *form-name*. The LP print service ensures that the form is mounted on the printer. If *form-name* is requested with a printer destination that cannot support the form, the request is rejected. If *form-name* has not been defined for the system, or if the user is not allowed to use the form, the request is rejected (see **lpforms(1M)**). When the **-d any** option is given, the request is printed on any printer that has the requested form mounted and can handle all other needs of the print request.

- H special-handling**
 Print the request according to the value of *special-handling*. Acceptable values for *special-handling* are defined below:
 - hold** Do not print the request until notified. If printing has already begun, stop it. Other print requests will go ahead of a held request until it is resumed.

 - resume** Resume a held request. If the request had begun to print when held, it will be the next request printed, unless it is superseded by an **immediate** request.

- immediate** (Available only to LP administrators.)
Print the request next. If more than one request is assigned the most recent request is printed next. If a request is currently printing on the desired printer, a hold request must be issued to allow the immediate request to print.
- m** Send mail (see **mail(1)**) after the files have been printed. By default, no mail is sent upon normal completion of the print request.
- n number** Print *number* copies (default is 1) of the output.
- o option** Specify printer-dependent *options*. Several such *options* may be collected by specifying the **-o** keyletter more than once (**-o option₁ -o option₂ ... -o option_n**), or by specifying the **-o** keyletter followed by a list of options enclosed in double quotes (that is, **-o "option₁ option₂ ... option_n"**). The standard interface recognizes the following options:
- nobanner** Do not print a banner page with this request. (The administrator can disallow this option at any time.)
- nofilebreak**
Do not insert a form feed between the files given, if submitting a job to print more than one file.
- length=scaled-decimal-number**
Print this request with pages *scaled-decimal-number* lines long. A *scaled-decimal-number* is an optionally scaled decimal number that gives a size in lines, columns, inches, or centimeters, as appropriate. The scale is indicated by appending the letter "i" for inches, or the letter "c" for centimeters. For length or width settings, an unscaled number indicates lines or columns; for line pitch or character pitch settings, an unscaled number indicates lines per inch or characters per inch (the same as a number scaled with "i"). For example, **length=66** indicates a page length of 66 lines, **length=11i** indicates a page length of 11 inches, and **length=27.94c** indicates a page length of 27.94 centimeters.
- This option may not be used with the **-f** option.
- width=scaled-decimal-number**
Print this request with page-width set to *scaled-decimal-number* columns wide. (See the explanation of *scaled-decimal-numbers* in the discussion of **length**, above.) This option may not be used with the **-f** option.

lpi=*scaled-decimal-number*

Print this request with the line pitch set to *scaled-decimal-number* lines per inch. This option may not be used with the **-f** option.

cpi=*scaled-decimal-number*

Print this request with the character pitch set to *scaled-decimal-number* characters per inch. Character pitch can also be set to **pica** (representing **10** characters per inch) or **elite** (representing **12** characters per inch), or it can be **compressed** (representing as many characters as a printer can handle). There is no standard number of characters per inch for all printers; see the Terminfo database (see **terminfo(4)**) for the default character pitch for your printer.

This option may not be used with the **-f** option.

stty=*'stty-option-list'*

A list of options valid for the **stty** command; enclose the list with single quotes if it contains blanks.

- P** *page-list* Print the pages specified in *page-list*. This option can be used only if there is a filter available to handle it; otherwise, the print request will be rejected.
The *page-list* may consist of range(s) of numbers, single page numbers, or a combination of both. The pages will be printed in ascending order.
- p** Enable notification on completion of the print request. Delivery of the notification is dependent on additional software.
- q** *priority-level* Assign this request *priority-level* in the printing queue. The values of *priority-level* range from **0**, the highest priority, to **39**, the lowest priority. If a priority is not specified, the default for the print service is used, as assigned by the system administrator. A priority limit may be assigned to individual users by the system administrator.
- s** Suppress messages from **lp** such as those that begin with "**request id is...**"
- S** *character-set* [**-d** *any*]
- S** *print-wheel* [**-d** *any*]
Print this request using the specified *character-set* or *print-wheel*. If a form was requested and it requires a character set or print wheel other than the one specified with the **-S** option, the request is rejected.
For printers that take print wheels: if the print wheel specified is not one listed by the administrator as acceptable for the printer specified in this request, the request is rejected unless the print wheel is already mounted on the printer.
For printers that use selectable or programmable character sets: if the *character-set* specified is not one defined in the Terminfo database for the printer (see **terminfo(4)**), or is not an alias defined by the administrator, the request is rejected.

- When the **-d any** option is used, the request is printed on any printer that has the print wheel mounted or any printer that can select the character set, and that can handle the needs of the request.
- t *title*** Print *title* on the banner page of the output. If *title* is not supplied the name of the file is printed on the banner page. Enclose *title* in quotes if it contains blanks.
- T *content-type* [-r]** Print the request on a printer that can support the specified *content-type*. If no printer accepts this type directly, a filter will be used to convert the content into an acceptable type. If the **-r** option is specified, a filter will not be used. If **-r** is specified, and no printer accepts the *content-type* directly, the request is rejected. If the *content-type* is not acceptable to any printer, either directly or with a filter, the request is rejected.
- w** Write a message on the user's terminal after the *files* have been printed. If the user is not logged in, then mail will be sent instead.
- y *mode-list*** Print this request according to the printing modes listed in *mode-list*. The allowed values for *mode-list* are locally defined. This option may be used only if there is a filter available to handle it; otherwise, the print request will be rejected.

Canceling a Print Request

The **cancel** command cancels requests for print jobs made with the **lp** command. The first form allows a user to specify one or more *request-ID* of print jobs to be canceled. Alternatively, the user can specify one or more *printer*, on which only the currently printing job will be canceled.

The second form of **cancel** permits a user to cancel all of his or her own jobs on all printers. In this form the *printer* option can be used to restrict the printer(s) on which the user's job(s) will be canceled. Note: In this form, when the *printer* option is used, all jobs queued for that printer will be canceled. A printer class is not a valid argument.

Users without special privileges can cancel only requests associated with their own login IDs. The system administrator can cancel jobs submitted by any user. The *login-ID-list* must be enclosed in quotes if it contains blanks.

For printers that take mountable print wheels or font cartridges, if you do not specify a particular print wheel or font with the **-S** option, the one mounted at the time your request is printed will be used. Use the **lpstat -p *printer* -l** command to see which print wheels are available on a particular printer, or the **lpstat -S -l** command to find out what print wheels are available and on which printers. For printers that have selectable character sets, you will get the standard character set if you don't use the **-S** option.

FILES /var/spool/lp/*

SEE ALSO **enable(1)**, **lpstat(1)**, **mail(1)**, **postprint(1)**, **pr(1)**, **accept(1M)**, **lpadmin(1M)**, **lpfilter(1M)**, **lpforms(1M)**, **lpsched(1M)**, **lpssystem(1M)**, **lpusers(1M)**, **terminfo(4)**

NOTES Printers for which requests are not being accepted will not be considered when the **lp** command is run and the destination is **any**. (Use the **lpstat -a** command to see which printers are accepting requests.) On the other hand, if a request is destined for a class of printers and the class itself is accepting requests, then *all* printers in the class will be considered, regardless of their acceptance status.

NAME	lpc – line printer control program
SYNOPSIS	/usr/ucb/lpc [<i>command</i> [<i>parameter</i> . . .]]
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>lpc controls the operation of the printer, or of multiple printers. lpc commands can be used to start or stop a printer, disable or enable a printer's spooling queue, rearrange the order of jobs in a queue, or display the status of each printer—along with its spooling queue and printer daemon.</p> <p>With no arguments, lpc runs interactively, prompting with 'lpc>'. If arguments are supplied, lpc interprets the first as a <i>command</i> to execute; each subsequent argument is taken as a <i>parameter</i> for that command. The standard input can be redirected so that lpc reads commands from a file.</p>
USAGE Commands	<p>Commands may be abbreviated to an unambiguous substring. Note: the <i>printer</i> parameter is specified just by the name of the printer (as lw), not as you would specify it to lpr(1B) or lpq(1B) (not as -Plw).</p> <p>? [command] . . . help [command] . . . Display a short description of each command specified in the argument list, or, if no arguments are given, a list of the recognized commands.</p> <p>abort [all [printer . . .]] Terminate an active spooling daemon on the local host immediately and then disable printing (preventing new daemons from being started by lpr(1B)) for the specified printers. The abort command can only be used by the privileged user.</p> <p>clean [all [printer . . .]] Remove all files created in the spool directory by the daemon from the specified printer queue(s) on the local machine. The clean command can only be used by the privileged user.</p> <p>disable [all [printer . . .]] Turn the specified printer queues off. This prevents new printer jobs from being entered into the queue by lpr(1B). The disable command can only be used by the privileged user.</p> <p>down [all [printer . . .] [message] Turn the specified printer queue off, disable printing and put <i>message</i> in the printer status file. The message does not need to be quoted, the remaining arguments are treated like echo(1). This is normally used to take a printer down and let others know why (lpq(1B) indicates that the printer is down, as does the status command).</p>

enable [**all** | [*printer* . . .]]

Enable spooling on the local queue for the listed printers, so that **lpr**(1B) can put new jobs in the spool queue. The **enable** command can only be used by the privileged user.

exit

quit Exit from **lpc**.

restart [**all** | [*printer* . . .]]

Attempt to start a new printer daemon. This is useful when some abnormal condition causes the daemon to die unexpectedly leaving jobs in the queue. This command can be run by any user.

start [**all** | [*printer* . . .]]

Enable printing and start a spooling daemon for the listed printers. The **start** command can only be used by the privileged user.

status [**all** | [*printer* . . .]]

Display the status of daemons and queues on the local machine. This command can be run by any user.

stop [**all** | [*printer* . . .]]

Stop a spooling daemon after the current job completes and disable printing. The **stop** command can only be used by the privileged user.

topq printer [*job#* . . .] [*user* . . .]

Move the print job(s) specified by *job#* or those job(s) belonging to *user* to the top (head) of the printer queue. The **topq** command can only be used by the privileged user.

up [**all** | [*printer* . . .]] Enable everything and start a new printer daemon. Undoes the effects of **down**.

FILES

*/var/spool/lp/**

/var/spool/lp/system/pstatus

SEE ALSO

echo(1), **lpq**(1B), **lpr**(1B), **lprm**(1B), **lpsched**(1M)

DIAGNOSTICS

?Ambiguous command

The abbreviation you typed matches more than one command.

?Invalid command

You typed a command or abbreviation that was not recognized.

?Privileged command

You used a command can be executed only by the privileged user.

lpc: printer: unknown printer to the print service

The **printer** was not found in the LP database. Usually this is a typing mistake; however, it may indicate that the printer does not exist on the system. Use **'lptstat -p'** to find the reason.

lpc: error on opening queue to spooler

The connection to **lpsched** on the local machine failed. This usually means the printer server started at boot time has died or is hung. Check if the printer spooler daemon **/usr/lib/lp/lpsched** is running.

lpc: Can't send message to LP print service**lpc: Can't receive message from LP print service**

These indicate that the LP print service has been stopped. Get help from the system administrator.

lpc: Received unexpected message from LP print service

It is likely there is an error in this software. Get help from system administrator.

NAME	lpq – display the queue of printer jobs				
SYNOPSIS	<code>/usr/ucb/lpq [-P <i>printer</i>] [-l] [+ [<i>interval</i>]] [<i>job#...</i>] [<i>username...</i>]</code>				
AVAILABILITY	SUNWscpu				
DESCRIPTION	<p>lpq displays the contents of a printer queue. It reports the status of jobs specified by <i>job#</i>, or all jobs owned by the user specified by <i>username</i>. lpq reports on all jobs in the default printer queue when invoked with no arguments.</p> <p>For each print job in the queue, lpq reports the user's name, current position, the names of input files comprising the job, the job number (by which it is referred to when using lprm(1B)) and the total size in bytes. Normally, only as much information as will fit on one line is displayed. Jobs are normally queued on a first-in-first-out basis. Filenames comprising a job may be unavailable, such as when lpr is used at the end of a pipeline; in such cases the filename field indicates the standard input.</p> <p>If lpq warns that there is no daemon present (that is, due to some malfunction), the lpc(1B) command can be used to restart a printer daemon.</p>				
OPTIONS	<p>-P <i>printer</i> Display information about the queue for the specified <i>printer</i>. In the absence of the -P option, the queue to the printer specified by the PRINTER variable in the environment is used. If the PRINTER variable is not set, and the LPDEST environment variable is not set, the queue for the default printer is used.</p> <p>-l Display queue information in long format; includes the name of the host from which the job originated.</p> <p>+ [<i>interval</i>] Display the spool queue periodically until it empties. This option clears the terminal screen before reporting on the queue. If an <i>interval</i> is supplied, lpq sleeps that number of seconds in between reports.</p>				
FILES	<table border="0"> <tr> <td><code>/var/spool/lp</code></td> <td>spooling directory</td> </tr> <tr> <td><code>/var/spool/lp/tmp/system_name/*-0</code></td> <td>request files specifying jobs</td> </tr> </table>	<code>/var/spool/lp</code>	spooling directory	<code>/var/spool/lp/tmp/system_name/*-0</code>	request files specifying jobs
<code>/var/spool/lp</code>	spooling directory				
<code>/var/spool/lp/tmp/system_name/*-0</code>	request files specifying jobs				
SEE ALSO	lp (1), lpr (1B), lprm (1B), lpc (1B), lpsched (1M)				
DIAGNOSTICS	<p>printer is printing The lpq program queries the spooler LPSCHED about the status of the printer. If the printer is disabled, the superuser can restart the spooler using lpc(1B).</p> <p>printer waiting for auto-retry (offline ?) The daemon could not open the printer device. The printer may be turned off-line. This message can also occur if a printer is out of paper, the paper is jammed, and so on. Another possible cause is that a process, such as an output filter, has exclusive use of the device. The only recourse in this case is to kill the offending process and restart the printer with lpc.</p>				

waiting for *host* to come up

A daemon is trying to connect to the remote machine named *host*, in order to send the files in the local queue. If the remote machine is up, **lpd** on the remote machine is probably dead or hung and should be restarted using **lpc**.

sending to *host*

The files are being transferred to the remote *host*, or else the local daemon has hung while trying to transfer the files.

printer disabled reason:

The printer has been marked as being unavailable with **lpc**.

lpq: The LP print service isn't running or can't be reached.

The **lpsched** process overseeing the spooling queue does not exist. This normally occurs only when the daemon has unexpectedly died. You can restart the printer daemon with **lpc**.

lpr: *printer*: unknown printer

The **printer** was not found in the System V LP database. Usually this is a typing mistake; however, it may indicate that the printer does not exist on the system. Use '**lptstat -p**' to find the reason.

lpr: error on opening queue to spooler

The connection to **lpsched** on the local machine failed. This usually means the printer server started at boot time has died or is hung. Check if the printer spooler daemon **/usr/lib/lpsched** is running.

lpr: Can't send message to LP print service**lpr: Can't receive message from LP print service**

These indicate that the LP print service has been stopped. Get help from the system administrator.

lpr: Received unexpected message from LP print service

It is likely there is an error in this software. Get help from system administrator.

NOTES

Output formatting is sensitive to the line length of the terminal; this can result in widely-spaced columns.

NAME	lpr – send a job to the printer
SYNOPSIS	<code>/usr/ucb/lpr [-P <i>printer</i>] [-# <i>copies</i>] [-C <i>class</i>] [-J <i>job</i>] [-T <i>title</i>] [-i [<i>indent</i>]] [-w <i>cols</i>] [-B] [-m] [-h] [-s] [-filter_option] [filename ...]</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	lpr forwards printer jobs to a spooling area for subsequent printing as facilities become available. Each printer job consists of copies of, or, with -s , complete pathnames of each <i>filename</i> you specify. The spool area is managed by the line printer spooler, lpsched . lpr reads from the standard input if no files are specified.
OPTIONS	<p>-P <i>printer</i> Send output to the named <i>printer</i>. In the absence of the -P option, the queue to the printer specified by the PRINTER variable in the environment is used. If the PRINTER variable is not set, and the LPDEST environment variable is not set, the queue for the default printer is used.</p> <p>-# <i>copies</i> Produce the number of <i>copies</i> indicated for each named file. For example:</p> <p style="padding-left: 40px;">lpr -#3 index.c lookup.c produces three copies of index.c, followed by three copies of lookup.c. On the other hand,</p> <p style="padding-left: 40px;">cat index.c lookup.c lpr -#3 generates three copies of the concatenation of the files.</p> <p>-C <i>class</i> Print <i>class</i> as the job classification on the burst page. For example,</p> <p style="padding-left: 40px;">lpr -C Operations new.index.c replaces the system name (the name returned by <i>hostname</i>) with Operations on the burst page, and prints the file new.index.c.</p> <p>-J <i>job</i> Print <i>job</i> as the job name on the burst page. Normally, lpr uses the first file's name.</p> <p>-T <i>title</i> Use <i>title</i> instead of the file name for the title used by pr(1).</p> <p>-i[<i>indent</i>] Indent output <i>indent</i> SPACE characters. Eight SPACE characters is the default.</p> <p>-w <i>cols</i> Use <i>cols</i> as the page width for pr.</p> <p>-m Send mail upon completion.</p> <p>-h Suppress printing the burst page.</p> <p>-s Use the full pathnames (not symbolic links) of the files to be printed rather than trying to copy them. This means the data files should not be modified or removed until they have been printed. -s only prevents copies of local files from being made. Jobs from remote hosts are copied</p>

anyway. **-s** only works with named data files; if the **lpr** command is at the end of a pipeline, the data is copied to the spool.

filter_option The following single letter options notify the line printer spooler that the files are not standard text files. The spooling daemon will use the appropriate filters to print the data accordingly.

- p** Use **pr** to format the files (**lpr -p** is very much like **pr | lpr**).
- l** Print control characters and suppress page breaks.
- t** The files contain **troff(1)** (cat phototypesetter) binary data.
- n** The files contain data from **ditroff** (device independent troff).
- d** The files contain data from **tex** (DVI format from Stanford).
- g** The files contain standard plot data as produced by the **plot(1B)** routines.
- v** The files contain a raster image. The printer must support an appropriate imaging model such as PostScript® in order to print the image.
- c** The files contain data produced by *cifplot*.
- f** Interpret the first character of each line as a standard FORTRAN carriage control character.

If no *filter_option* is given (and the printer can interpret PostScript), the string '%!' as the first two characters of a file indicates that it contains PostScript commands.

These filter options offer a standard user interface, and all options may not be available for, nor applicable to, all printers.

FILES	/etc/passwd	personal identification
	/usr/lib/lp/lpsched	System V line printer spooler
	/var/spool/lp/tmp/*	directories used for spooling
	/var/spool/lp/tmp/system/*-0	spooler control files
	/var/spool/lp/tmp/system/*-N	(N is an integer and > 0) data files specified in '*-0' files

SEE ALSO **lp(1), lpc(1B) lpq(1B), lprm(1B), plot(1B), pr(1), troff(1), lpsched(1M)**

DIAGNOSTICS **lpr: printer: unknown printer**

The **printer** was not found in the LP database. Usually this is a typing mistake; however, it may indicate that the printer does not exist on the system. Use **'lptstat -p'** to find the reason.

lpr: error on opening queue to spooler

The connection to **lpsched** on the local machine failed. This usually means the printer server started at boot time has died or is hung. Check if the printer spooler daemon **/usr/lib/lpsched** is running.

lpr: printer: printer queue is disabled

This means the queue was turned off with

/usr/etc/lpc disable printer

to prevent **lpr** from putting files in the queue. This is normally done by the system manager when a printer is going to be down for a long time. The printer can be turned back on by a privileged user with **lpc**.

lpr: Can't send message to the LP print service

lpr: Can't receive message from the LP print service

These indicate that the LP print service has been stopped. Get help from the system administrator.

lpr: Received unexpected message from LP print service

It is likely there is an error in this software. Get help from system administrator.

lpr: There is no filter to convert the file content

Use the '**lpstat -p -l**' command to find a printer that can handle the file type directly, or consult with your system administrator.

lpr: cannot access the file

Make sure file names are valid.

NOTES

lp is the preferred interface.

Command-line options cannot be combined into a single argument as with some other commands. The command:

lpr -fs

is not equivalent to

lpr -f -s

Placing the **-s** flag first, or writing each option as a separate argument, makes a link as expected.

lpr -p is not precisely equivalent to **pr | lpr**. **lpr -p** puts the current date at the top of each page, rather than the date last modified.

Fonts for **troff(1)** and **T_EX[®]** reside on the printer host. It is currently not possible to use local font libraries.

lpr objects to printing binary files.

The **-s** option, intended to use symbolic links in SunOS, does not use symbolic links in the compatibility package. Instead, the complete path names are used. Also, the copying is avoided only for print jobs that are run from the printer host itself. Jobs added to the queue from a remote host are always copied into the spool area. That is, if the printer does not reside on the host that **lpr** is run from, the spooling system makes a copy the file to print, and places it in the spool area of the printer host, regardless of **-s**.

NAME	<code>lprm</code> – remove jobs from the printer queue
SYNOPSIS	<code>/usr/ucb/lprm</code> [<code>-Pprinter</code>] [<code>-</code>] [<code>job # ...</code>] [<code>username ...</code>]
AVAILABILITY	SUNWscpu
DESCRIPTION	<p><code>lprm</code> removes a job or jobs from a printer's spooling queue. Since the spool directory is protected from users, using <code>lprm</code> is normally the only method by which a user can remove a job.</p> <p>Without any arguments, <code>lprm</code> deletes the job that is currently active, provided that the user who invoked <code>lprm</code> owns that job.</p> <p>When the privileged user specifies a <i>username</i>, <code>lprm</code> removes all jobs belonging to that user.</p> <p>You can remove a specific job by supplying its job number as an argument, which you can obtain using <code>lpq</code>(1B). For example:</p> <pre>example% lpq -Phost host is ready and printing Rank Owner Job Files Total Size active wendy 385 standard input 35501 bytes example% lprm -Phost 385</pre> <p><code>lprm</code> reports the names of any files it removes, and is silent if there are no applicable jobs to remove.</p> <p><code>lprm</code> Sends the request to cancel a job to the print spooler, <code>LPSCHED</code>.</p>
OPTIONS	<p><code>-Pprinter</code> Specify the queue associated with a specific printer. Otherwise the value of the <code>PRINTER</code> variable in the environment is used. If the <code>PRINTER</code> variable is not set, and the <code>LPDEST</code> environment variable is not set, the queue for the default printer is used.</p> <p><code>-</code> Remove all jobs owned by you. If invoked by the privileged user, all jobs in the spool are removed. Job ownership is determined by the user's login name and host name on the machine where the <code>lpr</code> command was executed.</p>
FILES	<code>/var/spool/lp/*</code> spooling directories
SEE ALSO	<code>lp</code> (1), <code>lpq</code> (1B), <code>lpr</code> (1B), <code>lpsched</code> (1M)
DIAGNOSTICS	<p><code>lprm: printer: unknown printer</code> The <code>printer</code> was not found in the System V LP database. Usually this is a typing mistake; however, it may indicate that the printer does not exist on the system. Use <code>'lptstat -p'</code> to find the reason.</p>

lprm: error on opening queue to spooler

The connection to **lpsched** on the local machine failed. This usually means the printer server started at boot time has died or is hung. Check if the printer spooler daemon **/usr/lib/lpsched** is running.

lprm: Can't send message to the LP print service**lprm: Can't receive message from the LP print service**

These indicate that the LP print service has been stopped. Get help from the system administrator.

lprm: Received unexpected message from the LP print service

It is likely there is an error in this software. Get help from system administrator.

lprm: Can't cancel request

You are not allowed to remove another's request.

NOTES

An active job may be incorrectly identified for removal by an **lprm** command issued with no arguments. During the interval between an **lpq(1B)** command and the execution of **lprm**, the next job in queue may have become active; that job may be removed unintentionally if it is owned by you. To avoid this, supply **lprm** with the job number to remove when a critical job that you own is next in line.

Only the privileged user can remove print jobs submitted from another host.

lp is the preferred interface.

NAME	lpstat – print information about the status of the LP print service
SYNOPSIS	lpstat [-d] [-r] [-R] [-s] [-t] [-a <i>[list]</i>] [-c <i>[list]</i>] [-f <i>[list]</i> [-l]] [-o <i>[list]</i>] [-p <i>[list]</i> [-D] [-l]] [-P] [-S <i>[list]</i> [-l]] [-u <i>[login-ID-list]</i>] [-v <i>[list]</i>]
AVAILABILITY	SUNWlpu
DESCRIPTION	<p>The lpstat command prints information about the current status of the LP print service. If no options are given, then lpstat prints the status of all the user's print requests made by lp (see lp(1)). Any arguments that are not <i>options</i> are assumed to be <i>request-IDs</i> as returned by lp. The lpstat command prints the status of such requests. The <i>options</i> may appear in any order and may be repeated and intermixed with other arguments. Some of the keyletters below may be followed by an optional <i>list</i> that can be in one of two forms: a list of items separated from one another by a comma, or a list of items separated from one another by spaces enclosed in quotes. For example:</p> <p style="padding-left: 40px;">example% lpstat -u "user1 user2 user3"</p> <p>Specifying all after any keyletter that takes <i>list</i> as an argument causes all information relevant to the keyletter to be printed. For example, the command:</p> <p style="padding-left: 40px;">example% lpstat -o all</p> <p>prints the status of all output requests.</p> <p>The omission of a <i>list</i> following such key letters causes all information relevant to the key letter to be printed. For example, the command:</p> <p style="padding-left: 40px;">example% lpstat -o</p> <p>prints the status of all output requests.</p>
OPTIONS	<p>-a <i>[list]</i> Reports whether print destinations are accepting requests. <i>list</i> is a list of intermixed printer names and class names.</p> <p>-c <i>[list]</i> Print name of all classes and their members. <i>list</i> is a list of class names.</p> <p>-d Print the system default destination for output requests.</p> <p>-f <i>[list]</i> [-l] Print a verification that the forms in <i>list</i> are recognized by the LP print service. <i>list</i> is a list of forms; the default is all. The -l option will list the form descriptions.</p> <p>-o <i>[list]</i> Print the status of output requests: <i>list</i> is a list of intermixed printer names, class names, and <i>request-IDs</i>. The keyletter -o may be omitted.</p> <p>-p <i>[list]</i> [-D] [-l] Print the status of printers. <i>list</i> is a list of printer names. If the -D option is given, a brief description is printed for each printer in <i>list</i>. If the -l option is given, and the printer is on the local machine, a full description of each printer's configuration is given, including the form mounted, the acceptable</p>

- content and printer types, a printer description, the interface used, and so on.
- P** Print the paper types.
 - r** Print the status of the LP request scheduler.
 - R** Print a number showing the position of the job in the print queue.
 - s** Print a status summary, including the status of the LP scheduler, the system default destination, a list of class names and their members, a list of printers and their associated devices, a list of the machines sharing print services, a list of all forms currently mounted, and a list of all recognized character sets and print wheels.
 - S [list] [-l]** Print a verification that the character sets or the print wheels specified in *list* are recognized by the LP print service. Items in *list* can be character sets or print wheels; the default for the list is **all**. If the **-l** option is given, each line is appended by a list of printers that can handle the print wheel or character set. The list also shows whether the print wheel or character set is mounted, or specifies the built-in character set into which it maps.
 - t** Print all status information. This includes all the information obtained with the **-s** option, plus the acceptance and idle/busy status of all printers.
 - u [login-ID-list]** Print the status of output requests for users. The *login-ID-list* argument may include any or all of the following constructs:
 - login-ID* a user on any system
 - system_name!login-ID* a user on system *system_name*
 - system_name!all* all users on system *system_name*
 - all!login-ID** a user on all systems
 - all** all users on all systems
 - v [list]** Print the names of printers and the path names of the devices associated with them or remote system names for network printers: *list* is a list of printer names.

FILES

*/etc/lp/**
*/var/spool/lp/**

ENVIRONMENT

Date and time format is based on locale specified by the LC_ALL, LC_TIME or LANG environments, in that order of priority.

SEE ALSO

enable(1), **lp(1)**

NAME	lptest – generate lineprinter ripple pattern
SYNOPSIS	/usr/ucb/lptest [<i>length</i> [<i>count</i>]]
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>lptest writes the traditional “ripple test” pattern on standard output. In 96 lines, this pattern will print all 96 printable ASCII characters in each position. While originally created to test printers, it is quite useful for testing terminals, driving terminal ports for debugging purposes, or any other task where a quick supply of random data is needed.</p> <p>The <i>length</i> argument specifies the output line length if the the default length of 79 is inappropriate.</p> <p>The <i>count</i> argument specifies the number of output lines to be generated if the default count of 200 is inappropriate.</p>
NOTES	<p>if <i>count</i> is to be specified, <i>length</i> must be also be specified.</p> <p>This command is obsolete.</p>

NAME	ls – list contents of directory
SYNOPSIS	ls [-abcCdffgilLmnopqrRstux1] [<i>names</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>For each directory argument, ls lists the contents of the directory; for each file argument, ls repeats its name and any other information requested. The output is sorted alphabetically by default. When no argument is given, the current directory is listed. When several arguments are given, the arguments are first sorted appropriately, but file arguments appear before directories and their contents.</p> <p>There are three major listing formats. The default format for output directed to a terminal is multi-column with entries sorted down the columns. The -1 option allows single column output and -m enables stream output format. In order to determine output formats for the -C, -x, and -m options, ls uses an environment variable, COLUMNS, to determine the number of character positions available on one output line. If this variable is not set, the terminfo(4) database is used to determine the number of columns, based on the environment variable TERM. If this information cannot be obtained, 80 columns are assumed.</p> <p>The mode printed under the -l option consists of ten characters. The first character may be one of the following:</p> <ul style="list-style-type: none"> d the entry is a directory; l the entry is a symbolic link; b the entry is a block special file; c the entry is a character special file; p the entry is a fifo (a.k.a. “named pipe”) special file; - the entry is an ordinary file. <p>The next 9 characters are interpreted as three sets of three bits each. The first set refers to the owner’s permissions; the next to permissions of others in the user-group of the file; and the last to all others. Within each set, the three characters indicate permission to read, to write, and to execute the file as a program, respectively. For a directory, “execute” permission is interpreted to mean permission to search the directory for a specified file.</p> <p>ls -l (the long list) prints its output as follows:</p> <pre style="margin-left: 40px;">-rwxrwxrwx 1 smith dev 10876 May 16 9:42 part2</pre> <p>Reading from right to left, you see that the current directory holds one file, named part2. Next, the last time that file’s contents were modified was 9:42 A.M. on May 16. The file contains 10,876 characters, or bytes. The owner of the file, or the user, belongs to the group dev (perhaps indicating “development”), and his or her login name is smith. The number, in this case 1, indicates the number of links to file part2; see cp(1). Finally, the dash and letters tell you that user, group, and others have permissions to read, write, and execute part2.</p>

The execute (**x**) symbol here occupies the third position of the three-character sequence. A **-** in the third position would have indicated a denial of execution permissions.

The permissions are indicated as follows:

- r** the file is readable
- w** the file is writable
- x** the file is executable
- the indicated permission is *not* granted
- l** mandatory locking occurs during access (the set-group-ID bit is on and the group execution bit is off)
- s** the set-user-ID or set-group-ID bit is on, and the corresponding user or group execution bit is also on
- S** undefined bit-state (the set-user-ID bit is on and the user execution bit is off)
- t** the 1000 (octal) bit, or sticky bit, is on (see **chmod(1)**), and execution is on
- T** the 1000 bit is turned on, and execution is off (undefined bit-state)

For user and group permissions, the third position is sometimes occupied by a character other than **x** or **-**. **s** also may occupy this position, referring to the state of the set-ID bit, whether it be the user's or the group's. The ability to assume the same ID as the user during execution is, for example, used during login when you begin as root but need to assume the identity of the user you login as.

In the case of the sequence of group permissions, **l** may occupy the third position. **l** refers to mandatory file and record locking. This permission describes a file's ability to allow other files to lock its reading or writing permissions during access.

For others permissions, the third position may be occupied by **t** or **T**. These refer to the state of the sticky bit and execution permissions.

OPTIONS

- a** List all entries, including those that begin with a dot (**.**), which are normally not listed.
- b** Force printing of non-printable characters to be in the octal **\ddd** notation.
- c** Use time of last modification of the i-node (file created, mode changed, and so forth) for sorting (**-t**) or printing (**-l**).
- C** Multi-column output with entries sorted down the columns. This is the default output format.
- d** If an argument is a directory, list only its name (not its contents); often used with **-l** to get the status of a directory.
- f** Force each argument to be interpreted as a directory and list the name found in each slot. This option turns off **-l**, **-t**, **-s**, and **-r**, and turns on **-a**; the order is the order in which entries appear in the directory.
- F** Put a slash (**/**) after each filename if the file is a directory, an asterisk (*****) if the file is an executable, and an at-sign (**@**) if the file is a symbolic link.
- g** The same as **-l**, except that the owner is not printed.
- i** For each file, print the i-node number in the first column of the report.

- l** List in long format, giving mode, number of links, owner, group, size in bytes, and time of last modification for each file (see above). If the file is a special file, the size field instead contains the major and minor device numbers rather than a size. If the file is a symbolic link, the filename is printed followed by “→” and the pathname of the referenced file.
- L** If an argument is a symbolic link, list the file or directory the link references rather than the link itself.
- m** Stream output format; files are listed across the page, separated by commas.
- n** The same as **-l**, except that the owner’s **UID** and group’s **GID** numbers are printed, rather than the associated character strings.
- o** The same as **-l**, except that the group is not printed.
- p** Put a slash (/) after each filename if the file is a directory.
- q** Force printing of non-printable characters in file names as the character question mark (?).
- r** Reverse the order of sort to get reverse alphabetic or oldest first as appropriate.
- R** Recursively list subdirectories encountered.
- s** Give size in blocks, including indirect blocks, for each entry.
- t** Sort by time stamp (latest first) instead of by name. The default is the last modification time. (See **-n** and **-c**.)
- u** Use time of last access instead of last modification for sorting (with the **-t** option) or printing (with the **-l** option).
- x** Multi-column output with entries sorted across rather than down the page.
- 1** Print one entry per line of output.

EXAMPLES

An example of a file’s permissions is:

```
-rwxr--r--
```

This describes a file that is readable, writable, and executable by the user and readable by the group and others.

Another example of a file’s permissions is:

```
-rwsr-xr-x
```

This describes a file that is readable, writable, and executable by the user, readable and executable by the group and others, and allows its user- ID to be assumed, during execution, by the user presently executing it.

Another example of a file’s permissions is:

```
-rw-rwl---
```

This describes a file that is readable and writable only by the user and the group and can be locked during access.

An example of a command line:

```
example% ls -a
```

This command prints the names of all files in the current directory, including those that begin with a dot (.), which normally do not print.

Another example of a command line:

```
example% ls -aisn
```

This command provides information on all files, including those that begin with a dot (a), the i-number—the memory address of the i-node associated with the file—printed in the left-hand column (i); the size (in blocks) of the files, printed in the column to the right of the i-numbers (s); finally, the report is displayed in the numeric version of the long list, printing the UID (instead of user name) and GID (instead of group name) numbers associated with the files.

When the sizes of the files in a directory are listed, a total count of blocks, including indirect blocks, is printed.

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

LC_TIME

Determines how **ls** handles date and time formats. In the "C" locale, date and time handling follows the U.S. rules.

FILES

/etc/group	group IDs for ls -l and ls -g
/etc/passwd	user IDs for ls -l and ls -o
/usr/share/lib/terminfo/?/*	terminal information database

SEE ALSO `chmod(1)`, `cp(1)`, `find(1)`, `terminfo(4)`, `environ(5)`

NOTES Unprintable characters in file names may confuse the columnar output options.
The total block count will be incorrect if there are hard links among the files.

NAME	ls – list the contents of a directory
SYNOPSIS	/usr/ucb/ls [-aAcCdffgilLqrRstu1] <i>filename</i> ...
AVAILABILITY	SUNWscpu
DESCRIPTION	For each <i>filename</i> which is a directory, ls lists the contents of the directory; for each <i>filename</i> which is a file, ls repeats its name and any other information requested. By default, the output is sorted alphabetically. When no argument is given, the current directory is listed. When several arguments are given, the arguments are first sorted appropriately, but file arguments are processed before directories and their contents.
Permissions Field	<p>The mode printed under the -l option contains 10 characters interpreted as follows. If the first character is:</p> <ul style="list-style-type: none"> d entry is a directory; b entry is a block-type special file; c entry is a character-type special file; l entry is a symbolic link; p entry is a FIFO (also known as “named pipe”) special file; s entry is an AF_UNIX address family socket, or – entry is a plain file. <p>The next 9 characters are interpreted as three sets of three bits each. The first set refers to owner permissions; the next refers to permissions to others in the same user-group; and the last refers to all others. Within each set the three characters indicate permission respectively to read, to write, or to execute the file as a program. For a directory, “execute” permission is interpreted to mean permission to search the directory. The permissions are indicated as follows:</p> <ul style="list-style-type: none"> r the file is readable; w the file is writable; x the file is executable; – the indicated permission is not granted. <p>The group-execute permission character is given as s if the file has the set-group-id bit set; likewise the owner-execute permission character is given as s if the file has the set-user-id bit set.</p> <p>The last character of the mode (normally x or ‘–’) is true if the 1000 bit of the mode is on. See chmod(1) for the meaning of this mode. The indications of set-ID and 1000 bits of the mode are capitalized (S and T respectively) if the corresponding execute permission is <i>not</i> set.</p> <p>When the sizes of the files in a directory are listed, a total count of blocks, including indirect blocks is printed.</p>

OPTIONS

- a** List all entries; in the absence of this option, entries whose names begin with a '.' are *not* listed (except for the privileged user, for whom **ls** normally prints even files that begin with a '.').
- A** Same as **-a**, except that '.' and '..' are not listed.
- c** Use time of last edit (or last mode change) for sorting or printing.
- C** Force multi-column output, with entries sorted down the columns; for **ls**, this is the default when output is to a terminal.
- d** If argument is a directory, list only its name (not its contents); often used with **-l** to get the status of a directory.
- f** Force each argument to be interpreted as a directory and list the name found in each slot. This option turns off **-l**, **-t**, **-s**, and **-r**, and turns on **-a**; the order is the order in which entries appear in the directory.
- F** Mark directories with a trailing slash ('/'), executable files with a trailing asterisk (*), symbolic links with a trailing at-sign (@), and **AF_UNIX** address family sockets with a trailing equals sign (=).
- g** For **ls**, show the group ownership of the file in a long output.
- i** For each file, print the i-node number in the first column of the report.
- l** List in long format, giving mode, number of links, owner, size in bytes, and time of last modification for each file. If the file is a special file the size field will instead contain the major and minor device numbers. If the time of last modification is greater than six months ago, it is shown in the format '*month date year*'; files modified within six months show '*month date time*'. If the file is a symbolic link the pathname of the linked-to file is printed preceded by '—>'.
-L If argument is a symbolic link, list the file or directory the link references rather than the link itself.
- q** Display non-graphic characters in filenames as the character ?; for **ls**, this is the default when output is to a terminal.
- r** Reverse the order of sort to get reverse alphabetic or oldest first as appropriate.
- R** Recursively list subdirectories encountered.
- s** Give size of each file, including any indirect blocks used to map the file, in kilobytes.
- t** Sort by time modified (latest first) instead of by name.
- u** Use time of last access instead of last modification for sorting (with the **-t** option) and/or printing (with the **-l** option).
- 1** Force one entry per line output format; this is the default when output is not to a terminal.

FILES **/etc/group** to get group ID for '**ls -g**'
 /etc/passwd to get user ID's for '**ls -l**' and '**ls -o**'

NOTES NEWLINE and TAB are considered printing characters in filenames.
 The output device is assumed to be 80 columns wide.
 The option setting based on whether the output is a teletype is undesirable as '**ls -s**' is much different than '**ls -s | lpr**'. On the other hand, not doing this setting would make old shell scripts which used **ls** almost certain losers.
 Unprintable characters in file names may confuse the columnar output options.

NAME	m4 – macro processor
SYNOPSIS	m4 [-e] [-s] [-Bint] [-Hint] [-Sint] [-Tint] [-Dname [=val]] [-Uname] [<i>filename ...</i>]
DESCRIPTION	<p>The m4 command is a macro processor intended as a front end for C, assembler, and other languages. Each of the argument files is processed in order; if there are no files, or if a file name is -, the standard input is read. The processed text is written on the standard output.</p> <p>Macro calls have the form:</p> <p style="padding-left: 40px;"><i>name(arg1,arg2, ..., argn)</i></p> <p>The (must immediately follow the name of the macro. If the name of a defined macro is not followed by a (, it is deemed to be a call of that macro with no arguments. Potential macro names consist of alphanumeric characters and underscore (_), where the first character is not a digit.</p> <p>Leading unquoted blanks, tabs, and new-lines are ignored while collecting arguments. Left and right single quotes are used to quote strings. The value of a quoted string is the string stripped of the quotes.</p> <p>When a macro name is recognized, its arguments are collected by searching for a matching right parenthesis. If fewer arguments are supplied than are in the macro definition, the trailing arguments are taken to be null. Macro evaluation proceeds normally during the collection of the arguments, and any commas or right parentheses that happen to turn up within the value of a nested call are as effective as those in the original input text. After argument collection, the value of the macro is pushed back onto the input stream and rescanned.</p>
OPTIONS	<p>The options and their effects are as follows:</p> <p>-e Operate interactively. Interrupts are ignored and the output is unbuffered.</p> <p>-s Enable line sync output for the C preprocessor (<i>#line ...</i>)</p> <p>-Bint Change the size of the push-back and argument collection buffers from the default of 4,096.</p> <p>-Hint Change the size of the symbol table hash array from the default of 199. The size should be prime.</p> <p>-Sint Change the size of the call stack from the default of 100 slots. Macros take three slots, and non-macro arguments take one.</p> <p>-Tint Change the size of the token buffer from the default of 512 bytes.</p> <p>To be effective, the above flags must appear before any file names and before any -D or -U flags:</p> <p>-Dname[=val] Defines <i>name</i> to <i>val</i> or to null in <i>val</i>'s absence.</p> <p>-Uname Undefines <i>name</i>.</p>

USAGE

m4 makes available the following built-in macros. These macros may be redefined, but once this is done the original meaning is lost. Their values are null unless otherwise stated.

define	The second argument is installed as the value of the macro whose name is the first argument. Each occurrence of $\$n$ in the replacement text, where n is a digit, is replaced by the n -th argument. Argument 0 is the name of the macro; missing arguments are replaced by the null string; $\$#$ is replaced by the number of arguments; $\$*$ is replaced by a list of all the arguments separated by commas; $\$@$ is like $\$*$, but each argument is quoted (with the current quotes).
undefine	Removes the definition of the macro named in its argument.
defn	Returns the quoted definition of its argument(s). It is useful for renaming macros, especially built-ins.
pushdef	Like define , but saves any previous definition.
popdef	Removes current definition of its argument(s), exposing the previous one, if any.
ifdef	If the first argument is defined, the value is the second argument, otherwise the third. If there is no third argument, the value is null. The word unix is predefined.
shift	Returns all but its first argument. The other arguments are quoted and pushed back with commas in between. The quoting nullifies the effect of the extra scan that will subsequently be performed.
changequote	Change quote symbols to the first and second arguments. The symbols may be up to five characters long. changequote without arguments restores the original values (that is, <code>`</code> and <code>~</code>).
changecom	Change left and right comment markers from the default <code>#</code> and new-line. With no arguments, the comment mechanism is effectively disabled. With one argument, the left marker becomes the argument and the right marker becomes new-line. With two arguments, both markers are affected. Comment markers may be up to five characters long.
divert	m4 maintains 10 output streams, numbered 0-9. The final output is the concatenation of the streams in numerical order; initially stream 0 is the current stream. The divert macro changes the current output stream to its (digit-string) argument. Output diverted to a stream other than 0 through 9 is discarded.
undivert	This macro causes immediate output of text from diversions named as arguments, or all diversions if no argument. Text may be undiverted into another diversion. Undiverting discards the diverted text.
divnum	Returns the value of the current output stream.
dnl	Reads and discards characters up to and including the next new-line.

ifelse	This macro has three or more arguments. If the first argument is the same string as the second, then the value is the third argument. If not, and if there are more than four arguments, the process is repeated with arguments 4, 5, 6 and 7. Otherwise, the value is either the fourth string, or, if it is not present, null.
incr	Returns the value of its argument incremented by 1. The value of the argument is calculated by interpreting an initial digit-string as a decimal number.
decr	Returns the value of its argument decremented by 1.
eval	Evaluates its argument as an arithmetic expression, using 32-bit arithmetic. Operators include +, -, *, /, %, ** (exponentiation), bitwise &, , ^, and ~; relationals; parentheses. Octal and hex numbers may be specified as in C. The second argument specifies the radix for the result; the default is 10. The third argument may be used to specify the minimum number of digits in the result.
len	Returns the number of characters in its argument.
index	Returns the position in its first argument where the second argument begins (zero origin), or -1 if the second argument does not occur.
substr	Returns a substring of its first argument. The second argument is a zero origin number selecting the first character; the third argument indicates the length of the substring. A missing third argument is taken to be large enough to extend to the end of the first string.
translit	Transliterates the characters in its first argument from the set given by the second argument to the set given by the third. No abbreviations are permitted.
include	Returns the contents of the file named in the argument.
sinclude	This macro is identical to include , except that it says nothing if the file is inaccessible.
syscmd	This macro executes the command given in the first argument. No value is returned.
sysval	This macro is the return code from the last call to syscmd .
maketemp	Fills in a string of XXXXXX in its argument with the current process ID.
m4exit	This macro causes immediate exit from m4 . Argument 1, if given, is the exit code; the default is 0.
m4wrap	Argument 1 will be pushed back at final EOF; example: m4wrap(- cleanup() ^)
errprint	Prints its argument on the diagnostic output file.
dumpdef	Prints current names and definitions, for the named items, or for all if no arguments are given.

traceon This macro with no arguments, turns on tracing for all macros (including built-ins). Otherwise, turns on tracing for named macros.

traceoff Turns off trace globally and for any macros specified. Macros specifically traced by **traceon** can be untraced only by specific calls to **traceoff**.

SEE ALSO

as(1)

NAME	mach – display the processor type of the current host
SYNOPSIS	/usr/ucb/mach
AVAILABILITY	SUNWscpu
DESCRIPTION	The mach command displays the processor-type of the current host.
SEE ALSO	arch(1B) , machid(1) , uname(1) , sysinfo(2) , uname(2)

NAME	machid, sun, iAPX286, i286, i386, i486, i860, pdp11, sparc, u3b, u3b2, u3b5, u3b15, vax, u370 – get processor type truth value																						
SYNOPSIS	sun iAPX286 i386 pdp11 sparc u3b u3b2 u3b5 u3b15 vax u370																						
DESCRIPTION	<p>The following commands will return a true value (exit code of 0) if you are using an instruction set that the command name indicates.</p> <table border="0"> <tr> <td style="padding-right: 10px;">sun</td> <td>True if you are on a Sun system.</td> </tr> <tr> <td>iAPX286</td> <td>True if you are on a computer using an iAPX286 processor.</td> </tr> <tr> <td>i386</td> <td>True if you are on a computer using an iAPX386 processor.</td> </tr> <tr> <td>pdp11</td> <td>True if you are on a PDP-11/45™ or PDP-11/70™.</td> </tr> <tr> <td>sparc</td> <td>True if you are on a computer using a SPARC-family processor.</td> </tr> <tr> <td>u3b</td> <td>True if you are on a 3B20 computer.</td> </tr> <tr> <td>u3b2</td> <td>True if you are on a 3B2 computer.</td> </tr> <tr> <td>u3b5</td> <td>True if you are on a 3B5 computer.</td> </tr> <tr> <td>u3b15</td> <td>True if you are on a 3B15 computer.</td> </tr> <tr> <td>vax</td> <td>True if you are on a VAX-11/750™ or VAX-11/780™.</td> </tr> <tr> <td>u370</td> <td>True if you are on an IBM® System/370™ computer.</td> </tr> </table> <p>The commands that do not apply will return a false (non-zero) value. These commands are often used within makefiles (see make(1S)) and shell procedures (see sh(1)) to increase portability.</p>	sun	True if you are on a Sun system.	iAPX286	True if you are on a computer using an iAPX286 processor.	i386	True if you are on a computer using an iAPX386 processor.	pdp11	True if you are on a PDP-11/45™ or PDP-11/70™.	sparc	True if you are on a computer using a SPARC-family processor.	u3b	True if you are on a 3B20 computer.	u3b2	True if you are on a 3B2 computer.	u3b5	True if you are on a 3B5 computer.	u3b15	True if you are on a 3B15 computer.	vax	True if you are on a VAX-11/750™ or VAX-11/780™.	u370	True if you are on an IBM® System/370™ computer.
sun	True if you are on a Sun system.																						
iAPX286	True if you are on a computer using an iAPX286 processor.																						
i386	True if you are on a computer using an iAPX386 processor.																						
pdp11	True if you are on a PDP-11/45™ or PDP-11/70™.																						
sparc	True if you are on a computer using a SPARC-family processor.																						
u3b	True if you are on a 3B20 computer.																						
u3b2	True if you are on a 3B2 computer.																						
u3b5	True if you are on a 3B5 computer.																						
u3b15	True if you are on a 3B15 computer.																						
vax	True if you are on a VAX-11/750™ or VAX-11/780™.																						
u370	True if you are on an IBM® System/370™ computer.																						
SEE ALSO	make(1S), sh(1), test(1), true(1), uname(1)																						
NOTES	The machid family of commands is obsolete. Use uname -p and uname -m instead.																						

NAME	mail, rmail – read mail or send mail to users
SYNOPSIS	
Sending mail	mail [<i>-tw</i>] [<i>-m message_type</i>] <i>recipient...</i> rmail [<i>-tw</i>] [<i>-m message_type</i>] <i>recipient...</i>
Reading mail	mail [<i>-ehpPqr</i>] [<i>-f filename</i>]
Debugging	mail [<i>-x debug_level</i>] [<i>other_mail_options</i>] <i>recipient...</i> mail [<i>-T mailsurr_file</i>] <i>recipient...</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>A <i>recipient</i> is usually a user name recognized by login(1). When <i>recipients</i> are named, mail assumes a message is being sent. It reads from the standard input up to an end-of-file (CTRL-D) or, if reading from a terminal device, until it reads a line consisting of just a period. When either of those indicators is received, mail adds the <i>letter</i> to the <i>mailfile</i> for each <i>recipient</i>.</p> <p>A <i>letter</i> is composed of some <i>header lines</i> followed by a blank line followed by the <i>message content</i>. The <i>header lines</i> section of the letter consists of one or more UNIX postmarks:</p> <p style="text-align: center;">From <i>sender date_and_time</i> [remote from <i>remote_system_name</i>]</p> <p>followed by one or more standardized message header lines of the form:</p> <p style="text-align: center;"><i>keyword-name</i>: [<i>printable text</i>]</p> <p>where <i>keyword-name</i> is comprised of any printable, non-whitespace characters other than colon (:). A Content-Length: header line, indicating the number of bytes in the <i>message content</i> will always be present unless the letter consists of only header lines with no message content. A Content-Type: header line that describes the type of the <i>message content</i> (such as text, binary, multipart, etc.) will also be present unless the letter consists of only header lines with no message content. Header lines may be continued on the following line if that line starts with white space.</p>
OPTIONS	
Sending mail	<p>The following command-line arguments affect sending mail:</p> <p>-m <i>message_type</i> A Message-Type: line is added to the message header with the value of <i>message_type</i>.</p> <p>-t A To: line is added to the message header for each of the intended recipients.</p>

-w A letter is sent to a remote recipient without waiting for the completion of the remote transfer program.

If a letter is found to be undeliverable, it is returned to the sender with diagnostics that indicate the location and nature of the failure. If **mail** is interrupted during input, the message is saved in the file **dead.letter** to allow editing and resending. **dead.letter** is always appended to, thus preserving any previous contents. The initial attempt to append to (or create) **dead.letter** will be in the current directory. If this fails, **dead.letter** will be appended to (or created in) the user's login directory. If the second attempt also fails, no **dead.letter** processing will be done.

rmail only permits the sending of mail; **uucp**(1C) uses **rmail** as a security precaution. Any application programs that generate mail messages should be sure to invoke **rmail** rather than **mail** for message transport and/or delivery.

If the local system has the Basic Networking Utilities installed, mail may be sent to a recipient on a remote system. There are numerous ways to address mail to recipients on remote systems depending on the transport mechanisms available to the local system. The two most prevalent addressing schemes are UUCP-style and Domain-style. With UUCP-style addressing, remote recipients are specified by prefixing the recipient name with the remote system name and an exclamation point, such as **sysa!user**. If **cs**h is the default shell, **sysa!user** should be used. A series of system names separated by exclamation points can be used to direct a letter through an extended network (such as **sysa!sysb!sysc!user** or **sysa!sysb!sysc!user**). With Domain-style addressing, remote recipients are specified by appending an '@' and domain (and possibly sub-domain) information to the recipient name (such as **user@sf.att.com**). (The local System Administrator should be consulted for details on which addressing conventions are available on the local system.)

Reading Mail

The following command-line arguments affect reading mail:

- e** Mail is not printed. An exit value of 0 is returned if the user has mail; otherwise, an exit value of 1 is returned.
- h** A window of headers are initially displayed rather than the latest message. The display is followed by the ? prompt.
- p** All messages are printed without prompting for disposition.
- P** All messages are printed with *all* header lines displayed, rather than the default selective header line display.
- q** **mail** terminates after interrupts. Normally an interrupt causes only the termination of the message being printed.
- r** Messages are printed in first-in, first-out order.
- f filename** **mail** uses *filename* (such as **mbox**) instead of the default *mailfile*.

mail, unless otherwise influenced by command-line arguments, prints a user's mail messages in last-in, first-out order. The default mode for printing messages is to display only those header lines of immediate interest. These include, but are not limited to, the UNIX **From** and **>From** postmarks, **From:**, **Date:**, **Subject:**, and **Content-Length:** header lines, and any recipient header lines such as **To:**, **Cc:**, **Bcc:**, and so forth. After the header lines have been displayed, **mail** will display the contents (body) of the message only if it

contains no unprintable characters. Otherwise, **mail** will issue a warning statement about the message having binary content and **not** display the content. (This may be overridden via the **p** command. See below.)

For each message, the user is prompted with a ? and a line is read from the standard input. The following commands are available to determine the disposition of the message:

#	Print the number of the current message.
-	Print previous message.
<new-line>,+, or n	Print the next message.
! <i>command</i>	Escape to the shell to do <i>command</i> .
a	Print message that arrived during the mail session.
d, or dp	Delete the current message and print the next message.
d n	Delete message number <i>n</i> . Do not go on to next message.
dq	Delete message and quit mail .
h	Display a window of headers around current message.
h n	Display a window of headers around message number <i>n</i> .
h a	Display headers of all messages in the user's <i>mailfile</i> .
h d	Display headers of messages scheduled for deletion.
m [<i>persons</i>]	Mail (and delete) the current message to the named <i>person(s)</i> .
n	Print message number <i>n</i> .
p	Print current message again, overriding any indications of binary (that is, unprintable) content.
P	Override default brief mode and print current message again, displaying all header lines.
q, or CTRL-D	Put undeleted mail back in the <i>mailfile</i> and quit mail .
r [<i>users</i>]	Reply to the sender, and other <i>user(s)</i> , then delete the message.
s [<i>files</i>]	Save message in the named <i>file(s)</i> (mbox is default) and delete the message.
u [<i>n</i>]	Undelete message number <i>n</i> (default is last read).
w [<i>files</i>]	Save message contents, without any header lines, in the named <i>files</i> (mbox is default) and delete the message.
x	Put all mail back in the <i>mailfile</i> unchanged and exit mail .
y [<i>files</i>]	Same as save.
?	Print a command summary.

When a user logs in, the presence of mail, if any, is usually indicated. Also, notification is made if new mail arrives while using **mail**.

The permissions of *mailfile* may be manipulated using **chmod** in two ways to alter the function of **mail**. The other permissions of the file may be read-write (0666), read-only (0664), or neither read nor write (0660) to allow different levels of privacy. If changed to other than the default (mode 0660), the file will be preserved even when empty to perpetuate the desired permissions. (The administrator may override this file preservation using the **DEL_EMPTY_MAILFILE** option of **mailcnfg**.)

The group id of the mailfile must be **mail** to allow new messages to be delivered, and the mailfile must be writable by group **mail**.

Debugging

The following command-line arguments cause **mail** to provide debugging information:

-T mailsurr_file mail displays how it will parse and interpret the **mailsurr** file.
-x debug_level mail creates a trace file containing debugging information.

The **-T** option requires an argument that will be taken as the pathname of a test **mailsurr** file. If NULL (as in **-T ""**), the system **mailsurr** file will be used. To use, type '**mail -T test_file recipient**' and some trivial message (like "testing"), followed by a line with either just a dot (".") or a CTRL-D. The result of using the **-T** option will be displayed on standard output and show the inputs and resulting transformations as **mailsurr** is processed by the **mail** command for the indicated recipient. Mail messages will never actually be sent or delivered when the **-T** option is used.

The **-x** option causes **mail** to create a file named **/tmp/MLDBGprocess_id** that contains debugging information relating to how **mail** processed the current message. The absolute value of *debug_level* controls the verbosity of the debug information. Zero implies no debugging. If *debug_level* is greater than zero, the debug file will be retained **only** if **mail** encountered some problem while processing the message. If *debug_level* is less than zero the debug file will always be retained. The *debug_level* specified via **-x** overrides any specification of **DEBUG** in **/etc/mail/mailcnfg**. The information provided by the **-x** option is esoteric and is probably only useful to System Administrators. The output produced by the **-x** option is a superset of that provided by the **-T** option.

Delivery Notification

Several forms of notification are available for mail by including one of the following lines in the message header.

Transport-Options: [/options]

Default-Options: [/options]

>**To:** *recipient* [*/options*]

Where the “*/options*” may be one or more of the following:

/delivery Inform the sender that the message was successfully delivered to the *recipient's* mailbox.

/nodelivery Do not inform the sender of successful deliveries.

/ignore Do not inform the sender of **unsuccessful** deliveries.

/return Inform the sender if mail delivery fails. Return the failed message to the sender.

/report Same as **/return** except that the original message is not returned.

The default is **/nodelivery/return**. If contradictory options are used, the first will be recognized and later, conflicting, terms will be ignored.

ENVIRONMENT

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

LC_CTYPE

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

FILES

dead.letter unmailable text

/etc/passwd to identify sender and locate recipients

/etc/mail/maillsurr routing / name translation information

/etc/mail/mailcnfg initialization information

\$HOME/mbox saved mail

\$MAIL variable containing path name of *mailfile*

/tmp/ma* temporary file

/tmp/MLDBG* debug trace file

/var/mail/*.lock lock for mail directory

/var/mail/:saved directory for holding temp files to prevent loss of data in the event of a system crash

/var/mail/user incoming mail for *user*; that is, the *mailfile*

SEE ALSO

chmod(1), **login(1)**, **mailx(1)**, **vacation(1)**, **write(1)**, **uucp(1C)**, **environ(5)**

Solaris Advanced User's Guide

NOTES

The interpretation and resulting action taken because of the header lines described in the Delivery Notifications section above will only occur if this version of **mail** is installed on the system where the delivery (or failure) happens. Earlier versions of **mail** may not support any types of delivery notification.

Conditions sometimes result in a failure to remove a lock file.

After an interrupt, the next message may not be printed; printing may be forced by typing a **p**.

NAME	mailstats – print statistics collected by sendmail
SYNOPSIS	mailstats [<i>filename</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>mailstats prints out the statistics collected by the sendmail(1M) program on mailer usage. These statistics are collected if the file indicated by the S configuration option of sendmail(1M) (defined in /etc/mail/sendmail.cf), exists. The default statistics file is /etc/mail/sendmail.st. mailstats first prints the time that the statistics file was created and the last time it was modified. It will then print a table with one row for each mailer specified in the configuration file. The first column is the mailer number, followed by the symbolic name of the mailer. The next two columns refer to the number of messages received by sendmail, and the last two columns refer to messages sent by sendmail. The number of messages and their total size (in 1024 byte units) is given. No numbers are printed if no messages were sent (or received) for any mailer.</p> <p>You might want to add an entry to /var/spool/cron/crontabs/root to reinitialize the statistics file once a night. Copy /dev/null into the statistics file or otherwise truncate it to reset the counters.</p>
FILES	<p>/dev/null /var/spool/cron/crontabs/root /etc/mail/sendmail.st default statistics file. /etc/mail/sendmail.cf sendmail(1M) configuration file.</p>
SEE ALSO	sendmail(1M)
NOTES	mailstats should read the configuration file instead of having a hard-wired table mapping mailer numbers to names.

NAME	mailx, mail, Mail – interactive message processing system
SYNOPSIS	<pre>mailx [-BdeHiInNUvV~] [-f [filename +folder]] [-T filename] [-u user] mailx [-BdFintUv~] [-b bcc] [-c cc] [-h number] [-r address] [-s subject] recipient ... /usr/ucb/mail ... /usr/ucb/Mail ...</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The mail utilities listed above provide a comfortable, flexible environment for sending and receiving mail messages electronically. The OPTIONS and USAGE documented below for mailx also apply to /usr/ucb/mail and /usr/ucb/Mail, except where noted.</p> <p>When reading mail, mailx provides commands to facilitate saving, deleting, and responding to messages. When sending mail, mailx allows editing, reviewing and other modification of the message as it is entered.</p> <p>Incoming mail is stored in a standard file for each user, called the mailbox for that user. When mailx is called to read messages, the mailbox is the default place to find them. As messages are read, they are marked to be moved to a secondary file for storage, unless specific action is taken, so that the messages need not be seen again. This secondary file is called the mbox and is normally located in the user's HOME directory (see MBOX in Environment Variables for a description of this file). Messages can be saved in other secondary files named by the user. Messages remain in a secondary file until forcibly removed.</p> <p>The user can access a secondary file by using the -f option of the mailx command. Messages in the secondary file can then be read or otherwise processed using the same Commands as in the primary mailbox. This gives rise within these pages to the notion of a current mailbox.</p>
OPTIONS	<p>On the command line, <i>options</i> start with a dash (-) and any other arguments are taken to be destinations (recipients). If no recipients are specified, mailx attempts to read messages from the mailbox.</p> <ul style="list-style-type: none"> -B Do not buffer standard input or standard output. -d Turn on debugging output. (Neither particularly interesting nor recommended.) -e Test for the presence of mail. mailx prints nothing and exits with a successful return code if there is mail to read. -F Record the message in a file named after the first recipient. Overrides the record variable, if set (see Environment Variables). -H Print header summary only. -i Ignore interrupts. See also ignore in Environment Variables.

- I** Include the newsgroup and article-id header lines when printing mail messages. This option requires the **-f** option to be specified.
- n** Do not initialize from the system default **mailx.rc** or **Mail.rc** file.
- N** Do not print initial header summary.
- U** Convert **uucp** style addresses to internet standards. Overrides the **conv** environment variable.
- v** Pass the **-v** flag to **sendmail(1M)**.
- V** Print the **mailx** version number and exit.
- b bcc** Set the blind carbon copy list to *bcc*. *bcc* should be enclosed in quotes if it contains more than one name.
- c cc** Set the carbon copy list to *cc*. *cc* should be enclosed in quotes if it contains more than one name.
- f [filename]** Read messages from *filename* instead of **mailbox**. If no *filename* is specified, the **mbox** is used.
- f [+folder]** Use the file *folder* in the folder directory (same as the **folder** command). The name of this directory is listed in the **folder** variable.
- h number** The number of network “hops” made so far. This is provided for network software to avoid infinite delivery loops. This option and its argument are passed to the delivery program.
- r address** Use *address* as the return address when invoking the delivery program. All tilde commands are disabled. This option and its argument is passed to the delivery program.
- s subject** Set the Subject header field to *subject*. *subject* should be enclosed in quotes if it contains embedded white space.
- t** Scan the input for “To:”, “Cc:”, and “Bcc:” fields. Any recipients on the command line will be ignored.
- T file** Message-id and article-id header lines are recorded in *file* after the message is read. This option also sets the **-I** option.
- u user** Read *user*’s **mailbox**. This is only effective if *user*’s **mailbox** is not read protected.
- ~** Interpret tilde escapes in the input even if not reading from a tty.

USAGE

Starting Mail

At start-up time, **mailx** tries to execute commands from the optional system-wide file (**/etc/mail/mailx.rc**) to initialize certain parameters, then from a private start-up file (**\$HOME/.mailrc**) for personalized variables. If invoked as **mail** or **Mail**, the system-wide start-up file **/etc/mail/Mail.rc** is used instead. With the exceptions noted below, regular commands are legal inside start-up files.

The most common use of a start-up file is to set up initial display options and alias lists. The following commands are not legal in the start-up file: **!**, **Copy**, **edit**, **followup**, **Followup**, **hold**, **mail**, **preserve**, **reply**, **Reply**, **shell**, and **visual**. An error in the start-up file

causes the remaining lines in the file to be ignored. The **mailrc** file is optional, and must be constructed locally.

At any time, the behavior of **mailx** is governed by a set of *environment variables*. These are flags and valued parameters which are set and cleared using the **set** and **unset** commands. See **Environment Variables** below for a summary of these parameters.

When reading mail, **mailx** is in *command mode*. A header summary of the first several messages is displayed, followed by a prompt indicating **mailx** can accept regular commands (see **Commands** below). When sending mail, **mailx** is in *input mode*. If no subject is specified on the command line, a prompt for the subject is printed.

As the message is typed, **mailx** reads the message and stores it in a temporary file. Commands may be entered by beginning a line with the tilde (~) escape character followed by a single command letter and optional arguments. See **Tilde Escapes** for a summary of these commands.

Reading Mail

Each message is assigned a sequential number, and there is at any time the notion of a current message, marked by a right angle bracket (>) in the header summary. Many commands take an optional list of messages (*message-list*) to operate on.

The default for *message-list* is the current message. A *message-list* is a list of message identifiers separated by spaces, which may include:

<i>n</i>	Message number <i>n</i> .										
.	The current message.										
^	The first undeleted message.										
\$	The last message.										
*	All messages.										
+	The next undeleted message.										
-	The previous undeleted message.										
<i>n-m</i>	An inclusive range of message numbers.										
<i>user</i>	All messages from <i>user</i> .										
<i>/string</i>	All messages with <i>string</i> in the Subject line (case ignored).										
<i>:c</i>	All messages of type <i>c</i> , where <i>c</i> is one of: <table> <tr> <td>d</td> <td>deleted messages</td> </tr> <tr> <td>n</td> <td>new messages</td> </tr> <tr> <td>o</td> <td>old messages</td> </tr> <tr> <td>r</td> <td>read messages</td> </tr> <tr> <td>u</td> <td>unread messages</td> </tr> </table>	d	deleted messages	n	new messages	o	old messages	r	read messages	u	unread messages
d	deleted messages										
n	new messages										
o	old messages										
r	read messages										
u	unread messages										

Note that the context of the command determines whether this type of message specification makes sense.

Other arguments are usually arbitrary strings whose usage depends on the command involved. Filenames, where expected, are expanded using the normal shell conventions (see **sh(1)**). Special characters are recognized by certain commands and are documented with the commands below.

Sending Mail

Recipients listed on the command line may be of three types: login names, shell commands, or alias groups. Login names may be any network address, including mixed network addressing. If mail is found to be undeliverable, an attempt is made to return it to the sender's **mailbox**. If the recipient name begins with a pipe symbol (|), the rest of the name is taken to be a shell command to pipe the message through. This provides an automatic interface with any program that reads the standard input, such as **lp(1)** for recording outgoing mail on paper. Alias groups are set by the **alias** command (see **Commands** below) or in a system start-up file (for example, **\$HOME/.mailrc**). Aliases are lists of recipients of any type.

Forwarding Mail

To forward a specific message, include it in a message to the desired recipients with the **~f** or **~m** tilde escapes. See **Tilde Escapes** below. To forward mail automatically, add a comma-separated list of addresses for additional recipients to the **.forward** file in your home directory. This is different from the format of the **alias** command, which takes a space-separated list instead. Note: forwarding addresses must be valid, or the messages will "bounce." You cannot, for instance, reroute your mail to a new host by forwarding it to your new address if it is not yet listed in the NIS aliases domain.

Commands

Regular commands are of the form
 [*command*] [*message-list*] [*arguments*]

In *input mode*, commands are recognized by the escape character, **~**, and lines not treated as commands are taken as input for the message.

If no command is specified in *command mode*, **next** is assumed.

The following is a complete list of **mailx** commands:

!shell-command

Escape to the shell. See **SHELL** in **Environment Variables**.

comment

Null command (comment). Useful in **mailrc** files.

= Print the current message number.

? Prints a summary of commands.

alias *alias name* ...

group *alias name* ...

Declare an alias for the given names. The names are substituted when *alias* is used as a recipient. Useful in the **mailrc** file. With no arguments, the command displays the list of defined aliases.

alternates *name* ...

Declare a list of alternate names for your login. When responding to a message, these names are removed from the list of recipients for the response. With no arguments, print the current list of alternate names. See also **allnet** in **Environment Variables**.

cd [*directory*]

chdir [*directory*]

Change directory. If *directory* is not specified, **\$HOME** is used.

copy [*filename*]
copy [*message-list*] *filename*
Copy messages to the file without marking the messages as saved. Otherwise equivalent to the **save** command.

Copy [*message-list*]
Save the specified messages in a file whose name is derived from the author of the message to be saved, without marking the messages as saved. Otherwise equivalent to the **Save** command.

delete [*message-list*]
Delete messages from the **mailbox**. If **autoprint** is set, the next message after the last one deleted is printed (see **Environment Variables**).

discard [*header-field...*]
ignore [*header-field...*]
Suppress printing of the specified header fields when displaying messages on the screen. Examples of header fields to ignore are **Status** and **Received**. The fields are included when the message is saved, unless the **alwaysignore** variable is set. The **More**, **Page**, **Print**, and **Type** commands override this command. If no header is specified, the current list of header fields being ignored is printed. See also the **undiscard** and **unignore** commands.

dp [*message-list*]
dt [*message-list*]
Delete the specified messages from the **mailbox** and print the next message after the last one deleted. Roughly equivalent to a **delete** command followed by a **print** command.

echo *string...*
Echo the given strings (like **echo(1)**).

edit [*message-list*]
Edit the given messages. The messages are placed in a temporary file and the **EDITOR** variable is used to get the name of the editor (see **Environment Variables**). Default editor is **ed(1)**.

exit
xit Exit from **mailx**, without changing the **mailbox**. No messages are saved in the **mbox** (see also **quit**).

field [*message-list*] *header-file*
Display the value of the header field in the specified message.

file [*filename*]
folder [*filename*]
Quit from the current file of messages and read in the specified file. Several special characters are recognized when used as file names:

%	the current mailbox .
%user	the mailbox for <i>user</i> .
#	the previous mail file.

& the current **mbox**.

+filename The named file in the *folder* directory (listed in the **folder** variable).
 With no arguments, print the name of the current mail file, and the number of messages and characters it contains.

folders Print the names of the files in the directory set by the **folder** variable (see **Environment Variables**).

Followup [*message*]
 Respond to a message, recording the response in a file whose name is derived from the author of the message. Overrides the **record** variable, if set. See also the **followup**, **Save**, and **Copy** commands and **outfolder** in **Environment Variables**.

followup [*message-list*]
 Respond to the first message in the *message-list*, sending the message to the author of each message in the *message-list*. The subject line is taken from the first message and the response is recorded in a file whose name is derived from the author of the first message. See also the **Followup**, **Save**, and **Copy** commands and **outfolder** in **Environment Variables**.

from [*message-list*]
 Print the header summary for the specified messages. If no messages are specified, print the header summary for the current message.

group *alias name* ...
alias *alias name* ...
 Declare an alias for the given names. The names are substituted when *alias* is used as a recipient. Useful in the **mailrc** file.

hheaders [*message*]
 Print the page of headers which includes the message specified. The **screen** variable sets the number of headers per page (see **Environment Variables**). See also the **z** command.

help Print a summary of commands.

hold [*message-list*]
preserve [*message-list*]
 Hold the specified messages in the **mailbox**.

if **s** | **r** | **t**
mail-commands
else
mail-commands
endif Conditional execution, where *s* executes following *mail-commands*, up to an **else** or **endif**, if the program is in *send* mode, *r* causes the *mail-commands* to be executed only in *receive* mode, and **t** causes the *mail-commands* to be executed only if **mailx** is being run from a terminal. Useful in the **mailrc** file.

inc Incorporate messages that arrive while you are reading the system mailbox. The new messages are added to the message list in the current **mail** session. This command does not commit changes made during the session, and prior

messages are not renumbered.

ignore [*header-field* ...]

discard [*header-field* ...]

Suppress printing of the specified header fields when displaying messages on the screen. Examples of header fields to ignore are **Status** and

.BR Cc . All fields are included when the message is saved. The **More**, **Page**, **Print** and **Type** commands override this command. If no header is specified, the current list of header fields being ignored is printed. See also the **undiscard** and **unignore** commands.

list Print all commands available. No explanation is given.

load [*message*] *filename*

The specified message is replaced by the message in the named file. *filename* should contain a single mail message including mail headers (as saved by the save command).

mail recipient ...

Mail a message to the specified recipients.

Mail recipient

Mail a message to the specified recipients, and record it in a file whose name is derived from the author of the message. Overrides the **record** variable, if set. See also the **Save** and **Copy** commands and **outfolder** in **Environment Variables**.

mbox [*message-list*]

Arrange for the given messages to end up in the standard **mbox** save file when **mailx** terminates normally. See **MBOX** in **Environment Variables** for a description of this file. See also the **exit** and **quit** commands.

more [*message-list*]

page [*message-list*]

Print the specified messages. If **crt** is set, the messages longer than the number of lines specified by the **crt** variable are paged through the command specified by the **PAGER** variable. The default command is **pg(1)** or if the **bsdcompat** variable is set, the default is **more(1)**. See **Environment Variables**. Same as the **print** and **type** commands.

More [*message-list*]

Page [*message-list*]

Print the specified messages on the screen, including all header fields. Overrides suppression of fields by the **ignore** command. Same as the **Print** and **Type** commands.

new [*message-list*]

New [*message-list*]

unread [*message-list*]

Unread [*message-list*]

Take a message list and mark each message as *not* having been read.

next [*message*]

Go to the next message matching *message*. If *message* is not supplied, this command finds the next message that was not deleted or saved. A *message-list* may be specified, but in this case the first valid message in the list is the only one used. This is useful for jumping to the next message from a specific user, since the name would be taken as a command in the absence of a real command. See the discussion of *message-list* above for a description of possible message specifications.

pipe [*message-list*] [*shell-command*]

| [*message-list*] [*shell-command*]

Pipe the message through the given *shell-command*. The message is treated as if it were read. If no arguments are given, the current message is piped through the command specified by the value of the **cmd** variable. If the **page** variable is set, a form feed character is inserted after each message (see **Environment Variables**).

preserve [*message-list*]

hold [*message-list*]

Preserve the specified messages in the **mailbox**.

print [*message-list*]

type [*message-list*]

Print the specified messages. If **crt** is set, the messages longer than the number of lines specified by the **crt** variable are paged through the command specified by the **PAGER** variable. The default command is **pg(1)** or if the **bsdcompat** variable is set, the default is **more(1)**. See **Environment Variables**. Same as the **more** and **page** commands.

Print [*message-list*]

Type [*message-list*]

Print the specified messages on the screen, including all header fields. Overrides suppression of fields by the **ignore** command. Same as the **More** and **Page** commands.

put [*filename*]

put [*message-list*] *filename*

Save the specified message in the given filename. Use the same conventions as the **print** command for which header fields are ignored.

Put [*filename*]

Put [*message-list*] *filename*

Save the specified message in the given filename. Overrides suppression of fields by the **ignore** command.

quit Exit from **mailx**, storing messages that were read in **mbox** and unread messages in the **mailbox**. Messages that have been explicitly saved in a file are deleted unless the **keepsave** variable is set.

reply [*message-list*]

respond [*message-list*]

replysender [*message-list*]

Send a response to the author of each message in the *message-list*. The subject line is taken from the first message. If **record** is set to a filename, a copy of the reply is added to that file. If the **replyall** variable is set, the actions of **Reply/Respond** and **reply/respond** are reversed. The **replysender** command is not affected by the **replyall** variable, but sends each reply only to the sender of each message.

Reply [*message*]

Respond [*message*]

replyall [*message*]

Reply to the specified message, including all other recipients of that message. If the variable **record** is set to a filename, a copy of the reply added to that file. If the **replyall** variable is set, the actions of **Reply/Respond** and **reply/respond** are reversed. The **replyall** command is not affected by the **replyall** variable, but always sends the reply to all recipients of the message.

retain Add the list of header fields named to the *retained list*. Only the header fields in the retain list are shown on your terminal when you print a message. All other header fields are suppressed. The set of retained fields specified by the **retain** command overrides any list of ignored fields specified by the **ignore** command. The **Type** and **Print** commands can be used to print a message in its entirety. If **retain** is executed with no arguments, it lists the current set of retained fields.

Save [*message-list*]

Save the specified messages in a file whose name is derived from the author of the first message. The name of the file is taken to be the author's name with all network addressing stripped off. See also the **Copy**, **followup**, and **Followup** commands and **outfolder** in **Environment Variables**.

save [*filename*]

save [*message-list*] *filename*

Save the specified messages in the given file. The file is created if it does not exist. The file defaults to **mbox**. The message is deleted from the **mailbox** when **mailx** terminates unless **keepsave** is set (see also **Environment Variables** and the **exit** and **quit** commands).

set

set *variable*

set *variable=string*

set *variable=number*

Define a *variable*. To assign a *value* to *variable*, separate the variable name from the value by an '=' (there must be no space before or after the '='). A variable may be given a null, string, or numeric *value*. To embed SPACE characters within a *value* enclose it in quotes.

With no arguments, **set** displays all defined variables and any values they might have. See **Environment Variables** for a description of all predefined **mail** variables.

shell Invoke an interactive shell. See also **SHELL** in **Environment Variables**.

size [*message-list*]

Print the size in characters of the specified messages.

source *filename*

Read commands from the given file and return to command mode.

top [*message-list*]

Print the top few lines of the specified messages. If the **toplines** variable is set, it is taken as the number of lines to print (see **Environment Variables**). The default is 5.

touch [*message-list*]

Touch the specified messages. If any message in *message-list* is not specifically saved in a file, it is placed in the **mbox**, or the file specified in the **MBOX** environment variable, upon normal termination. See **exit** and **quit**.

Type [*message-list*]

Print [*message-list*]

Print the specified messages on the screen, including all header fields. Overrides suppression of fields by the **ignore** command.

type [*message-list*]

print [*message-list*]

Print the specified messages. If **prt** is set, the messages longer than the number of lines specified by the **prt** variable are paged through the command specified by the **PAGER** variable. The default command is **pg(1)** See **Environment Variables**.

unalias [*alias*] ...

ungroup [*alias*] ...

Remove the definitions of the specified aliases.

undelete [*message-list*]

Restore the specified deleted messages. Will only restore messages deleted in the current mail session. If **autoprint** is set, the last message of those restored is printed (see **Environment Variables**).

undiscard [*header-field...*]

unignore [*header-field...*]

Remove the specified header fields from the list being ignored. If no header fields are specified, all header fields are removed from the list being ignored.

unretain [*header-field...*]

Remove the specified header fields from the list being retained. If no header fields are specified, all header fields are removed from the list being retained.

- unread** [*message-list*]
Unread [*message-list*]
 Same as the **new** command.
- unset variable** . . .
 Erase the specified variables. If the variable was imported from the environment (that is, an environment variable or exported shell variable), it cannot be unset from within **mailx**.
- version** Print the current version and release date of the **mailx** utility.
- visual** [*message-list*]
 Edit the given messages with a screen editor. The messages are placed in a temporary file and the **VISUAL** variable is used to get the name of the editor (see **Environment Variables**).
- write** [*message-list*] *filename*
 Write the given messages on the specified file, minus the header and trailing blank line. Otherwise equivalent to the **save** command.
- xit**
exit Exit from **mailx**, without changing the **mailbox**. No messages are saved in the **mbox** (see also **quit**).
- z**[+ | -] Scroll the header display forward or backward one screen–full. The number of headers displayed is set by the **screen** variable (see **Environment Variables**).

Tilde Escapes

The following *tilde escape* commands can be used when composing mail to send. These may be entered only from *input mode*, by beginning a line with the tilde escape character (~). See **escape** in **Environment Variables** for changing this special character. The escape character can be entered as text by typing it twice.

- ~!*shell-command*
 Escape to the shell. If present, run *shell-command*.
- ~.
 Simulate end of file (terminate message input).
- ~:*mail-command*
 ~_ *mail-command*
 Perform the command-level request. Valid only when sending a message while reading mail.
- ~?
 Print a summary of tilde escapes.
- ~**A** Insert the autograph string **Sign** into the message (see **Environment Variables**).
- ~**a** Insert the autograph string **sign** into the message (see **Environment Variables**).
- ~**b** *name* . . .
 Add the *names* to the blind carbon copy (**Bcc**) list. This is like the carbon copy (**Cc**) list, except that the names in the **Bcc** list are not shown in the header of the mail message.

- ~c** *name* ...
Add the *names* to the carbon copy (Cc) list.
- ~d** Read in the **dead-letter** file. See **DEAD** in **Environment Variables** for a description of this file.
- ~e** Invoke the editor on the partial message. See also **EDITOR** in **Environment Variables**.
- ~f** [*message-list*]
Forward the specified message, or the current message being read. Valid only when sending a message while reading mail. The messages are inserted into the message without alteration (as opposed to the **~m** escape).
- ~F** [*message-list*]
Forward the specified message, or the current message being read, including all header fields. Overrides the suppression of fields by the **ignore** command.
- ~h** Prompt for Subject line and To, Cc, and Bcc lists. If the field is displayed with an initial value, it may be edited as if you had just typed it.
- ~i** *variable*
Insert the value of the named variable into the text of the message. For example, **~A** is equivalent to '**~i Sign.**' Environment variables set and exported in the shell are also accessible by **~i**.
- ~m** [*message-list*]
Insert the listed messages, or the current message being read into the letter. Valid only when sending a message while reading mail. The text of the message is shifted to the right, and the string contained in the **indentprefix** variable is inserted as the leftmost characters of each line. If **indentprefix** is not set, a TAB character is inserted into each line.
- ~M** [*message-list*]
Insert the listed messages, or the current message being read, including the header fields, into the letter. Valid only when sending a message while reading mail. The text of the message is shifted to the right, and the string contained in the **indentprefix** variable is inserted as the leftmost characters of each line. If **indentprefix** is not set, a TAB character is inserted into each line. Overrides the suppression of fields by the **ignore** command.
- ~p** Print the message being entered.
- ~q** Quit from input mode by simulating an interrupt. If the body of the message is not null, the partial message is saved in **dead-letter**. See **DEAD** in **Environment Variables** for a description of this file.
- ~r** *filename*
~< *filename*
~<! *shell-command*
Read in the specified file. If the argument begins with an exclamation point (!), the rest of the string is taken as an arbitrary shell command and is executed, with the standard output inserted into the message.

- ~s** *string* ...
Set the subject line to *string*.
- ~t** *name* ...
Add the given *names* to the To list.
- ~v** Invoke a preferred screen editor on the partial message. The default visual editor is **vi(1)**. See also **VISUAL** in **Environment Variables**.
- ~w** *filename*
Write the message into the given file, without the header.
- ~x** Exit as with **~q** except the message is not saved in **dead-letter**.
- ~|** *shell-command*
Pipe the body of the message through the given *shell-command*. If the *shell-command* returns a successful exit status, the output of the command replaces the message.

Environment Variables

The following are environment variables taken from the execution environment and are not alterable within **mailx**.

HOME=*directory*

The user's home directory.

MAIL=*filename*

The name of the initial mailbox file to read (in lieu of the standard system mailbox). The default is **/var/mail/username**.

MAILRC=*filename*

The name of the start-up file. Default is **\$HOME/.mailrc**.

The following variables are internal **mailx** variables. They may be imported from the execution environment or set using the **set** command at any time. The **unset** command may be used to erase variables.

allnet All network names whose last component (login name) match are treated as identical. This causes the *message-list* message specifications to behave similarly. Disabled by default. See also the **alternates** command and the **metoo** variable.

alwaysignore

Ignore header fields with **ignore** everywhere, not just during **print** or **type**. Affects the **save**, **Save**, **copy**, **Copy**, **top**, **pipe**, and **write** commands, and the **~m** and **~f** tilde escapes. Enabled by default.

append

Upon termination, append messages to the end of the **mbox** file instead of prepending them. Disabled by default, but **append** is set in the global start-up file (which can be suppressed with the **-n** command line option).

askbcc Prompt for the Bcc list after the Subject is entered if it is not specified on the command line with the **-b** option. Disabled by default.

askcc Prompt for the Cc list after the Subject is entered if it is not specified on the command line with the **-c** option. Disabled by default.

- asksub** Prompt for subject if it is not specified on the command line with the **-s** option. Enabled by default.
- autoinc**
Automatically incorporate new messages into the current session as they arrive. This has an affect similar to issuing the **inc** command every time the command prompt is displayed. Disabled by default, but **autoinc** is set in the default system start-up file for **mailx**; it is not set for **/usr/ucb/mail** or **/usr/ucb/Mail**.
- autoprint**
Enable automatic printing of messages after **delete** and **undelete** commands. Disabled by default.
- bang** Enable the special-casing of exclamation points (!) in shell escape command lines as in **vi(1)**. Disabled by default.
- bsdcompat**
Set automatically if **mailx** is invoked as **mail** or **Mail**. Causes **mailx** to use **/etc/mail/Mail.rc** as the system start-up file. Changes the default pager to **more(1)**.
- cmd=shell-command**
Set the default command for the **pipe** command. No default value.
- conv=conversion**
Convert **uucp** addresses to the specified address style, which can be either:
internet
This requires a mail delivery program conforming to the RFC822 standard for electronic mail addressing.
optimize
Remove loops in **uucp(1C)** address paths (typically generated by the **reply** command). No rerouting is performed; **mail** has no knowledge of UUCP routes or connections.
Conversion is disabled by default. See also **sendmail(1M)** and the **-U** command-line option.
- crt[=number]**
Pipe messages having more than *number* lines through the command specified by the value of the **PAGER** variable (**pg(1)** or **more(1)** by default). If *number* is not specified, the current window size is used. Disabled by default.
- DEAD=filename**
The name of the file in which to save partial letters in case of untimely interrupt. Default is **\$HOME/dead.letter**.
- debug** Enable verbose diagnostics for debugging. Messages are not delivered. Disabled by default.
- dot** Take a period on a line by itself, or EOF during input from a terminal as end-of-file. Disabled by default, but **dot** is set in the global start-up file (which can be suppressed with the **-n** command line option).

EDITOR=*shell-command*

The command to run when the **edit** or **~e** command is used. Default is **ed**(1).

escape=*c*

Substitute *c* for the **~** escape character. Takes effect with next message sent.

folder=*directory*

The directory for saving standard mail files. User-specified file names beginning with a plus (+) are expanded by preceding the file name with this directory name to obtain the real file name. If *directory* does not start with a slash (/), **\$HOME** is prepended to it. There is no default for the **folder** variable. See also **outfolder** below.

header Enable printing of the header summary when entering **mailx**. Enabled by default.

hold Preserve all messages that are read in the **mailbox** instead of putting them in the standard **mbox** save file. Disabled by default.

ignore Ignore interrupts while entering messages. Handy for noisy dial-up lines. Disabled by default.

ignoreeof

Ignore end-of-file during message input. Input must be terminated by a period (.) on a line by itself or by the **~.** command. See also **dot** above. Disabled by default.

indentprefix=*string*

When **indentprefix** is set, *string* is used to mark indented lines from messages included with **~m**. The default is a TAB character.

keep When the **mailbox** is empty, truncate it to zero length instead of removing it. Disabled by default.

iprompt=*string*

The specified prompt string is displayed before each line on input is requested when sending a message.

keepsave

Keep messages that have been saved in other files in the **mailbox** instead of deleting them. Disabled by default.

LISTER=*shell-command*

The command (and options) to use when listing the contents of the **folder** directory. The default is **ls**(1).

MAILX_HEAD=*string*

The specified string is included at the beginning of the body of each message that is sent.

MAILX_TAIL=*string*

The specified string is included at the end of the body of each message that is sent.

makeremote

When replying to all recipients of a message, if an address does not include a machine name, it is assumed to be relative to the sender of the message. Normally not needed when dealing with hosts that support RFC822.

MBOX=filename

The name of the file to save messages which have been read. The **exit** command overrides this function, as does saving the message explicitly in another file. Default is **\$HOME/mbox**.

metoo If your login appears as a recipient, do not delete it from the list. Disabled by default.

mustbang

Force all mail addresses to be in bang format.

onehop

When responding to a message that was originally sent to several recipients, the other recipient addresses are normally forced to be relative to the originating author's machine for the response. This flag disables alteration of the recipients' addresses, improving efficiency in a network where all machines can send directly to all other machines (that is, one hop away). Disabled by default.

outfolder

Locate the files used to record outgoing messages in the directory specified by the **folder** variable unless the path name is absolute. Disabled by default. See **folder** above and the **Save**, **Copy**, **followup**, and **Followup** commands.

page Used with the **pipe** command to insert a form feed after each message sent through the pipe. Disabled by default.

PAGER=shell-command

The command to use as a filter for paginating output. This can also be used to specify the options to be used. Default is **pg(1)**, or if the **bsdcompat** variable is set, the default is **more(1)**. See **Environment Variables**.

postmark

Your "real name" to be included in the From line of messages you send. By default this is derived from the comment field in your passwd file entry.

prompt=string

Set the *command mode* prompt to *string*. Default is "? ", unless the **bsdcompat** variable is set, then the default is "&".

quiet Refrain from printing the opening message and version when entering **mailx**. Disabled by default.

record=filename

Record all outgoing mail in *filename*. Disabled by default. See also **outfolder** above. If you have the **record** and **outfolder** variables set but the **folder** variable not set, messages are saved in *+filename* instead of *filename*.

replyall

Reverse the effect of the **reply** and **Reply** commands.

save Enable saving of messages in **dead-letter** on interrupt or delivery error. See **DEAD** for a description of this file. Enabled by default.

screen=number

Sets the number of lines in a screen-full of headers for the **hheaders** command. *number* must be a positive number.

The default is set according to baud rate or window size. With a baud rate less than 1200, *number* defaults to 5, if baud rate is exactly 1200, it defaults to 10. If you are in a window, *number* defaults to the default window size minus 4. Otherwise, the default is 20.

sendmail=shell-command

Alternate command for delivering messages. Note: in addition to the expected list of recipients, **mail** also passes the **-i** and **-m**, flags to the command. Since these flags are not appropriate to other commands, you may have to use a shell script that strips them from the arguments list before invoking the desired command. Default is **/usr/bin/rmail**.

sendwait

Wait for background mailer to finish before returning. Disabled by default.

SHELL=shell-command

The name of a preferred command interpreter. Default is **sh(1)**.

showname

Causes the message header display to show the sender's real name (if known) rather than their mail address. Disabled by default, but **showname** is set in the **/etc/mail/mailx.rc** system start-up file for **mailx**.

showto

When displaying the header summary and the message is from you, print the recipient's name instead of the author's name.

sign=string

The variable inserted into the text of a message when the **~a** (autograph) command is given. No default (see also **~i** in **Tilde Escapes**).

Sign=string

The variable inserted into the text of a message when the **~A** command is given. No default (see also **~i** in **Tilde Escapes**).

toplines=number

The number of lines of header to print with the **top** command. Default is 5.

verbose

Invoke **sendmail(1M)** with the **-v** flag.

translate

The name of a program to translate mail addresses. The program receives mail addresses as arguments. The program produces, on the standard output, lines containing the following data, in this order:

- the postmark for the sender (see the postmark variable)
- translated mail addresses, one per line, corresponding to the program's arguments. Each translated address will replace the corresponding address in the mail message being sent.
- a line containing only "y" or "n". if the line contains "y" the user will be asked to confirm that the message should be sent.

The translate program will be invoked for each mail message to be sent. If the program exits with a non-zero exit status, or fails to produce enough output, the message is not sent.

VISUAL=shell-command

The name of a preferred screen editor. Default is **vi(1)**.

FILES

\$HOME/.mailrc	personal start-up file
\$HOME/mbox	secondary storage file
/etc/mail/mailx.rc	optional global start-up file for mailx only
/etc/mail/Mail.rc	BSD compatibility system-wide start-up file for /usr/ucb/mail and /usr/ucb/Mail
/tmp/R[emqsx]*	temporary files
/usr/share/lib/mailx/mailx.help*	help message files
/var/mail/*	post office directory

SEE ALSO

biff(1B), **echo(1)**, **ed(1)**, **ex(1)**, **fmt(1)**, **ls(1)**, **mail(1)**, **more(1)**, **newaliases(1)**, **pg(1)**, **sh(1)**, **vacation(1)**, **vi(1)**, **uucp(1C)**, **sendmail(1M)**, **aliases(4)**

NOTES

Where *shell-command* is shown as valid, arguments are not always allowed. Experimentation is recommended.

Internal variables imported from the execution environment cannot be **unset**.

The full internet addressing is not fully supported by **mailx**. The new standards need some time to settle down.

Attempts to send a message having a line consisting only of a "." are treated as the end of the message by **mail(1)** (the standard mail delivery program).

Replies do not always generate correct return addresses. Try resending the errant reply with **onehop** set.

mailx does not lock your record file. So, if you use a record file and send two or more messages simultaneously, lines from the messages may be interleaved in the record file.

The format for the **alias** command is a space-separated list of recipients, while the format for an alias in either the **.forward** or **/etc/aliases** is a comma-separated list.

NAME	make – maintain, update, and regenerate related programs and files
SYNOPSIS	<code>/usr/ccs/bin/make [-f <i>makefile</i>] ... [-d] [-dd] [-D] [-DD] [-e] [-i] [-k] [-n] [-p] [-P] [-q] [-r] [-s] [-S] [-t] [<i>target</i> ...] [<i>macro=value</i> ...]</code>
DESCRIPTION	<p>make executes a list of shell commands associated with each <i>target</i>, typically to create or update a file of the same name. <i>makefile</i> contains entries that describe how to bring a target up to date with respect to those on which it depends, which are called <i>dependencies</i>. Since each dependency is a target, it may have dependencies of its own. Targets, dependencies, and sub-dependencies comprise a tree structure that make traces when deciding whether or not to rebuild a <i>target</i>.</p> <p>make recursively checks each <i>target</i> against its dependencies, beginning with the first target entry in <i>makefile</i> if no <i>target</i> argument is supplied on the command line. If, after processing all of its dependencies, a target file is found either to be missing, or to be older than any of its dependencies, make rebuilds it. Optionally with this version of make, a target can be treated as out-of-date when the commands used to generate it have changed since the last time the target was built.</p> <p>To build a given target, make executes the list of commands, called a <i>rule</i>. This rule may be listed explicitly in the target's makefile entry, or it may be supplied implicitly by make.</p> <p>When no <i>makefile</i> is specified with a <code>-f</code> option:</p> <ul style="list-style-type: none"> • If there is a file named makefile in the working directory, make uses that file. If, however, there is an SCCS history file (SCCS/s.makefile) which is newer, make attempts to retrieve and use the most recent version. • In the absence of the above file(s), if a file named Makefile is present in the working directory, make attempts to use it. If there is an SCCS history file (SCCS/s.Makefile) that is newer, make attempts to retrieve and use the most recent version. <p>If no <i>target</i> is specified on the command line, make uses the first target defined in <i>makefile</i>. If a <i>target</i> has no makefile entry, or if its entry has no rule, make attempts to derive a rule by each of the following methods, in turn, until a suitable rule is found. (Each method is described under USAGE below.)</p> <ul style="list-style-type: none"> • Pattern matching rules. • Implicit rules, read in from a user-supplied makefile. • Standard implicit rules (also known as suffix rules), typically read in from the file <code></usr/share/lib/make/make.rules></code>. • SCCS retrieval. make retrieves the most recent version from the SCCS history file (if any). See the description of the .SCCS_GET: special-function target for details.

- The rule from the `.DEFAULT:` target entry, if there is such an entry in the makefile.

If there is no makefile entry for a *target*, if no rule can be derived for building it, and if no file by that name is present, **make** issues an error message and halts.

OPTIONS

`-f makefile`

Use the description file *makefile*. A `-` as the *makefile* argument denotes the standard input. The contents of *makefile*, when present, override the standard set of implicit rules and predefined macros. When more than one `-f makefile` argument pair appears, **make** uses the concatenation of those files, in order of appearance.

- `-d` Display the reasons why **make** chooses to rebuild a target; **make** displays any and all dependencies that are newer. In addition, **make** displays options read in from the `MAKEFLAGS` environment variable.
- `-dd` Display the dependency check and processing in vast detail.
- `-D` Display the text of the makefiles read in.
- `-DD` Display the text of the makefiles, **make.rules** file, the state file, and all hidden-dependency reports.
- `-e` Environment variables override assignments within makefiles.
- `-i` Ignore error codes returned by commands. Equivalent to the special-function target `.IGNORE:`.
- `-k` When a nonzero error status is returned by a rule, or when **make** cannot find a rule, abandon work on the current target, but continue with other dependency branches that do not depend on it.
- `-n` No execution mode. Print commands, but do not execute them. Even lines beginning with an `@` are printed. However, if a command line contains a reference to the `$(MAKE)` macro, that line is always executed (see the discussion of `MAKEFLAGS` in **Reading Makefiles and the Environment**).
- `-p` Print out the complete set of macro definitions and target descriptions.
- `-P` Merely report dependencies, rather than building them.
- `-q` Question mode. **make** returns a zero or nonzero status code depending on whether or not the target file is up to date.
- `-r` Do not read in the default makefile `</usr/share/lib/make/make.rules>`.
- `-s` Silent mode. Do not print command lines before executing them. Equivalent to the special-function target `.SILENT:`.
- `-S` Undo the effect of the `-k` option. Stop processing when a non-zero exit status is returned by a command.
- `-t` Touch the target files (bringing them up to date) rather than performing their rules. *This can be dangerous when files are maintained by more than one person.* When the `.KEEP_STATE:` target appears in the makefile, this option updates the state file just as if the rules had been performed.

macro=value

Macro definition. This definition overrides any regular definition for the specified macro within the makefile itself, or in the environment. However, this definition can still be overridden by conditional macro assignments.

USAGE

Refer to **make** in *Programming Utilities Guide* for tutorial information.

Reading Makefiles and the Environment

When **make** first starts, it reads the **MAKEFLAGS** environment variable to obtain any the following options specified present in its value: **-d**, **-D**, **-e**, **-i**, **-k**, **-l**, **-n**, **-p**, **-q**, **-r**, **-s**, **-S**, or **-t**. (Within the **MAKEFLAGS** value, the leading ‘**-**’ character for the option string is omitted.) **make** then reads the command line for additional options, which also take effect.

Next, **make** reads in a default makefile that typically contains predefined macro definitions, target entries for implicit rules, and additional rules, such as the rule for retrieving SCCS files. If present, **make** uses the file **make.rules** in the current directory; otherwise it reads the file `</usr/share/lib/make/make.rules>`, which contains the standard definitions and rules.

Use the directive:

```
include </usr/share/lib/make/make.rules>
```

in your local **make.rules** file to include them.

Next, **make** imports variables from the environment (unless the **-e** option is in effect), and treats them as defined macros. Because **make** uses the most recent definition it encounters, a macro definition in the makefile normally overrides an environment variable of the same name. When **-e** is in effect, however, environment variables are read in *after* all makefiles have been read. In that case, the environment variables take precedence over definitions in the makefile.

Next, **make** reads the state file, **.make.state** in the local directory if it exists, and then any makefiles you specify with **-f**, or one of **makefile** or **Makefile** as described above.

Next, (after reading the environment if **-e** is in effect), **make** reads in any macro definitions supplied as command line arguments. These override macro definitions in the makefile and the environment both, but only for the **make** command itself.

make exports environment variables, using the most recently defined value. Macro definitions supplied on the command line are not normally exported, unless the macro is also an environment variable.

make does not export macros defined in the makefile. If an environment variable is set, and a macro with the same name is defined on the command line, **make** exports its value as defined on the command line. Unless **-e** is in effect, macro definitions within the makefile take precedence over those imported from the environment.

The macros **MAKEFLAGS**, **MAKE**, **SHELL**, **HOST_ARCH**, **HOST_MACH**, and **TARGET_MACH** are special cases. See **Special-Purpose Macros**, below, for details.

Makefile Target Entries

A target entry has the following format:

```
target... [: | ::] [dependency] ... [; command] ...
      [command]
...

```

The first line contains the name of a target, or a space-separated list of target names, terminated with a colon or double colon. If a list of targets is given, this is equivalent to having a separate entry of the same form for each target. The colon(s) may be followed by a *dependency*, or a dependency list. **make** checks this list before building the target. The dependency list may be terminated with a semicolon (;), which in turn can be followed by a single Bourne shell command. Subsequent lines in the target entry begin with a TAB, and contain Bourne shell commands. These commands comprise the rule for building the target.

Shell commands may be continued across input lines by escaping the NEWLINE with a backslash (\). The continuing line must also start with a TAB.

To rebuild a target, **make** expands macros, strips off initial TAB characters and either executes the command directly (if it contains no shell metacharacters), or passes each command line to a Bourne shell for execution.

The first line that does not begin with a TAB or # begins another target or macro definition.

Special Characters
Global

Start a comment. The comment ends at the next NEWLINE. If the # follows the TAB in a command line, that line is passed to the shell (which also treats # as the start of a comment).

include filename

If the word **include** appears as the first seven letters of a line and is followed by a SPACE or TAB, the string that follows is taken as a filename to interpolate at that line. **include** files can be nested to a depth of no more than about 16. If *filename* is a macro reference, it is expanded.

Targets and Dependencies

: Target list terminator. Words following the colon are added to the dependency list for the target or targets. If a target is named in more than one colon-terminated target entry, the dependencies for all its entries are added to form that target's complete dependency list.

:: Target terminator for alternate dependencies. When used in place of a ':' the double-colon allows a target to be checked and updated with respect to alternate dependency lists. When the target is out-of-date with respect to dependencies listed in the first alternate, it is built according to the rule for that entry. When out-of-date with respect to dependencies in another alternate, it is built according to the rule in that other entry. Implicit rules do not apply to double-colon targets; you must supply a rule for each entry. If no dependencies are specified, the rule is always performed.

	<i>target</i> [+ <i>target...</i>] :	Target group. The rule in the target entry builds all the indicated targets as a group. It is normally performed only once per make run, but is checked for command dependencies every time a target in the group is encountered in the dependency scan.
	%	Pattern matching wild card metacharacter. Like the '*' shell wild card, '%' matches any string of zero or more characters in a target name or dependency, in the target portion of a conditional macro definition, or within a pattern replacement macro reference. Note: only one '%' can appear in a target, dependency-name, or pattern-replacement macro reference.
	<i>./pathname</i>	make ignores the leading './' characters from targets with names given as pathnames relative to "dot," the working directory.
Macros	=	Macro definition. The word to the left of this character is the macro name; words to the right comprise its value. Leading and trailing white space characters are stripped from the value. A word break following the = is implied.
	\$	Macro reference. The following character, or the parenthesized or bracketed string, is interpreted as a macro reference: make expands the reference (including the \$) by replacing it with the macro's value.
	()	
	{ }	Macro-reference name delimiters. A parenthesized or bracketed word appended to a \$ is taken as the name of the macro being referred to. Without the delimiters, make recognizes only the first character as the macro name.
	\$\$	A reference to the dollar-sign macro, the value of which is the character '\$'. Used to pass variable expressions beginning with \$ to the shell, to refer to environment variables which are expanded by the shell, or to delay processing of dynamic macros within the dependency list of a target, until that target is actually processed.
	\\$	Escaped dollar-sign character. Interpreted as a literal dollar sign within a rule.
	+=	When used in place of '=', appends a string to a macro definition (must be surrounded by white space, unlike '=').
	:=	Conditional macro assignment. When preceded by a list of targets with explicit target entries, the macro definition that follows takes effect when processing only those targets, and their dependencies.
	:sh =	Define the value of a macro to be the output of a command (see Command Substitutions , below).
	:sh	In a macro reference, execute the command stored in the macro, and replace the reference with the output of that command (see Command Substitutions).

- Rules*
- **make** ignores any nonzero error code returned by a command line for which the first non-TAB character is a '-'. This character is not passed to the shell as part of the command line. **make** normally terminates when a command returns nonzero status, unless the **-i** or **-k** options, or the **.IGNORE:** special-function target is in effect.
 - @ If the first non-TAB character is a @, **make** does not print the command line before executing it. This character is not passed to the shell.
 - ? Escape command-dependency checking. Command lines starting with this character are not subject to command dependency checking.
 - ! Force command-dependency checking. Command-dependency checking is applied to command lines for which it would otherwise be suppressed. This checking is normally suppressed for lines that contain references to the '?' dynamic macro (for example, '\$?').

When any combination of '-', '@', '?', or '!' appear as the first characters after the TAB, all that are present apply. None are passed to the shell.

**Special-Function
Targets**

When incorporated in a makefile, the following target names perform special-functions:

.DEFAULT:

If it has an entry in the makefile, the rule for this target is used to process a target when there is no other entry for it, no rule for building it, and no SCCS history file from which to retrieve a current version. **make** ignores any dependencies for this target.

.DONE: If defined in the makefile, **make** processes this target and its dependencies after all other targets are built. This target is also performed when **make** halts with an error, unless the **.FAILED** target is defined.

.FAILED:

This target, along with its dependencies, is performed instead of **.DONE** when defined in the makefile and **make** halts with an error.

.IGNORE:

Ignore errors. When this target appears in the makefile, **make** ignores non-zero error codes returned from commands.

.INIT: If defined in the makefile, this target and its dependencies are built before any other targets are processed.

.KEEP_STATE:

If this target appears in the makefile, **make** updates the state file, **.make.state**, in the current directory. This target also activates command dependencies, and hidden dependency checks. If either the **.KEEP_STATE:** target appears in the makefile, or the environment variable **KEEP_STATE** is set ("setenv KEEP_STATE"), **make** will rebuild everything in order to collect dependency information, even if all the targets were up to date due to previous **make** runs. (See also the **ENVIRONMENT** section.)

.MAKE_VERSION:

A target-entry of the form:

.MAKE_VERSION: VERSION-number

enables version checking. If the version of **make** differs from the version indicated, **make** issues a warning message.

.NO_PARALLEL:

Currently, this target has no effect, it is, however, reserved for future use.

.PARALLEL:

Currently of no effect, but reserved for future use.

.PRECIOUS:

List of files not to delete. **make** does not remove any of the files listed as dependencies for this target when interrupted. **make** normally removes the current target when it receives an interrupt.

.SCCS_GET:

This target contains the rule for retrieving the current version of an SCCS file from its history file. To suppress automatic retrieval, add an entry for this target with an empty rule to your makefile.

.SILENT:

Run silently. When this target appears in the makefile, **make** does not echo commands before executing them.

.SUFFIXES:

The suffixes list for selecting implicit rules (see **The Suffixes List**).

.WAIT: Currently of no effect, but reserved for future use.

Clearing Special Targets

In this version of **make**, you can clear the definition of the following special targets by supplying entries for them with no dependencies and no rule:

.DEFAULT, **.SCCS_GET**, and **.SUFFIXES**

Command Dependencies

When the **.KEEP_STATE:** target appears in the makefile, **make** checks the command for building a target against the state file, **.make.state**. If the command has changed since the last **make** run, **make** rebuilds the target.

Hidden Dependencies

When the **.KEEP_STATE:** target appears in the makefile, **make** reads reports from **cpp(1)** and other compilation processors for any "hidden" files, such as **#include** files. If the target is out of date with respect to any of these files, **make** rebuilds it.

Macros

Entries of the form

macro=value

define macros. *macro* is the name of the macro, and *value*, which consists of all characters up to a comment character or unescaped NEWLINE, is the value. **make** strips both leading and trailing white space in accepting the value.

Subsequent references to the macro, of the forms $\$(name)$ or $\${name}$ are replaced by *value*. The parentheses or brackets can be omitted in a reference to a macro with a single-character name.

Macro references can contain references to other macros, in which case nested references are expanded first.

Suffix Replacement
Macro References

Substitutions within macros can be made as follows:

$\$(name:string1=string2)$

where *string1* is either a suffix, or a word to be replaced in the macro definition, and *string2* is the replacement suffix or word. Words in a macro value are separated by SPACE, TAB, and escaped NEWLINE characters.

Pattern Replacement
Macro References

Pattern matching replacements can also be applied to macros, with a reference of the form:

$\$(name:op%os= np%ns)$

where *op* is the existing (old) prefix and *os* is the existing (old) suffix, *np* and *ns* are the new prefix and new suffix, respectively, and the pattern matched by % (a string of zero or more characters), is carried forward from the value being replaced. For example:

```
PROGRAM=fabricate
DEBUG= $(PROGRAM:%=tmp/%-g)
```

sets the value of **DEBUG** to **tmp/fabricate-g**.

Note: pattern replacement macro references cannot be used in the dependency list of a pattern matching rule; the % characters are not evaluated independently. Also, any number of % metacharacters can appear after the equal-sign.

Appending to a Macro

Words can be appended to macro values as follows:

macro += word . . .

Special-Purpose
Macros

When the **MAKEFLAGS** variable is present in the environment, **make** takes options from it, in combination with options entered on the command line. **make** retains this combined value as the **MAKEFLAGS** macro, and exports it automatically to each command or shell it invokes.

Note: flags passed by way of **MAKEFLAGS** are only displayed when the **-d**, or **-dd** options are in effect.

The **MAKE** macro is another special case. It has the value **make** by default, and temporarily overrides the **-n** option for any line in which it is referred to. This allows nested invocations of **make** written as:

$\$(MAKE) . . .$

to run recursively, with the **-n** flag in effect for all commands but **make**. This lets you use '**make -n**' to test an entire hierarchy of makefiles.

For compatibility with the 4.2 BSD **make**, the **MFLAGS** macro is set from the **MAKEFLAGS** variable by prepending a '-'. **MFLAGS** is not exported automatically.

The **SHELL** macro, when set to a single-word value such as **/usr/bin/csh**, indicates the name of an alternate shell to use. The default is **/bin/sh**. Note: **make** executes commands that contain no shell metacharacters itself. Built-in commands, such as **dirs** in the C shell, are not recognized unless the command line includes a metacharacter (for instance, a semicolon). This macro is neither imported from, nor exported to the environment, regardless of **-e**. To be sure it is set properly, you must define this macro within every makefile that requires it.

The following macros are provided for use with cross-compilation:

HOST_ARCH

The machine architecture of the host system. By default, this is the output of the **arch(1B)** command prepended with '-'. Under normal circumstances, this value should never be altered by the user.

HOST_MACH

The machine architecture of the host system. By default, this is the output of the **mach(1B)**, prepended with '-'. Under normal circumstances, this value should never be altered by the user.

TARGET_ARCH

The machine architecture of the target system. By default, the output of **mach**, prepended with '-'.

Dynamic Macros

There are several dynamically maintained macros that are useful as abbreviations within rules. They are shown here as references; if you were to define them, **make** would simply override the definition.

- \$*** The basename of the current target, derived as if selected for use with an implicit rule.
- \$<** The name of a dependency file, derived as if selected for use with an implicit rule.
- \$@** The name of the current target. This is the only dynamic macro whose value is strictly determined when used in a dependency list. (In which case it takes the form '\$\$@'.)
- \$?** The list of dependencies that are newer than the target. Command-dependency checking is automatically suppressed for lines that contain this macro, just as if the command had been prefixed with a '?'. See the description of '?', under **Makefile Special Tokens**, above. You can force this check with the **!** command-line prefix.
- \$%** The name of the library member being processed. (See **Library Maintenance**, below.)

To refer to the `$$` dynamic macro within a dependency list, precede the reference with an additional `$` character (as in, `$$$`). Because **make** assigns `$$` and `$$*` as it would for implicit rules (according to the suffixes list and the directory contents), they may be unreliable when used within explicit target entries.

These macros can be modified to apply either to the filename part, or the directory part of the strings they stand for, by adding an upper case **F** or **D**, respectively (and enclosing the resulting name in parentheses or braces). Thus, `$(@D)` refers to the directory part of the string `$$`; if there is no directory part, `.` is assigned. `$(@F)` refers to the filename part.

Conditional Macro Definitions

A macro definition of the form:

```
target-list := macro = value
```

indicates that when processing any of the targets listed *and their dependencies*, *macro* is to be set to the *value* supplied. Note that if a conditional macro is referred to in a dependency list, the `$` must be delayed (use `$$` instead). Also, *target-list* may contain a `%` pattern, in which case the macro will be conditionally defined for all targets encountered that match the pattern. A pattern replacement reference can be used within the *value*.

You can temporarily append to a macro's value with a conditional definition of the form:

```
target-list := macro += value
```

Predefined Macros

make supplies the macros shown in the table that follows for compilers and their options, host architectures, and other commands. Unless these macros are read in as environment variables, their values are not exported by **make**. If you run **make** with any of these set in the environment, it is a good idea to add commentary to the makefile to indicate what value each is expected to take. If `-r` is in effect, **make** does not read the default makefile (`./make.rules` or `</usr/share/lib/make/make.rules>`) in which these macro definitions are supplied.

<i>Table of Predefined Macros</i>		
<i>Use</i>	<i>Macro</i>	<i>Default Value</i>
Library <i>Archives</i>	AR ARFLAGS	ar rv
<i>Assembler</i> <i>Commands</i>	AS ASFLAGS COMPILE.s COMPILE.S	as \$(AS) \$(ASFLAGS) \$(CC) \$(ASFLAGS) \$(CPPFLAGS) -c
<i>C Compiler</i> <i>Commands</i>	CC CFLAGS CPPFLAGS COMPILE.c LINK.c	cc \$(CC) \$(CFLAGS) \$(CPPFLAGS) -c \$(CC) \$(CFLAGS) \$(CPPFLAGS) \$(LDFLAGS)
<i>C++</i> <i>Compiler</i> <i>Commands</i>	CCC CCFLAGS CPPFLAGS COMPILE.cc LINK.cc	CC CFLAGS \$(CCC) \$(CCFLAGS) \$(CPPFLAGS) -c \$(CCC) \$(CCFLAGS) \$(CPPFLAGS) \$(LDFLAGS)
<i>FORTRAN 77</i> <i>Compiler</i> <i>Commands</i>	FC FFLAGS COMPILE.f LINK.f COMPILE.F LINK.F	f77 \$(FC) \$(FFLAGS) -c \$(FC) \$(FFLAGS) \$(LDFLAGS) \$(FC) \$(FFLAGS) \$(CPPFLAGS) -c \$(FC) \$(FFLAGS) \$(CPPFLAGS) \$(LDFLAGS)
<i>Link Editor</i> <i>Command</i>	LD LDFLAGS	ld
<i>lex</i> <i>Command</i>	LEX LFLAGS LEX.l	lex \$(LEX) \$(LFLAGS) -t
<i>lint</i> <i>Command</i>	LINT LINTFLAGS LINT.c	lint \$(LINT) \$(LINTFLAGS) \$(CPPFLAGS)
<i>Modula 2</i> <i>Commands</i>	M2C M2FLAGS MODFLAGS DEFFLAGS COMPILE.def COMPILE.mod	m2c \$(M2C) \$(M2FLAGS) \$(DEFFLAGS) \$(M2C) \$(M2FLAGS) \$(MODFLAGS)
<i>Pascal</i> <i>Compiler</i> <i>Commands</i>	PC PFLAGS COMPILE.p LINK.p	pc \$(PC) \$(PFLAGS) \$(CPPFLAGS) -c \$(PC) \$(PFLAGS) \$(CPPFLAGS) \$(LDFLAGS)

<i>Table of Predefined Macros</i>		
<i>Use</i>	<i>Macro</i>	<i>Default Value</i>
Ratfor <i>Compilation</i> <i>Commands</i>	RFLAGS COMPILE.r LINK.r	\$(FC) \$(FFLAGS) \$(RFLAGS) -c \$(FC) \$(FFLAGS) \$(RFLAGS) \$(LDFFLAGS)
<i>rm</i> Command	RM	rm -f
<i>secs</i> <i>Command</i>	SCCSFLAGS SCCSGETFLAGS	-s
<i>yacc</i> <i>Command</i>	YACC YFLAGS YACC.y	yacc \$(YACC) \$(YFLAGS)
<i>Suffixes</i> List	SUFFIXES	.o .c .c~ .cc .cc~ .y .y~ .I .I~ .s .s~ .sh .sh~ .S .S~ .ln .h .h~ .f .f~ .F .F~ .mod .mod~ .sym .def .def~ .p .p~ .r .r~ .cps .cps~ .C .C~ .Y .Y~ .L .L~

Implicit Rules

When a target has no entry in the makefile, **make** attempts to determine its class (if any) and apply the rule for that class. An implicit rule describes how to build any target of a given class, from an associated dependency file. The class of a target can be determined either by a pattern, or by a suffix; the corresponding dependency file (with the same basename) from which such a target might be built. In addition to a predefined set of implicit rules, **make** allows you to define your own, either by pattern, or by suffix.

Pattern Matching Rules

A target entry of the form:

```
tp%ts: dp%ds
         rule
```

is a pattern matching rule, in which *tp* is a target prefix, *ts* is a target suffix, *dp* is a dependency prefix, and *ds* is a dependency suffix (any of which may be null). The % stands for a basename of zero or more characters that is matched in the target, and is used to construct the name of a dependency. When **make** encounters a match in its search for an implicit rule, it uses the rule in that target entry to build the target from the dependency file. Pattern-matching implicit rules typically make use of the \$@ and \$< dynamic macros as placeholders for the target and dependency names. Other, regular dependencies may occur in the dependency list. An entry of the form:

```
tp%ts: [dependency...] dp%ds [dependency...]
         rule
```

is a valid pattern matching rule.

Suffix Rules

When no pattern matching rule applies, **make** checks the target name to see if it ends with a suffix in the known suffixes list. If so, **make** checks for any suffix rules, as well as a dependency file with same root and another recognized suffix, from which to build it.

The target entry for a suffix rule takes the form:

```
DsTs: rule
```

where *Ts* is the suffix of the target, *Ds* is the suffix of the dependency file, and *rule* is the rule for building a target in the class. Both *Ds* and *Ts* must appear in the suffixes list. (A suffix need not begin with a '.' to be recognized.)

A suffix rule with only one suffix describes how to build a target having a null (or no) suffix from a dependency file with the indicated suffix. For instance, the `.c` rule could be used to build an executable program named **file** from a C source file named '**file.c**'. If a target with a null suffix has an explicit dependency, **make** omits the search for a suffix rule.

<i>Table of Standard Implicit (Suffix) Rules</i>		
<i>Use</i>	<i>Implicit Rule Name</i>	<i>Command Line</i>
Assembly Files	.s.o	<code>\$(COMPILE.s) -o \$@ \$<</code>
	.s.a	<code>\$(COMPILE.s) -o \$% \$< \$(AR) \$(ARFLAGS) \$@ \$% \$(RM) \$%</code>
	.S.o	<code>\$(COMPILE.S) -o \$@ \$<</code>
	.S.a	<code>\$(COMPILE.S) -o \$% \$< \$(AR) \$(ARFLAGS) \$@ \$% \$(RM) \$%</code>
C Files	.c	<code>\$(LINK.c) -o \$@ \$<\$(LDLIBS)</code>
	.c.ln	<code>\$(LINT.c) \$(OUTPUT_OPTION) -i \$<</code>
	.c.o	<code>\$(COMPILE.c) \$(OUTPUT_OPTION) \$<</code>
	.c.a	<code>\$(COMPILE.c) -o \$% \$< \$(AR) \$(ARFLAGS) \$@ \$% \$(RM) \$%</code>
C++ Files	.cc	<code>\$(LINK.cc) -o \$@ \$<\$(LDLIBS)</code>
	.cc.o	<code>\$(COMPILE.cc) \$(OUTPUT_OPTION) \$<</code>
	.cc.a	<code>\$(COMPILE.cc) -o \$% \$< \$(AR) \$(ARFLAGS) \$@ \$% \$(RM) \$%</code>
FORTAN 77 Files	.f	<code>\$(LINK.f) -o \$@ \$<\$(LDLIBS)</code>
	.f.o	<code>\$(COMPILE.f) \$(OUTPUT_OPTION) \$<</code>
	.f.a	<code>\$(COMPILE.f) -o \$% \$< \$(AR) \$(ARFLAGS) \$@ \$% \$(RM) \$%</code>
	.F	<code>\$(LINK.F) -o \$@ \$<\$(LDLIBS)</code>
	.F.o	<code>\$(COMPILE.F) \$(OUTPUT_OPTION) \$<</code>
	.F.a	<code>\$(COMPILE.F) -o \$% \$< \$(AR) \$(ARFLAGS) \$@ \$% \$(RM) \$%</code>

<i>Table of Standard Implicit (Suffix) Rules</i>		
<i>Use</i>	<i>Implicit Rule Name</i>	<i>Command Line</i>
lex <i>Files</i>	.l	\$(RM) \$*.c \$(LEX.l) \$< > \$*.c \$(LINK.c) -o \$@ \$*.c \$(LDLIBS) \$(RM) \$*.c
	.lc	\$(RM) \$@ \$(LEX.l) \$< > \$@
	.ln	\$(RM) \$*.c \$(LEX.l) \$< > \$*.c \$(LINT.c) -o \$@ -i \$*.c \$(RM) \$*.c
	.lo	\$(RM) \$*.c \$(LEX.l) \$< > \$*.c \$(COMPILE.c) -o \$@ \$*.c \$(RM) \$*.c
<i>Modula 2</i> <i>Files</i>	.mod	\$(COMPILE.mod) -o \$@ -e \$@ \$<
	.mod.o	\$(COMPILE.mod) -o \$@ \$<
	.def.sym	\$(COMPILE.def) -o \$@ \$<
<i>NeWS</i>	.cps.h	cps \$*.cps
<i>Pascal</i> Files	.p	\$(LINK.p) -o \$@ \$< \$(LDLIBS)
	.p.o	\$(COMPILE.p) \$(OUTPUT_OPTION) \$<
<i>Ratfor</i> Files	.r	\$(LINK.r) -o \$@ \$< \$(LDLIBS)
	.r.o	\$(COMPILE.r) \$(OUTPUT_OPTION) \$<
	.r.a	\$(COMPILE.r) -o \$% \$< \$(AR) \$(ARFLAGS) \$@ \$% \$(RM) \$%
<i>SCCS</i> Files	.SCCS_GET	scs \$(SCCSFLAGS) get \$(SCCSGETFLAGS) \$@ -G\$@
<i>Shell</i> <i>Scripts</i>	.sh	cat \$< > \$@ chmod +x \$@
<i>yacc</i> <i>Files</i>	.y	\$(YACC.y) \$< \$(LINK.c) -o \$@ y.tab.c \$(LDLIBS) \$(RM) y.tab.c
	.y.c	\$(YACC.y) \$< mv y.tab.c \$@
	.y.ln	\$(YACC.y) \$< \$(LINT.c) -o \$@ -i y.tab.c \$(RM) y.tab.c
	.y.o	\$(YACC.y) \$< \$(COMPILE.c) -o \$@ y.tab.c \$(RM) y.tab.c

make reads in the standard set of implicit rules from the file `</usr/share/lib/make/make.rules>`, unless `-r` is in effect, or there is a `make.rules` file in the local directory that does not **include** that file.

The Suffixes List

The suffixes list is given as the list of dependencies for the `.SUFFIXES:` special-function target. The default list is contained in the `SUFFIXES` macro (See *Table of Predefined Macros* for the standard list of suffixes). You can define additional `.SUFFIXES:` targets; a `.SUFFIXES` target with no dependencies clears the list of suffixes. Order is significant within the list; **make** selects a rule that corresponds to the target's suffix and the first dependency-file suffix found in the list. To place suffixes at the head of the list, clear the list and replace it with the new suffixes, followed by the default list:

```
.SUFFIXES:
.SUFFIXES: suffixes $(SUFFIXES)
```

A tilde (~) indicates that if a dependency file with the indicated suffix (minus the ~) is under SCCS its most recent version should be retrieved, if necessary, before the target is processed.

Library Maintenance

A target name of the form:

```
lib(member . . .)
```

refers to a member, or a space-separated list of members, in an `ar(1)` library.

The dependency of the library member on the corresponding file must be given as an explicit entry in the makefile. This can be handled by a pattern matching rule of the form:

```
lib(%s): %s
```

where `.s` is the suffix of the member; this suffix is typically `.o` for object libraries.

A target name of the form

```
lib((symbol))
```

refers to the member of a randomized object library that defines the entry point named *symbol*.

Command Execution

Command lines are executed one at a time, *each by its own process or shell*. Shell commands, notably `cd`, are ineffectual across an unescaped NEWLINE in the makefile. A line is printed (after macro expansion) just before being executed. This is suppressed if it starts with a '@', if there is a `.SILENT:` entry in the makefile, or if **make** is run with the `-s` option. Although the `-n` option specifies printing without execution, lines containing the macro `$(MAKE)` are executed regardless, and lines containing the @ special character are printed. The `-t` (touch) option updates the modification date of a file without executing any rules. This can be dangerous when sources are maintained by more than one person.

make invokes the shell with the `-e` (exit-on-errors) argument. Thus, with semicolon-separated command sequences, execution of the later commands depends on the success of the former. This behavior can be overridden by starting the command line with a '-', or by writing a shell script that returns a non-zero status only as it finds appropriate.

Bourne Shell Constructs

To use the Bourne shell **if** control structure for branching, use a command line of the form:

```
if expression ; \
then command ; \
```

```

    ...; \
else command ; \
    ...; \
fi

```

Although composed of several input lines, the escaped NEWLINE characters insure that **make** treats them all as one (shell) command line.

To use the Bourne shell **for** control structure for loops, use a command line of the form:

```

for var in list ; \
    do command; \
    ...; \
done

```

To refer to a shell variable, use a double-dollar-sign (\$\$). This prevents expansion of the dollar-sign by **make**.

Command Substitutions

To incorporate the standard output of a shell command in a macro, use a definition of the form:

```
MACRO :sh =command
```

The command is executed only once, standard error output is discarded, and NEWLINE characters are replaced with SPACES. If the command has a non-zero exit status, **make** halts with an error.

To capture the output of a shell command in a macro reference, use a reference of the form:

```
$(MACRO :sh)
```

where *MACRO* is the name of a macro containing a valid Bourne shell command line. In this case, the command is executed whenever the reference is evaluated. As with shell command substitutions, the reference is replaced with the standard output of the command. If the command has a non-zero exit status, **make** halts with an error.

Signals

INT, SIGTERM, and QUIT signals received from the keyboard halt **make** and remove the target file being processed unless that target is in the dependency list for **.PRECIOUS**.

EXAMPLES

This makefile says that **pgm** depends on two files **a.o** and **b.o**, and that they in turn depend on their corresponding source files (**a.c** and **b.c**) along with a common file **incl.h**:

```

pgm: a.o b.o
    $(LINK.c) -o $@ a.o b.o
a.o: incl.h a.c
    cc -c a.c
b.o: incl.h b.c
    cc -c b.c

```

The following makefile uses implicit rules to express the same dependencies:

```
pgm: a.o b.o
    cc a.o b.o -o pgm
a.o b.o: incl.h
```

ENVIRONMENT

KEEP_STATE

This environment variable has the same effect as the **.KEEP_STATE**: special-function target. It enables command dependencies, hidden dependencies and writing of the state file.

USE_SVR4_MAKE

This environment variable causes **make** to invoke the generic System V version of **make** (**/usr/ccs/lib/sv4.make**). See **sysV-make(1)**.

FILES

/usr/ccs/make	make
makefile	
Makefile	current version(s) of make description file
SCCS/s.makefile	
SCCS/s.Makefile	SCCS history files for the above makefile(s)
make.rules	default file for user-defined targets, macros, and implicit rules
</usr/share/lib/make/make.rules>	makefile for standard implicit rules and macros (not read if make.rules is)
.make.state	state file in the local directory

SEE ALSO

ar(1), **cd(1)**, **passwd(4)**

Solaris Advanced User's Guide

Programming Utilities Guide

DIAGNOSTICS

make returns an exit status of **1** when it halts as a result of an error. Otherwise it returns an exit status of **0**.

Don't know how to make target 'target'

There is no makefile entry for *target*, and none of **make**'s implicit rules apply (there is no dependency file with a suffix in the suffixes list, or the target's suffix is not in the list).

*** target removed.

make was interrupted while building *target*. Rather than leaving a partially-completed version that is newer than its dependencies, **make** removes the file named *target*.

*** target not removed.

make was interrupted while building *target* and *target* was not present in the directory.

***** target could not be removed, reason**

make was interrupted while building *target*, which was not removed for the indicated reason.

Read of include file 'file' failed

The makefile indicated in an **include** directive was not found, or was inaccessible.

Loop detected when expanding macro value 'macro'

A reference to the macro being defined was found in the definition.

Could not write state file 'file'

You used the **.KEEP_STATE:** target, but do not have write permission on the state file.

***** Error code n**

The previous shell command returned a nonzero error code.

***** signal message**

The previous shell command was aborted due to a signal. If **'- core dumped'** appears after the message, a **core** file was created.

Conditional macro conflict encountered

Displayed only when **-d** is in effect, this message indicates that two or more parallel targets currently being processed depend on a target which is built differently for each by virtue of conditional macros. Since the target cannot simultaneously satisfy both dependency relationships, it is conflicted.

BUGS

Some commands return nonzero status inappropriately; to overcome this difficulty, prefix the offending command line in the rule with a **'-'**.

Filenames with the characters **'='**, **':'**, or **'@'**, do not work.

You cannot build **file.o** from **lib(file.o)**.

Options supplied by **MAKEFLAGS** should be reported for nested **make** commands. Use the **-d** option to find out what options the nested command picks up from **MAKEFLAGS**.

This version of **make** is incompatible in certain respects with previous versions:

- The **-d** option output is much briefer in this version. **-dd** now produces the equivalent voluminous output.
- **make** attempts to derive values for the dynamic macros **'\$*'**, **'\$<'**, and **'\$?'**, while processing explicit targets. It uses the same method as for implicit rules; in some cases this can lead either to unexpected values, or to an empty value being assigned. (Actually, this was true for earlier versions as well, even though the documentation stated otherwise.)
- **make** no longer searches for SCCS history "(s.)" files.
- Suffix replacement in macro references are now applied after the macro is expanded.

There is no guarantee that makefiles created for this version of **make** will work with earlier versions.

If there is no **make.rules** file in the current directory, and the file `</usr/share/lib/make/make.rules>` is missing, **make** stops before processing any targets. To force **make** to run anyway, create an empty **make.rules** file in the current directory.

Once a dependency is made, **make** assumes the dependency file is present for the remainder of the run. If a rule subsequently removes that file and future targets depend on its existence, unexpected errors may result.

When hidden dependency checking is in effect, the `$?` macro's value includes the names of hidden dependencies. This can lead to improper filename arguments to commands when `$?` is used in a rule.

Pattern replacement macro references cannot be used in the dependency list of a pattern matching rule.

Unlike previous versions, this version of **make** strips a leading `./` from the value of the `$@` dynamic macro.

With automatic SCCS retrieval, this version of **make** does not support tilde suffix rules.

The only dynamic macro whose value is strictly determined when used in a dependency list is `$@` (takes the form `$$@`).

make invokes the shell with the `-e` argument. This cannot be inferred from the syntax of the rule alone.

NAME	man – find and display reference manual pages
SYNOPSIS	<pre>man [-] [-adFlrt] [-M path] [-T macro-package] [-s section] name ... man [-M path] -k keyword ... man [-M path] -f filename ...</pre>
AVAILABILITY	SUNWdoc
DESCRIPTION	<p>man displays information from the reference manuals. It displays complete manual pages that you select by <i>name</i>, or one-line summaries selected either by <i>keyword</i> (-k), or by the name of an associated file (-f).</p> <p>A <i>section</i>, when given, applies to the <i>names</i> that follow it on the command line (up to the next <i>section</i>, if any). man looks in the indicated section of the manual (or in all sections, if none is specified) for those <i>names</i>; see Search Paths below for an explanation of how man conducts its search. If no manual page is located, man prints an error message.</p> <p>The reference page sources are typically located in the /usr/share/man/man* or /usr/man/man* directories, with each directory corresponding to a section of the manual. Since these directories are optionally installed, they may not reside on your host; you may have to mount /usr/share/man from a host on which they do reside. If there are preformatted, up-to-date versions in the corresponding cat* or fmt* directories, man simply displays or prints those versions. If the preformatted version of interest is out of date or missing, man reformats it prior to display and will store the preformatted version if cat? or fmt? is writable. The windex database is not updated. See catman(1M). If directories for the preformatted versions are not provided, man reformats a page whenever it is requested; it uses a temporary file to store the formatted text during display.</p> <p>If the standard output is not a terminal, or if the - flag is given, man pipes its output through cat(1); otherwise, man pipes its output through more(1) to handle paging and underlining on the screen.</p>
OPTIONS	<ul style="list-style-type: none"> -a Show all manual pages matching <i>name</i> within the MANPATH search path. Manual pages are displayed in the order found. -d Debug. Displays what a section-specifier evaluates to, method used for searching, and paths searched by man. -F Force man to search all directories specified by MANPATH or the man.cf file, rather than using the windex lookup database. This is useful if the database is not up to date. If the windex database does not exist, this option is assumed. -l List all manual pages found matching <i>name</i> within the search path. -r Reformat the manual page, but do not display it. This replaces the man - -t name combination. -t man arranges for the specified manual pages to be troffed to a suitable raster output device (see troff(1)). If both the - and -t flags are given, man updates the troffed versions of each named <i>name</i> (if necessary), but does not display them.

-M path

Specify an alternate search path for manual pages. *path* is a colon-separated list of directories that contain manual page directory subtrees. For example, if *path* is `/usr/share/man:/usr/local/man`, **man** searches for *name* in the standard location, and then `/usr/local/man`. When used with the `-k` or `-f` options, the `-M` option must appear first. Each directory in the *path* is assumed to contain subdirectories of the form `man*`, one for each section. This option overrides the **MANPATH** environment variable.

-s section ...

Specify sections of the manual for **man** to search. The directories searched for *name* is limited to those specified by *section*. *section* can be a digit (perhaps followed by one or more letters), a word (for example: local, new, old, public), or a letter. To specify multiple sections, separate each section with a comma. This option overrides the **MANPATH** environment variable and the `man.cf` file. See **Search Path** under USAGE for a complete explanation of the default search path order.

-T macro-package

Format manual pages using *macro-package* rather than the standard `-man` macros defined in `/usr/share/lib/tmac/an`.

-k keyword ...

Print out one-line summaries from the **windex** database (table of contents) that contain any of the given *keywords*. The **windex** database is created using **catman**(1M).

-f filename ...

man attempts to locate manual pages related to any of the given *filenames*. It strips the leading pathname components from each *filename*, and then prints one-line summaries containing the resulting basename or names. This option also uses the **windex** database.

USAGE
Sections

Entries in the reference manuals are organized into *sections*. A section name consists of a major section name, typically a single digit, optionally followed by a subsection name, typically one or more letters. An unadorned major section name acts as an abbreviation for the section of the same name along with all of its subsections. Each section contains descriptions apropos to a particular reference category, with subsections refining these distinctions. See the **intro** manual pages for an explanation of the classification used in this release.

Search Path

Before searching for a given *name*, **man** constructs a list of candidate directories and sections. **man** searches for *name* in the directories specified by the **MANPATH** environment variable. If this variable is not set, `/usr/share/man` is searched by default.

Within the manual page directories, **man** confines its search to the sections specified in the following order:

- *sections* specified on the command line with the `-s` option

- *sections* embedded in the **MANPATH** environment variable
- *sections* specified in the **man.cf** file for each directory specified in the **MANPATH** environment variable

If none of the above exist, **man** searches each directory in the manual page path, and displays the first matching manual page found.

The **man.cf** file has the following format:

```
MANSECTS=section[,section]...
```

Lines beginning with '#' and blank lines are considered comments, and are ignored. Each directory specified in **MANPATH** can contain a manual page configuration file, specifying the default search order for that directory.

Formatting Manual Pages

Manual pages are **troff(1)** or **nroff(1)** source files prepared with the **-man** macro package. Please refer to **man(5)** for more information.

Preprocessing Manual Pages

When formatting a manual page, **man** examines the first line to determine whether it requires special processing. If the first line is a string of the form:

```
`\" X
```

where *X* is separated from the '"' by a single SPACE and consists of any combination of characters in the following list, **man** pipes its input to **troff(1)** or **nroff(1)** through the corresponding preprocessors.

```

e      eqn(1), or neqn for nroff
r      refer(1)
t      tbl(1)
v      vgrind(1)

```

If **eqn** or **neqn** is invoked, it will automatically read the file **/usr/pub/eqnchar** (see **eqnchar(5)**). If **nroff(1)** is invoked, **col(1)** is automatically used.

Referring to Other Manual Pages

If the first line of the manual page is a reference to another manual page entry fitting the pattern:

```
.so man*/sourcefile
```

man processes the indicated file in place of the current one. The reference must be expressed as a pathname relative to the root of the manual page directory subtree.

When the second or any subsequent line starts with **.so**, **man** ignores it; **troff(1)** or **nroff(1)** processes the request in the usual manner.

ENVIRONMENT

MANPATH A colon-separated list of directories; each directory can be followed by a comma-separated list of sections. If set, its value overrides **/usr/share/man** as the default directory search path, and the **man.cf** file as the default section search path. The **-M** and **-s** flags, in turn, override these values.)

PAGER A program to use for interactively delivering **man**'s output to the screen. If not set, **'more -s'** (see **more(1)**) is used.

	TCAT	The name of the program to use to display troffed manual pages.
	TROFF	The name of the formatter to use when the -t flag is given. If not set, troff(1) is used.
FILES	/usr/share/man	root of the standard manual page directory subtree
	/usr/share/man/man?/*	unformatted manual entries
	/usr/share/man/cat?/*	nroffed manual entries
	/usr/share/man/fmt?/*	troffed manual entries
	/usr/share/man/windex	table of contents and keyword database
	/usr/share/lib/tmac/an	standard -man macro package
	/usr/share/lib/pub/eqnchar	standard definitions for eqn and neqn
	man.cf	default search order by section
SEE ALSO	apropos(1) , cat(1) , col(1) , eqn(1) , more(1) , nroff(1) , refer(1) , tbl(1) , troff(1) , vgrind(1) , whatis(1) , catman(1M) , eqnchar(5) , man(5)	
NOTES	Because troff is not 8-bit clean, man has not been made 8-bit clean. The -f and -k options use the /usr/share/man/windex database, which is created by catman(1M) .	
BUGS	The manual is supposed to be reproducible either on a phototypesetter or on an ASCII terminal. However, on a terminal some information (indicated by font changes, for instance) is lost. Some dumb terminals cannot process the vertical motions produced by the e (see eqn(1)) preprocessing flag. To prevent garbled output on these terminals, when you use e also use t , to invoke col(1) implicitly. This workaround has the disadvantage of eliminating superscripts and subscripts — even on those terminals that can display them. CTRL-Q will clear a terminal that gets confused by eqn(1) output.	

NAME	mconnect – connect to SMTP mail server socket
SYNOPSIS	mconnect [-p <i>port</i>] [-r] [<i>hostname</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	mconnect opens a connection to the mail server on a given host, so that it can be tested independently of all other mail software. If no host is given, the connection is made to the local host. Servers expect to speak the Simple Mail Transfer Protocol (SMTP) on this connection. Exit by typing the quit command. Typing EOF sends an end of file to the server. An interrupt closes the connection immediately and exits.
OPTIONS	-p <i>port</i> Specify the port number instead of the default SMTP port (number 25) as the next argument. -r “Raw” mode: disable the default line buffering and input handling. This produces an effect similar to telnet to port number 25.
FILES	/etc/mail/sendmail.hf help file for SMTP commands
SEE ALSO	sendmail(1M) Postel, Jonathan B <i>Simple Mail Transfer Protocol</i> , RFC821 August 1982, SRI Network Information Center

NAME	mcs – manipulate the comment section of an object file
SYNOPSIS	mcs [-a <i>string</i>] [-c] [-d] [-n <i>name</i>] [-p] [-V] <i>filename</i> . . .
DESCRIPTION	The mcs command is used to manipulate a section, by default the .comment section, in an ELF object file. It is used to add to, delete, print, and compress the contents of a section in an ELF object file, and only print the contents of a section in a COFF object file. mcs must be given one or more of the options described below. It applies each of the options in order to each file.
OPTIONS	<p>-a <i>string</i> Append <i>string</i> to the comment section of the ELF object files. If <i>string</i> contains embedded blanks, it must be enclosed in quotation marks.</p> <p>-c Compress the contents of the comment section of the ELF object files. All duplicate entries are removed. The ordering of the remaining entries is not disturbed.</p> <p>-d Delete the contents of the comment section from the ELF object files. The section header for the comment section is also removed.</p> <p>-n <i>name</i> Specify the name of the comment section to access if other than .comment. By default, mcs deals with the section named .comment. This option can be used to specify another section.</p> <p>-p Print the contents of the comment section on the standard output. Each section printed is tagged by the name of the file from which it was extracted, using the format <i>filename</i>[<i>member_name</i>]: for archive files; and <i>filename</i>: for other files.</p> <p>-V Print, on standard error, the version number of mcs.</p> <p>If the input file is an archive (see ar(4)), the archive is treated as a set of individual files. For example, if the -a option is specified, the string is appended to the comment section of each ELF object file in the archive; if the archive member is not an ELF object file, then it is left unchanged.</p> <p>If mcs is executed on an archive file the archive symbol table will be removed, unless only the -p option has been specified. The archive symbol table must be restored by executing the ar command with the -s option before the archive can be linked by the ld command. mcs will produce appropriate warning messages when this situation arises.</p>
EXAMPLES	<p>The following example:</p> <pre>example% mcs -p filename</pre> <p>prints filename's comment section.</p> <p>The next example:</p> <pre>example% mcs -a string filename</pre> <p>appends string to filename's comment section.</p>

FILES	/tmp/mcs* temporary files
SEE ALSO	ar(1), as(1), ld(1), tmpnam(3S), a.out(4), ar(4)
NOTES	mcs cannot add to, delete or compress the contents of a section that is contained within a segment.

NAME	mesg – permit or deny messages
SYNOPSIS	mesg [-n] [-y]
AVAILABILITY	SUNWcsu
DESCRIPTION	mesg with argument -n forbids messages via write(1) by revoking non-user write permission on the user's terminal. mesg with argument -y reinstates permission. All by itself, mesg reports the current state without changing it.
OPTIONS	-n Forbids messages via write(1) . -y Reinstates permission.
FILES	/dev/tty*
SEE ALSO	write(1)
DIAGNOSTICS	Exit status is 0 if messages are receivable, 1 if not, 2 on error.

NAME	message – puts its arguments on FMLI message line
SYNOPSIS	message [-t] [-b [num]] [-o] [-w] [string] message [-f] [-b [num]] [-o] [-w] [string] message [-p] [-b [num]] [-o] [-w] [string]
DESCRIPTION	<p>The message command puts <i>string</i> out on the FMLI message line. If there is no string, the <i>stdin</i> input to message will be used. The output of message has a duration (length of time it remains on the message line). The default duration is "transient": it or one of two other durations can be requested with the mutually-exclusive options below.</p> <p>Messages displayed with message -p will replace (change the value of) any message currently displayed or stored via use of the permanentmsg descriptor. Likewise, message -f will replace any message currently displayed or stored via use of the framemsg descriptor. If more than one message in a frame definition file is specified with the -p option, the last one specified will be the permanent duration message.</p> <p>The <i>string</i> argument should always be the last argument.</p>
OPTIONS	<p>-t Explicitly defines a message to have transient duration. Transient messages remain on the message line only until the user presses another key or a CHECKWORLD occurs. The descriptors itemmsg, fieldmsg, invalidmsg, choicemsg, the default-if-not-defined value of oninterrupt, and FMLI generated error messages (that is, from syntax errors) also output transient duration messages. Transient messages take precedence over both frame messages and permanent messages.</p> <p>-f Defines a message to have "frame" duration. Frame messages remain on the message line as long as the frame in which they are defined is current. The descriptor framemsg also outputs a frame duration message. Frame messages take precedence over permanent messages.</p> <p>-p Defines a message to have "permanent" duration. Permanent messages remain on the message line for the length of the FMLI session, unless explicitly replaced by another permanent message or temporarily superseded by a transient message or frame message. A permanent message is not affected by navigating away from, or by closing, the frame which generated the permanent message. The descriptor permanentmsg also outputs a permanent duration message.</p> <p>-b[num] Rings the terminal bell <i>num</i> times, where <i>num</i> is an integer from 1 to 10. The default value is 1. If the terminal has no bell, the screen will flash <i>num</i> times instead, if possible.</p> <p>-o Forces message to duplicate its message to <i>stdout</i>.</p> <p>-w Turns on the working indicator.</p>

EXAMPLES

When a value entered in a field is invalid, ring the bell 3 times and then display **Invalid Entry: Try again!** on the message line:

```
invalidmsg=`message -b 3 "Invalid Entry: Try again!"`
```

Display a message that tells the user what is being done:

```
done=`message EDITOR has been set in your environment` close
```

Display a message on the message line and *stdout* for each field in a form (a pseudo-"field duration" message).

```
fieldmsg=""`message -o -f "Enter a filename."`"
```

Display a blank transient message (effect is to "remove" a permanent or frame duration message).

```
done=`message ""` nop
```

SEE ALSO

sleep(1)

NOTES

If **message** is coded more than once on a single line, it may appear that only the right-most instance is interpreted and displayed. Use **sleep(1)** between uses of **message** in this case, to display multiple messages.

message -f should not be used in a stand-alone backquoted expression or with the **init** descriptor because the frame is not yet current when these are evaluated.

In cases where ``message -f "string"``` is part of a stand-alone backquoted expression, the context for evaluation of the expression is the previously current frame. The previously current frame can be the frame that issued the **open** command for the frame containing the backquoted expression, or it can be a frame given as an argument when **fml** was invoked. That is, the previously current frame is the one whose frame message will be modified.

Permanent duration messages are displayed when the user navigates to the command line.

NAME	mkdir – make directories
SYNOPSIS	mkdir [-m mode] [-p] <i>dirname</i> . . .
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>mkdir creates the named directories in mode 777 (possibly altered by umask(1)). Standard entries in a directory (for instance, the files ., for the directory itself, and .., for its parent) are made automatically. mkdir cannot create these entries by name. Creation of a directory requires write permission in the parent directory.</p> <p>The owner-ID and group-ID of the new directories are set to the process's effective user-ID and group-ID, respectively. The behavior mimics the mkdir(2) system call.</p>
OPTIONS	<p>-m mode This option allows users to specify the mode to be used for new directories. Choices for modes can be found in chmod(1).</p> <p>-p With this option, mkdir creates dirname by creating all the non-existing parent directories first.</p>
EXAMPLES	<p>The following example:</p> <pre>example% mkdir -p ltr/jd/jan</pre> <p>creates the subdirectory structure ltr/jd/jan.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of mkdir for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how mkdir behaves.</p> <p>LC_CTYPE</p> <p>Determines how mkdir handles characters. When LC_CTYPE is set to a valid value, mkdir can display and handle text and filenames containing valid characters for that locale. mkdir can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. mkdir can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p> <p>LC_MESSAGES</p> <p>Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).</p>

SEE ALSO **rm(1), sh(1), umask(1), intro(2), mkdir(2), environ(5)**

DIAGNOSTICS **mkdir** returns exit code **0** if all directories given in the command line were made successfully. Otherwise, it prints a diagnostic and returns non-zero.

NAME	mkmsgs – create message files for use by gettext
SYNOPSIS	mkmsgs [-o] [-i locale] <i>inputstrings</i> <i>msgfile</i>
AVAILABILITY	SUNWloc
DESCRIPTION	<p>The mkmsgs utility is used to create a file of text strings that can be accessed using the text retrieval tools (see gettext(1), srchtxt(1), exstr(1), and gettext(3C)). It will take as input a file of text strings for a particular geographic locale (see setlocale(3C)) and create a file of text strings in a format that can be retrieved by both gettext(1) and gettext(3C). By using the -i option, you can install the created file under the /usr/lib/locale/locale/LC_MESSAGES directory (<i>locale</i> corresponds to the language in which the text strings are written).</p> <p><i>inputstrings</i> is the name of the file that contains the original text strings. <i>msgfile</i> is the name of the output file where mkmsgs writes the strings in a format that is readable by gettext(1) and gettext(3C). The name of <i>msgfile</i> can be up to 14 characters in length, but may not contain either \0 (null) or the ASCII code for / (slash) or : (colon).</p> <p>The input file contains a set of text strings for the particular geographic locale. Text strings are separated by a newline character. Nongraphic characters must be represented as alphabetic escape sequences. Messages are transformed and copied sequentially from <i>inputstrings</i> to <i>msgfile</i>. To generate an empty message in <i>msgfile</i>, leave an empty line at the correct place in <i>inputstrings</i>.</p> <p>Strings can be changed simply by editing the file <i>inputstrings</i>. New strings must be added only at the end of the file; then a new <i>msgfile</i> file must be created and installed in the correct place. If this procedure is not followed, the retrieval function will retrieve the wrong string and software compatibility will be broken.</p>
OPTIONS	<p>-o Overwrite <i>msgfile</i>, if it exists.</p> <p>-i locale Install <i>msgfile</i> in the /usr/lib/locale/locale/LC_MESSAGES directory. Only someone who is super-user or a member of group bin can create or overwrite files in this directory. Directories under /usr/lib/locale will be created if they do not exist.</p>
EXAMPLES	<p>The following example shows an input message source file C.str:</p> <pre> File %s:\t cannot be opened\n %s: Bad directory\n . . . write error\n . . </pre>

The following command uses the input strings from **C.str** to create text strings in the appropriate format in the file **UX** in the current directory:

example% mkmsgs C.str UX

The following command uses the input strings from **FR.str** to create text strings in the appropriate format in the file **UX** in the directory **/usr/lib/locale/french/LC_MESSAGES/UX**:

example% mkmsgs -i french FR.str UX

These text strings would be accessed if you had set the environment variable **LC_MESSAGES=french** and then invoked one of the text retrieval tools listed at the beginning of the **DESCRIPTION** section.

FILES **/usr/lib/locale/locale/LC_MESSAGES/*** message files created by **mkmsgs**

SEE ALSO **exstr(1), gettext(1), srchtxt(1), gettext(3C), setlocale(3C)**

NAME	mkstr – create an error message file by massaging C source files
SYNOPSIS	<code>/usr/ucb/mkstr [-] messagefile prefix filename...</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>mkstr creates files of error messages. You can use mkstr to make programs with large numbers of error diagnostics much smaller, and to reduce system overhead in running the program — as the error messages do not have to be constantly swapped in and out.</p> <p>mkstr processes each of the specified <i>filenames</i>, placing a massaged version of the input file in a file with a name consisting of the specified <i>prefix</i> and the original source file name. A typical example of using mkstr would be:</p> <pre style="margin-left: 40px;">mkstr pistrings processed *.c</pre> <p>This command would cause all the error messages from the C source files in the current directory to be placed in the file pistrings and processed copies of the source for these files to be placed in files whose names are prefixed with <i>processed</i>.</p> <p>To process the error messages in the source to the message file, mkstr keys on the string 'error(' in the input stream. Each time it occurs, the C string starting at the '(' is placed in the message file followed by a null character and a NEWLINE character; the null character terminates the message so it can be easily used when retrieved, the NEWLINE character makes it possible to sensibly cat the error message file to see its contents. The massaged copy of the input file then contains a lseek pointer into the file which can be used to retrieve the message, that is:</p>

```

char efilename[] = "/usr/lib/pi_strings";
int efil = -1;

error(a1, a2, a3, a4)
{
    char
    buf[256];
    if (efil < 0) {
        efil = open(efilename, 0);
        if (efil < 0) {
oops:
            perror (efilename);
            exit (1);
        }
    }
}

```

```
        }  
    }  
    if (lseek(efil, (long) a1, 0) || read(efil, buf, 256) <= 0)  
        goto oops;  
    printf(buf, a2, a3, a4);  
}
```

OPTIONS – Place error messages at the end of the specified message file for recompiling part of a large **mkstred** program.

SEE ALSO **xstr(1)**

NAME	more, page – browse or page through a text file
SYNOPSIS	more [-cdfirsuw] [-lines] [+linenumber] [+/pattern] [filename ...] page [-cdfirsuw] [-lines] [+linenumber] [+/pattern] [filename ...]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>more is a filter that displays the contents of a text file on the terminal, one screenful at a time. It normally pauses after each screenful, and prints --More-- at the bottom of the screen. more provides a two-line overlap between screens for continuity. If more is reading from a file rather than a pipe, the percentage of characters displayed so far is also shown.</p> <p>more scrolls up to display one more line in response to a RETURN character; it displays another screenful in response to a SPACE character. Other commands are listed below.</p> <p>page clears the screen before displaying the next screenful of text; it only provides a one-line overlap between screens.</p> <p>more sets the terminal to <i>noecho</i> mode, so that the output can be continuous. Commands that you type do not normally show up on your terminal, except for the / and ! commands.</p> <p>If the standard output is not a terminal, more acts just like cat(1), except that a header is printed before each file in a series.</p>
OPTIONS	<p>The following options are available with more:</p> <ul style="list-style-type: none"> -c Clear before displaying. Redrawing the screen instead of scrolling for faster displays. This option is ignored if the terminal does not have the ability to clear to the end of a line. -d Display error messages rather than ringing the terminal bell if an unrecognized command is used. This is helpful for inexperienced users. -f Do not fold long lines. This is useful when lines contain nonprinting characters or escape sequences, such as those generated when nroff(1) output is piped through ul(1). -l Do not treat FORMFEED characters (CTRL-L) as “page breaks.” If -l is not used, more pauses to accept commands after any line containing a ^L character (CTRL-L). Also, if a file begins with a FORMFEED, the screen is cleared before the file is printed. -r Normally, more ignores control characters that it does not interpret in some way. The -r option causes these to be displayed as ^C where C stands for any such control character. -s Squeeze. Replace multiple blank lines with a single blank line. This is helpful when viewing nroff(1) output, on the screen.

- u** Suppress generation of underlining escape sequences. Normally, **more** handles underlining, such as that produced by **nroff(1)**, in a manner appropriate to the terminal. If the terminal can perform underlining or has a stand-out mode, **more** supplies appropriate escape sequences as called for in the text file.
- w** Normally, **more** exits when it comes to the end of its input. With **-w**, however, **more** prompts and waits for any key to be struck before exiting.
- lines** Display the indicated number of *lines* in each screenful, rather than the default (the number of lines in the terminal screen less two).
- +linenumber**
Start up at *linenumber*.
- +/*pattern*** Start up two lines above the line containing the regular expression *pattern*.
Note: Unlike editors, this construct should *not* end with a *'*. If it does, then the trailing slash is taken as a character in the search pattern.

USAGE**Environment**

more uses the terminal's **/usr/share/lib/terminfo** entry to determine its display characteristics, and looks in the environment variable **MORE** for any preset options. For instance, to page through files using the **-c** mode by default, set the value of this variable to **-c**. (Normally, the command sequence to set up this environment variable is placed in the **.login** or **.profile** file).

Commands

The commands take effect immediately; it is not necessary to type a carriage return. Up to the time when the command character itself is given, the user may type the line kill character to cancel the numerical argument being formed. In addition, the user may type the erase character to redisplay the **'--More--(xx%)'** message.

In the following commands, *i* is a numerical argument (**1** by default).

- i*SPACE** Display another screenful, or *i* more lines if *i* is specified.
- i*RETURN** Display another line, or *i* more lines, if specified.
- i*D** (CTRL-D) Display (scroll down) 11 more lines. If *i* is given, the scroll size is set to *i*.
- i*d** Same as **^D**.
- i*z** Same as **SPACE**, except that *i*, if present, becomes the new default number of lines per screenful.
- i*s** Skip *i* lines and then print a screenful.
- i*f** Skip *i* screenfuls and then print a screenful.
- i*B** (CTRL-B) Skip back *i* screenfuls and then print a screenful.
- b*** Same as **^B** (CTRL-B).
- q**
- Q** Exit from **more**.
- =** Display the current line number.

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

FILES **/usr/lib/more.help** help file
 /usr/share/lib/terminfo/* terminal data base

SEE ALSO **cat(1), csh(1), man(1), script(1), sh(1), environ(4), environ(5)**

NOTES Skipping backwards is too slow on large files.

NAME	msgfmt – create a message object from a message file
SYNOPSIS	msgfmt [-v] [-o <i>output-file</i>] <i>filename.po</i> ...
AVAILABILITY	SUNWloc
DESCRIPTION	<p>msgfmt creates message object files from portable object files (<i>filename.po</i>), without changing the portable object files.</p> <p>The .po file contains messages displayed to users by system commands or by application programs. .po files can be edited, and the messages in them can be rewritten in any language supported by the system.</p> <p>The xgettext(1) command can be used to create .po files from script or programs.</p>
Portable Object Files	<p>Formats for all .po files are the same. Each .po file contains one or more lines, with each line containing either a comment or a statement. Comments start the line with a hash mark (#) and end with the newline character. All comments are ignored. The format of a statement is:</p> <p style="padding-left: 40px;"><i>directive value</i></p> <p>Each directive starts at the beginning of the line and is separated from <i>value</i> by white space (such as one or more space or tab characters). <i>value</i> consists of one or more quoted strings separated by white space. Use any of the following types of directives:</p> <p style="padding-left: 40px;">domain <i>domainname</i></p> <p style="padding-left: 40px;">msgid <i>message_identifier</i></p> <p style="padding-left: 40px;">msgstr <i>message_string</i></p> <p>The behavior of the domain directive is affected by the options used. See OPTIONS for the behavior when the -o option is specified. If the -o option is not specified, the behavior of the domain directive is as follows:</p> <ul style="list-style-type: none"> • All <i>msgids</i> from the beginning of each .po file to the first domain directive are put into a default message object file, messages.mo. • When msgfmt encounters a domain <i>domainname</i> directive in the .po file, all following <i>msgids</i> until the next domain directive are put into the message object file <i>domainname.mo</i>. • Duplicate <i>msgids</i> are defined in the scope of each domain. That is, a <i>msgid</i> is considered a duplicate only if the identical <i>msgid</i> exists in the same domain. • All duplicate <i>msgids</i> are ignored. <p>The msgid directive specifies the value of a message identifier associated with the directive that follows it. The <i>message_identifier</i> string identifies a target string to be used at retrieval time. Each statement containing a msgid directive must be followed by a statement containing a msgstr directive.</p> <p>The msgstr directive specifies the target string associated with the <i>message_identifier</i> string declared in the immediately preceding msgid directive.</p>

Message strings can contain the escape sequences `\n` for newline, `\t` for tab, `\v` for vertical tab, `\b` for backspace, `\r` for carriage return, `\f` for formfeed, `\\` for backslash, `\"` for double quote, `\ddd` for octal bit pattern, and `\xDD` for hexadecimal bit pattern.

OPTIONS

- `-v` Verbose. List duplicate message identifiers. Message strings are not redefined.
- `-o output-file` Specify output file name as *output-file*. All **domain** directives and duplicate *msgid*s in the **.po** file are ignored.

EXAMPLES

In this example **module1.po** and **module2.po** are portable message objects files.

```
example% cat module1.po
# default domain "messages.mo"
msgid "msg 1"
msgstr "msg 1 translation"
#
domain "help_domain"
msgid "help 2"
msgstr "help 2 translation"
#
domain "error_domain"
msgid "error 3"
msgstr "error 3 translation"
```

```
example% cat module2.po
# default domain "messages.mo"
msgid "mesg 4"
msgstr "mesg 4 translation"
#
domain "error_domain"
msgid "error 5"
msgstr "error 5 translation"
#
domain "window_domain"
msgid "window 6"
msgstr "window 6 translation"
```

The following command will produce the output files, **messages.mo**, **help_domain.mo**, and **error_domain.mo**.

```
example% msgfmt module1.po
```

The following command will produce the output files, **messages.mo**, **help_domain.mo**, **error_domain.mo**, and **window_domain.mo**.

```
example% msgfmt module1.po module2.po
```

The following example will produce the output file **hello.mo**.

```
example% msgfmt -o hello.mo module1.po module2.po
```

Install message object files in `/usr/lib/locale/locale/LC_MESSAGES/domain.mo` where *locale* is the message locale as set by `setlocale(3C)`, and *domain* is text domain as set by `textdomain()`. The `/usr/lib/locale` portion can optionally be changed by calling `bindtextdomain()`. See `gettext(3I)`.

SEE ALSO `xgettext(1)`, `gettext(3I)`

NOTES Neither `msgfmt` nor any `gettext(3I)` routine imposes a limit on the total length of a message. However, each line in the `*.po` file is limited to `MAX_INPUT` (512) bytes.

Installing message catalogs under the C locale is pointless, since they are ignored for the sake of efficiency.

NAME	mt – magnetic tape control	
SYNOPSIS	mt [-f <i>tapename</i>] <i>command...</i> [<i>count</i>]	
AVAILABILITY	SUNWcsu	
DESCRIPTION	<p>mt sends commands to a magnetic tape drive. If <i>tapename</i> is not specified, the environment variable TAPE is used. If TAPE does not exist, mt uses the device /dev/rmt/0. <i>tapename</i> refers to a raw tape device. By default, mt performs the requested operation once; multiple operations may be performed by specifying <i>count</i>.</p> <p>The available <i>commands</i> are listed below. Only as many characters as are required to uniquely identify a <i>command</i> need be specified.</p> <p>mt returns a 0 exit status when the operation(s) were successful, 1 if the command was unrecognized or if mt was unable to open the specified tape drive, and 2 if an operation failed.</p>	
mt Commands	eof, weof	Write <i>count</i> EOF marks at the current position on the tape.
	fsf	Forward space over <i>count</i> EOF marks. The tape is positioned on the first block of the file.
	fsr	Forward space <i>count</i> records.
	bsf	Back space over <i>count</i> EOF marks. The tape is positioned on the beginning-of-tape side of the EOF mark.
	bsr	Back space <i>count</i> records.
	nbsf	Back space <i>count</i> files. The tape is positioned on the first block of the file. This is equivalent to <i>count+1</i> bsf 's followed by one fsf .
	asf	Absolute space to <i>count</i> file number. This is equivalent to a rewind followed by a fsf <i>count</i> .
	For the following commands, <i>count</i> is ignored:	
	eom	Space to the end of recorded media on the tape. This is useful for appending files onto previously written tapes.
	rewind	Rewind the tape.
	offline, rewoffl	Rewind the tape and, if appropriate, take the drive unit off-line by unloading the tape.
	status	Print status information about the tape unit.
	retension	Rewind the cartridge tape completely, then wind it forward to the end of the reel and back to beginning-of-tape to smooth out tape tension.
	erase	Erase the entire tape.

FILES /dev/rmt/* magnetic tape interface
 /dev/rmt/*b
 /dev/rmt/*bn
 /dev/rmt/*c
 /dev/rmt/*cb
 /dev/rmt/*cbn
 /dev/rmt/*cn
 /dev/rmt/*h
 /dev/rmt/*hb
 /dev/rmt/*hbn
 /dev/rmt/*hn
 /dev/rmt/*l
 /dev/rmt/*lb
 /dev/rmt/*lbn
 /dev/rmt/*ln
 /dev/rmt/*m
 /dev/rmt/*mb
 /dev/rmt/*mbn
 /dev/rmt/*mn
 /dev/rmt/*n
 /dev/rmt/*u
 /dev/rmt/*ub
 /dev/rmt/*ubn
 /dev/rmt/*un

SEE ALSO tar(1), tcopy(1), ar(4), environ(4), mtio(7), st(7)

BUGS Not all devices support all options. Some options are hardware-dependent. Refer to the corresponding device manual page.
 mt is architecture sensitive. Heterogeneous operation (that is, Sun3 to Sun4 or visa versa) is not supported.

NAME	mv – move files
SYNOPSIS	mv [-f] [-i] <i>filename1</i> [<i>filename2</i> ...] <i>target</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The mv command moves <i>filenamen</i> to <i>target</i>. <i>filenamen</i> and <i>target</i> may not have the same name. (Care must be taken when using sh(1) metacharacters). If <i>target</i> is not a directory, only one file may be specified before it; if it is a directory, more than one file may be specified. If <i>target</i> does not exist, mv creates a file named <i>target</i>. If <i>target</i> exists and is not a directory, its contents are overwritten. If <i>target</i> is a directory the file(s) are moved to that directory.</p> <p>If mv determines that the mode of <i>target</i> forbids writing, it will print the mode (see chmod(2)), ask for a response, and read the standard input for one line. If the line begins with y, the mv occurs, if permissible; otherwise, the command exits. When the parent directory of <i>filenamen</i> is writable and has the sticky bit set, one or more of the following conditions must be true:</p> <ul style="list-style-type: none"> the user must own the file the user must own the directory the file must be writable by the user the user must be a privileged user <p>If <i>filenamen</i> is a directory, <i>target</i> must be a directory in the same physical file system. <i>target</i> and <i>filenamen</i> do not have to share the same parent directory.</p> <p>If <i>filenamen</i> is a file and <i>target</i> is a link to another file with links, the other links remain and <i>target</i> becomes a new file.</p>
OPTIONS	<p>-f mv will move the file(s) without prompting even if it is writing over an existing <i>target</i>. This option overrides the -i option. Note that this is the default if the standard input is not a terminal.</p> <p>-i mv will prompt for confirmation whenever the move would overwrite an existing <i>target</i>. A y answer means that the move should proceed. Any other answer prevents mv from overwriting the <i>target</i>.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of mv for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how mv behaves.</p>

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

chmod(1), **cp(1)**, **cpio(1)**, **ln(1)**, **rm(1)**, **environ(5)**

NOTES

If *filenamen* and *target* are on different file systems, **mv** copies the file and deletes the original; any links to other files are lost.

A **--** permits the user to mark explicitly the end of any command line options, allowing **mv** to recognize filename arguments that begin with a **-**. As an aid to BSD migration, **mv** will accept **-** as a synonym for **--**. This migration aid may disappear in a future release. If a **--** and a **-** both appear on the same command line, the second will be interpreted as a filename.

NAME	nawk – pattern scanning and processing language
SYNOPSIS	nawk [-F <i>re</i>] [-v <i>var=value</i>] [' <i>prog</i> '] [<i>filename ...</i>] nawk [-F <i>re</i>] [-v <i>var=value</i>] [-f <i>program-file</i>] [<i>filename ...</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	nawk scans each input <i>filename</i> for lines that match any of a set of patterns specified in <i>prog</i> . The <i>prog</i> string must be enclosed in single quotes (') to protect it from the shell. For each pattern in <i>prog</i> there may be an associated action performed when a line of a <i>filename</i> matches the pattern. The set of pattern-action statements may appear literally as <i>prog</i> or in a file specified with the -f <i>program-file</i> option. Input files are read in order; if there are no files, the standard input is read. The file name '-' means the standard input.
OPTIONS	-F <i>re</i> Specifies the input field separator. <i>re</i> can be either a character or a regular expression. -v <i>var=value</i> Allows the setting of a variable to <i>value</i> , before executing any part of the nawk script. The -v must be specified for each <i>variable</i> to be set. -f <i>program-file</i> Specify a file, <i>program-file</i> , to be used as the source for the program.
USAGE Input Lines	Each input line is matched against the pattern portion of every pattern-action statement; the associated action is performed for each matched pattern. Any <i>filename</i> of the form <i>var=value</i> is treated as an assignment, not a filename, and is executed at the time it would have been opened if it were a filename. <i>Variables</i> assigned in this manner are not available inside a BEGIN rule, and are assigned after previously specified files have been read. The option -v followed by <i>var=value</i> is an assignment to be done before <i>prog</i> is executed; any number of -v options may be present. An input line is normally made up of fields separated by white spaces. (This default can be changed by using the FS built-in variable or the -F <i>re</i> option.) The default is to ignore leading blanks and to separate fields by blanks and/or tab characters. However, if FS is assigned a value that does not include any of the white spaces, then leading blanks are not ignored. The fields are denoted \$1 , \$2 , ...; \$0 refers to the entire line.
Pattern-action Statements	A pattern-action statement has the form: <i>pattern</i> { <i>action</i> } Either pattern or action may be omitted. If there is no action, the matching line is printed. If there is no pattern, the action is performed on every input line. Pattern-action statements are separated by newlines or semicolons. Patterns are arbitrary Boolean combinations (!, , &&, and parentheses) of relational expressions and regular expressions. A relational expression is one of the following:

```

expression relop expression
expression matchop regular_expression
expression in array-name
(expression,expression, ... ) in array-name

```

where a *relop* is any of the six relational operators in C, and a *matchop* is either `~` (contains) or `!` (does not contain). An *expression* is an arithmetic expression, a relational expression, the special expression

```
var in array
```

or a Boolean combination of these.

Regular expressions are as in **egrep**(1). In patterns they must be surrounded by slashes. Isolated regular expressions in a pattern apply to the entire line. Regular expressions may also occur in relational expressions. A pattern may consist of two patterns separated by a comma; in this case, the action is performed for all lines between the occurrence of the first pattern to the occurrence of the second pattern.

The special patterns **BEGIN** and **iEND** may be used to capture control before the first input line has been read and after the last input line has been read respectively. These keywords do not combine with any other patterns.

Built-in Variables

Built-in variables include:

ARGC	command line argument count
ARGV	command line argument array
ENVIRON	array of environment variables; subscripts are names
FILENAME	name of the current input file
FNR	ordinal number of the current record in the current file
FS	input field separator regular expression (default blank and tab)
NF	number of fields in the current record
NR	ordinal number of the current record
OFMT	output format for numbers (default %6g)
OFS	output field separator (default blank)
ORS	output record separator (default new-line)
RLENGTH	length of string matched with the match() function
RS	input record separator (default new-line)
RSTART	starting position of string matched with the match() function
SUBSEP	separates multiple subscripts (default is 034)

An action is a sequence of statements. A statement may be one of the following:

```

if ( expression ) statement [ else statement ]
while ( expression ) statement
do statement while ( expression )
for ( expression ; expression ; expression ) statement

```

```

for ( var in array ) statement
delete array[subscript] #delete an array element
break
continue
{ [ statement ] ... }
expression      # commonly variable = expression
print [ expression-list ] [ >expression ]
printf format [ , expression-list ] [ >expression ]
next             # skip remaining patterns on this input line
exit [expr]    # skip the rest of the input; exit status is expr
return [expr]

```

Statements are terminated by semicolons, newlines, or right braces. An empty expression-list stands for the whole input line. Expressions take on string or numeric values as appropriate, and are built using the operators +, -, *, /, %, ^ and concatenation (indicated by a blank). The operators ++ -- += -= *= /= %= ^= > >= < <= == != ?: are also available in expressions. Variables may be scalars, array elements (denoted x[i]), or fields. Variables are initialized to the null string or zero. Array subscripts may be any string, not necessarily numeric; this allows for a form of associative memory. Multiple subscripts such as [i,j,k] are permitted; the constituents are concatenated, separated by the value of SUBSEP. String constants are quoted (""), with the usual C escapes recognized within.

The **print** statement prints its arguments on the standard output, or on a file if *>expression* is present, or on a pipe if ' | *cmd*' is present. The output resulted from the print statement is terminated by the output record separator with each argument separated by the current output field separator. The **printf** statement formats its expression list according to the format (see **printf(3S)**). The built-in function **close**(*expr*) closes the file or pipe *expr*.

The mathematical functions: **atan2**, **cos**, **exp**, **log**, **sin**, **sqrt**, are built-in.

Other built-in functions include:

gsub(*for,repl,in*) behaves like **sub** (see below), except that it replaces successive occurrences of the regular expression (like the **ed** global substitute command).

index(*s, t*) returns the position in string *s* where string *t* first occurs, or 0 if it does not occur at all.

int(*s*) truncates *s* to an integer value. If *s* is not specified, \$0 is used.

length(*s*) returns the length of its argument taken as a string, or of the whole line if there is no argument.

match(*s, re*) returns the position in string *s* where the regular expression *re* occurs, or 0 if it does not occur at all. **RSTART** is set to the starting position (which is the same as the returned value), and **RLENGTH** is set to the length of the matched string.

rand random number on (0, 1).

split(*s, a, fs*) splits the string *s* into array elements *a[1]*, *a[2]*, ... *a[n]*, and returns *n*.

The separation is done with the regular expression *fs* or with the field separator **FS** if *fs* is not given.

srand sets the seed for **rand**.

sprintf(*fmt, expr, expr, ...*) formats the expressions according to the **printf**(3S) format given by *fmt* and returns the resulting string.

sub(*for, repl, in*) substitutes the string *repl* in place of the first instance of the regular expression *for* in string *in* and returns the number of substitutions. If *in* is omitted, **nawk** substitutes in the current record (**\$0**).

substr(*s, m, n*) returns the *n*-character substring of *s* that begins at position *m*.

The input/output built-in functions are:

close(*filename*) closes the file or pipe named *filename*.

cmd | **getline** pipes the output of *cmd* into **getline** each successive call to **getline** returns the next line of output from *cmd*.

getline sets **\$0** to the next input record from the current input file.

getline <*filename*
sets **\$0** to the next record from *file*.

getline *x* sets variable *x* instead.

getline *x* <*filename*
sets *x* from the next record of *file*.

system(*cmd*) executes *cmd* and returns its exit status.

All forms of **getline** return 1 for successful input, 0 for end of file, and -1 for an error.

nawk also provides user-defined functions. Such functions may be defined (in the pattern position of a pattern-action statement) as

```
function name(args,...) { stmts }
```

Function arguments are passed by value if scalar and by reference if array name. Argument names are local to the function; all other variable names are global. Function calls may be nested and functions may be recursive. The **return** statement may be used to return a value.

EXAMPLES

Print lines longer than 72 characters:

```
length > 72
```

Print first two fields in opposite order:

```
{ print $2, $1 }
```

Same, with input fields separated by comma and/or blanks and tabs:

```
BEGIN { FS = "[ \t]*| [ \t]+"
```

```
      { print $2, $1 }
```

Add up first column, print sum and average:

```
    { s += $1 }
  END  { print "sum is", s, " average is", s/NR }
```

Print fields in reverse order:

```
  { for (i = NF; i > 0; --i) print $i }
```

Print all lines between start/stop pairs:

```
  /start/, /stop/
```

Print all lines whose first field is different from previous one:

```
  $1 != prev { print; prev = $1 }
```

Simulate `echo(1)`:

```
  BEGIN {
    for (i = 1; i < ARGV; i++)
      printf "%s", ARGV[i]
    printf "\n"
    exit
  }
```

Print a file, filling in page numbers starting at 5:

```
  /Page/ { $2 = n++; }
        { print }
```

Assuming this program is in a file named `prog`, the following command line prints the file `input` numbering its pages starting at 5: `nawk -f prog n=5 input`.

SEE ALSO `egrep(1)`, `grep(1)`, `sed(1)`, `printf(3S)`

The `awk` chapter in the *Solaris Advanced User's Guide*.

A. V. Aho, B. W. Kerningham, P. J. Weinberger, *The AWK Programming Language*
Addison-Wesley, 1988.

NOTES `nawk` is a new version of `awk` that provides capabilities unavailable in previous versions. This version will become the default version of `awk` in the next major release.

Input white space is not preserved on output if fields are involved.

There are no explicit conversions between numbers and strings. To force an expression to be treated as a number add 0 to it; to force it to be treated as a string concatenate the null string ("") to it.

NAME	newaliases – rebuild the data base for the mail aliases file
SYNOPSIS	newaliases
AVAILABILITY	SUNWnisu
DESCRIPTION	newaliases rebuilds the random access data base for the mail aliases file /etc/aliases . It is run automatically by sendmail(1M) (in the default configuration) whenever a message is sent.
FILES	/etc/aliases
SEE ALSO	sendmail(1M) , aliases(4)

NAME	newform – change the format of a text file
SYNOPSIS	newform [-s] [-itabspec] [-otabspec] [-bn] [-en] [-pn] [-an] [-f] [-cchar] [-ln] [<i>filename...</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	<p>newform reads lines from the named <i>filenames</i>, or the standard input if no input file is named, and reproduces the lines on the standard output. Lines are reformatted in accordance with command line options in effect.</p> <p>Except for -s, command line options may appear in any order, may be repeated, and may be intermingled with the optional <i>filenames</i>. Command line options are processed in the order specified. This means that option sequences like “-e15 -l60” will yield results different from “-l60 -e15”. Options are applied to all <i>filenames</i> on the command line.</p>
OPTIONS	<p>-s Shears off leading characters on each line up to the first tab and places up to 8 of the sheared characters at the end of the line. If more than 8 characters (not counting the first tab) are sheared, the eighth character is replaced by a * and any characters to the right of it are discarded. The first tab is always discarded.</p> <p>An error message and program exit will occur if this option is used on a file without a tab on each line. The characters sheared off are saved internally until all other options specified are applied to that line. The characters are then added at the end of the processed line.</p> <p>For example, to convert a file with leading digits, one or more tabs, and text on each line, to a file beginning with the text, all tabs after the first expanded to spaces, padded with spaces out to column 72 (or truncated to column 72), and the leading digits placed starting at column 73, the command would be:</p> <p style="padding-left: 40px;">newform -s -i -l -a -e file-name</p> <p>-itabspec Input tab specification: expands tabs to spaces, according to the tab specifications given. <i>Tabspec</i> recognizes all tab specification forms described in tabs(1). In addition, <i>tabspec</i> may be --, in which newform assumes that the tab specification is to be found in the first line read from the standard input (see fspec(4)). If no <i>tabspec</i> is given, <i>tabspec</i> defaults to -8. A <i>tabspec</i> of -0 expects no tabs; if any are found, they are treated as -1.</p> <p>-otabspec Output tab specification: replaces spaces by tabs, according to the tab specifications given. The tab specifications are the same as for -itabspec. If no <i>tabspec</i> is given, <i>tabspec</i> defaults to -8. A <i>tabspec</i> of -0 means that no spaces will be converted to tabs on output.</p> <p>-bn Truncate <i>n</i> characters from the beginning of the line when the line length is greater than the effective line length (see -ln). Default is to truncate the number of characters necessary to obtain the effective line length. The default value is used when -b with no <i>n</i> is used. This option can be used to delete the</p>

sequence numbers from a COBOL program as follows:

newform -l1 -b7 file-name

- en** Same as **-bn** except that characters are truncated from the end of the line.
- pn** Prefix *n* characters (see **-cchar**) to the beginning of a line when the line length is less than the effective line length. Default is to prefix the number of characters necessary to obtain the effective line length.
- an** Same as **-pn** except characters are appended to the end of a line.
- f** Write the tab specification format line on the standard output before any other lines are output. The tab specification format line which is printed will correspond to the format specified in the *last* **-o** option. If no **-o** option is specified, the line which is printed will contain the default specification of **-8**.
- cchar** Change the prefix/append character to *char*. Default character for *char* is a space.
- ln** Set the effective line length to *n* characters. If *n* is not entered, **-l** defaults to 72. The default line length without the **-l** option is 80 characters. Note: Tabs and backspaces are considered to be one character (use **-i** to expand tabs to spaces).

The **-l1** must be used to set the effective line length shorter than any existing line in the file so that the **-b** option is activated.

SEE ALSO

csplit(1), tabs(1), fspec(4)

DIAGNOSTICS

All diagnostics are fatal.

usage: ...

newform was called with a bad option.

"not -s format"

There was no tab on one line.

"can't open file"

Self-explanatory.

"internal line too long"

A line exceeds 512 characters after being expanded in the internal work buffer.

"tabspec in error"

A tab specification is incorrectly formatted, or specified tab stops are not ascending.

"tabspec indirection illegal"

A *tabspec* read from a file (or standard input) may not contain a *tabspec* referencing another file (or standard input).

0 – normal execution

1 – for any error

NOTES

newform normally only keeps track of physical characters; however, for the **-i** and **-o** options, **newform** will keep track of backspaces in order to line up tabs in the appropriate logical columns.

newform will not prompt the user if a *tabspec* is to be read from the standard input (by use of **-i** or **-o**).

If the **-f** option is used, and the last **-o** option specified was **-o**, and was preceded by either a **-o** or a **-i**, the tab specification format line will be incorrect.

NAME	newgrp – shell built-in function to allow new group permissions to the user
SYNOPSIS	
sh	newgrp [<i>argument</i>]
ksh	† newgrp [<i>arg...</i>]
DESCRIPTION	
sh	Equivalent to exec newgrp <i>argument</i> . See newgrp(1M) for usage and description.
ksh	Equivalent to exec /bin/newgrp <i>arg...</i> On this man page, ksh(1) commands that are preceded by one or two † (daggers) are treated specially in the following ways: <ol style="list-style-type: none">1. Variable assignment lists preceding the command remain in effect when the command completes.2. I/O redirections are processed after variable assignments.3. Errors cause a script that contains them to abort.4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.
SEE ALSO	ksh(1) , sh(1) , newgrp(1M)

NAME	news – print news items
SYNOPSIS	news [-a] [-n] [-s] [<i>items</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	<p>news is used to keep the user informed of current events. By convention, these events are described by files in the directory /var/news.</p> <p>When invoked without arguments, news prints the contents of all current files in /var/news, most recent first, with each preceded by an appropriate header. news stores the “currency” time as the modification date of a file named .news_time in the user’s home directory (the identity of this directory is determined by the environment variable \$HOME); only files more recent than this currency time are considered “current.”</p>
OPTIONS	<p>-a news prints all items, regardless of currency. In this case, the stored time is not changed.</p> <p>-n news reports the names of the current items without printing their contents, and without changing the stored time.</p> <p>-s news reports how many current items exist, without printing their names or contents, and without changing the stored time. It is useful to include such an invocation of news in one’s .profile file, or in the system’s /etc/profile.</p> <p>All other arguments are assumed to be specific news items that are to be printed. If a <i>delete</i> is typed during the printing of a news item, printing stops and the next item is started. Another <i>delete</i> within one second of the first causes the program to terminate.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of news for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how news behaves.</p> <p>LC_CTYPE</p> <p>Determines how news handles characters. When LC_CTYPE is set to a valid value, news can display and handle text and filenames containing valid characters for that locale. news can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. news can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p>

FILES /etc/profile
/var/news/*
\$HOME/.news_time

SEE ALSO profile(4), environ(5)

NAME	nice – run a command at low priority
SYNOPSIS	/usr/bin/nice [<i>-increment</i>] <i>command</i> [<i>arguments</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>/usr/bin/nice executes <i>command</i> with a lower CPU scheduling priority. The priocntl command is a more general interface to scheduler functions.</p> <p>The invoking process (generally the user's shell) must be in a scheduling class that supports the /usr/bin/nice command.</p>
OPTIONS	<p><i>-increment</i> If the <i>increment</i> argument (in the range 1–19) is given, it is used; if not, an increment of 10 is assumed.</p> <p>The super-user may run commands with priority higher than normal by using a negative increment, for example, —10. A negative increment assigned by an unprivileged user is ignored.</p> <p>If using the cs, the syntax on this man page does not apply, unless /usr/bin/nice is invoked at the beginning of the command line. Instead, refer to cs(1) for a different nice syntax, if in the cs.</p>
SEE ALSO	cs(1) , nohup(1) , priocntl(1) , shell_builtins(1) , nice(2)
DIAGNOSTICS	/usr/bin/nice returns the exit status of <i>command</i> .
NOTES	An <i>increment</i> larger than 19 is equivalent to 19.

NAME	nis+, NIS+, nis – a new version of the network information name service
DESCRIPTION	<p>NIS+ is a new version of the network information nameservice. This version differs in several significant ways from version 2, which is referred to as NIS or YP in earlier releases. Specific areas of enhancement include the ability to scale to larger networks, security, and the administration of the service.</p> <p>The man pages for NIS+ are broken up into three basic categories. Those in section 1 are the user commands that are most often executed from a shell script or directly from the command line. Section 1M man pages describe utility commands that can be used by the network administrator to administer the service itself. The NIS+ programming API is described by man pages in section 3N.</p> <p>All commands and functions that use NIS version 2 are prefixed by the letters yp as in ypmatch(1), ypcat(1), yp_match(3N), and yp_first(3N). Commands and functions that use the new replacement software NIS+ are prefixed by the letters nis as in nismatch(1), nischown(1), nis_list(3N), and nis_add_entry(3N). A complete list of NIS+ commands is in the LIST OF COMMANDS section.</p> <p>This man page introduces the NIS+ terminology. It also describes the NIS+ namespace, authentication, and authorization policies.</p>
NIS+ NAMESPACE	<p>The naming model of NIS+ is based upon a tree structure. Each node in the tree corresponds to an NIS+ object. There are six types of NIS+ objects: <i>directory</i>, <i>table</i>, <i>group</i>, <i>link</i>, <i>entry</i>, and <i>private</i>.</p>
NIS+ Directory Object	<p>Each NIS+ namespace will have at least one NIS+ directory object. An NIS+ directory is like a UNIX file system directory which contains other NIS+ objects including NIS+ directories. The NIS+ directory that forms the root of the NIS+ namespace is called the root directory. There are two special NIS+ directories: org_dir and groups_dir. The org_dir directory consists of all the system-wide administration tables, such as passwd, hosts, and mail_aliases. The groups_dir directory consists of NIS+ group objects which are used for access control. The collection of org_dir, groups_dir and their parent directory is referred to as an NIS+ domain. NIS+ directories can be arranged in a tree-like structure so that the NIS+ namespace can match the organizational or administrative hierarchy.</p>
NIS+ Table Object	<p>NIS+ tables (not files), contained within NIS+ directories, store the actual information about some particular type. For example, the hosts system table stores information about the IP address of the hosts in that domain. NIS+ tables are multicolumn and the tables can be searched through any of the searchable columns. Each table object defines the schema for its table. The NIS+ tables consist of NIS+ entry objects. For each entry in the NIS+ table, there is an NIS+ entry object. NIS+ entry objects conform to the schema defined by the NIS+ table object.</p>

NIS+ Group Object	NIS+ group objects are used for access control at group granularity. NIS+ group objects, contained within the groups_dir directory of a domain, contain a list of all the NIS+ principals within a certain NIS+ group. An NIS+ principal is a user or a machine making NIS+ requests.
NIS+ Link Object	NIS+ link objects are like UNIX symbolic file-system links—they are typically used for shortcuts in the NIS+ namespace. Refer to nis_objects(3N) for more information about the NIS+ objects.
NIS+ NAMES	The NIS+ service defines two forms of names, <i>simple</i> names and <i>indexed</i> names. Simple names are used by the service to identify NIS+ objects contained within the NIS+ namespace. Indexed names are used to identify NIS+ entries contained within NIS+ tables. Furthermore, entries within NIS+ tables are returned to the caller as NIS+ objects of type <i>entry</i> . NIS+ objects are implemented as a union structure which is described in the file <code><rpcsvc/nis_object.x></code> . The differences between the various types and the meanings of the components of these objects are described in nis_objects(3N) .
Simple Names	<p>Simple names consist of a series of labels that are separated by the '.' (dot) character. Each label is composed of printable characters from the ISO Latin 1 set. Each label can be of any nonzero length, provided that the fully qualified name is fewer than NIS_MAXNAMELEN octets including the separating dots. (See <code><rpcsvc/nis.h></code> for the actual value of NIS_MAXNAMELEN in the current release.) Labels that contain special characters (see Grammar) must be quoted.</p> <p>The NIS+ namespace is organized as a singly rooted tree. Simple names identify nodes within this tree. These names are constructed such that the leftmost label in a name identifies the leaf node and all of the labels to the right of the leaf identify that object's parent node. The parent node is referred to as the leaf's <i>directory</i>. This is a naming directory and should not be confused with a file system directory.</p> <p>For example, the name <i>example.simple.name</i> is a simple name with three labels, where <i>example</i> is the leaf node in this name, the directory of this leaf is <i>simple.name</i> which by itself is a simple name. The leaf of which is <i>simple</i> and its directory is simply <i>name</i>.</p> <p>The function nis_leaf_of(3N) returns the first label of a simple name. The function nis_domain_of(3N) returns the name of the directory that contains the leaf. Iterative use of these two functions can break a simple name into each of its label components.</p> <p>The name '.' (dot) is reserved to name the <i>global root</i> of the namespace. For systems that are connected to the Internet, this global root will be served by a Domain Name Service. When an NIS+ server is serving a root directory whose name is not '.' (dot) this directory is referred to as a <i>local root</i>.</p> <p>NIS+ names are said to be <i>fully qualified</i> when the name includes all of the labels identifying all of the directories, up to the global root. Names without the trailing dot are called <i>partially qualified</i>.</p>

Indexed Names

Indexed names are compound names that are composed of a search criterion and a simple name. The search criterion component is used to select entries from a table; the simple name component is used to identify the NIS+ table that is to be searched. The search criterion is a series of column names and their desired values enclosed in bracket '[' characters. These criteria take the following form:

[*column_name=value, column_name=value, ...*]

A search criterion is combined with a simple name to form an indexed name by concatenating the two parts, separated by a ',' (comma) character as follows.

[*search-criterion*],*table.directory*.

When multiple column name/value pairs are present in the search criterion, only those entries in the table that have the appropriate value in all columns specified are returned. When no column name/value pairs are specified in the search criterion, [], *all* entries in the table are returned.

Grammar

The following text represents a context-free grammar that defines the set of legal NIS+ names. The terminals in this grammar are the characters '.' (dot), '[' (open bracket), ']' (close bracket), ',' (comma), '=' (equals) and whitespace. Angle brackets ('<' and '>'), which delineate non-terminals, are not part of the grammar. The character '|' (vertical bar) is used to separate alternate productions and should be read as "this production **OR** this production".

```

name ::= . | <simple name> | <indexed name>
simple name ::= <string>. | <string>.<simple name>
indexed name ::= <search criterion>,<simple name>
search criterion ::= [ <attribute list> ]
attribute list ::= <attribute> | <attribute>,<attribute list>
attribute ::= <string> = <string>
string ::= ISO Latin 1 character set except the
character '/' (slash). The initial character
may not be a terminal character or the
characters '@' (at), '+' (plus), or '-'
hyphen.

```

Terminals that appear in strings must be quoted with "" (double quote). The "" character may be quoted by quoting it with itself "" "".

Name Expansion

The NIS+ service only accepts fully qualified names. However, since such names may be unwieldy, the NIS+ commands in section 1 employ a set of standard expansion rules that will attempt to fully qualify a partially qualified name. This expansion is actually done by the NIS+ library function **nis_getnames(3N)** which generates a list of names using the default NIS+ directory search path or the **NIS_PATH** environment variable. The default NIS+ directory search path includes all the names in its path. **nis_getnames()** is invoked by the functions **nis_lookup(3N)** and **nis_list(3N)** when the **EXPAND_NAME** flag is used.

The **NIS_PATH** environment variable contains an ordered list of simple names. The names are separated by the ':' (colon) character. If any name in the list contains colons, the colon should be quoted as described in the **Grammar** section. When the list is exhausted, the resolution function returns the error **NIS_NOTFOUND**. This may mask the fact that the name existed but a server for it was unreachable. If the name presented to the list or lookup interface is fully qualified, the **EXPAND_NAME** flag is ignored.

In the list of names from the **NIS_PATH** environment variable, the '\$' (dollar sign) character is treated specially. Simple names that end with the label '\$' have this character replaced by the default directory (see **nis_local_directory(3N)**). Using "\$" as a name in this list results in this name being replaced by the list of directories between the default directory and the global root that contain at least two labels.

Below is an example of this expansion. Given the default directory of *some.long.domain.name.*, and the **NIS_PATH** variable set to **fred.bar:org_dir:\$**. This path is initially broken up into the list:

- 1 **fred.bar.**
- 2 **org_dir.\$**
- 3 **\$**

The dollar sign in the second component is replaced by the default directory. The dollar sign in the third component is replaced with the names of the directories between the default directory and the global root that have at least two labels in them. The effective path value becomes:

- 1 **fred.bar.**
- 2a **org_dir.some.long.domain.name.**
- 3a **some.long.domain.name.**
- 3b **long.domain.name.**
- 3c **domain.name.**

Each of these simple names is appended to the partially qualified name that was passed to the **nis_lookup(3N)** or **nis_list(3N)** interface. Each is tried in turn until **NIS_SUCCESS** is returned or the list is exhausted.

If the **NIS_PATH** variable is not set, the path "\$" is used.

The library function **nis_getnames(3N)** can be called from user programs to generate the list of names that would be attempted. The program **nisdefaults(1)** with the **-s** option can also be used to show the fully expanded path.

Concatenation Path

Normally all the entries for a certain type of information are stored within the table itself. However, there are times when it is desirable for the table to point to other tables where entries can be found. For example, you may want to store all the IP addresses in the host table for their own domain, and yet want to be able to resolve hosts in some other domain without explicitly specifying the new domain name. NIS+ provides a mechanism for concatenating different but related tables with a "NIS+ Concatenation Path". With a concatenation path, you can create a sort of flat namespace from a hierarchical

structure. You can also create a table with no entries and just point the hosts or any other table to its parent domain. Note that with such a setup, you are moving the administrative burden of managing the tables to the parent domain. The concatenation path will slow down the request response time because more tables and more servers are searched. It will also decrease the availability if all the servers are incapacitated for a particular directory in the table path.

The NIS+ Concatenation Path is also referred to as the "table path". This path is set up at table creation time through **nistbladm**(1). You can specify more than one table to be concatenated and they will be searched in the given order. Note that the NIS+ client libraries, by default, will not follow the concatenation path set in site-specific tables. Refer to **nis_list**(3N) for more details.

Namespaces

The NIS+ service defines two additional *disjoint* namespaces for its own use. These namespaces are the NIS+ *Principal* namespace, and the NIS+ *Group* namespace. The names associated with the group and principal namespaces are syntactically identical to simple names. However, the information they represent *cannot* be obtained by directly presenting these names to the NIS+ interfaces. Instead, special interfaces are defined to map these names into NIS+ names so that they may then be resolved.

Principal Names

NIS+ principal names are used to uniquely identify users and machines that are making NIS+ requests. These names have the form:

principal.domain

Here *domain* is the fully qualified name of an NIS+ directory where the named principal's credentials can be found. See **Directories and Domains** for more information on domains. Note that in this name, *principal*, is not a leaf in the NIS+ namespace.

Credentials are used to map the identity of a host or user from one context such as a process UID into the NIS+ context. They are stored as records in an NIS+ table named *cred*, which always appears in the *org_dir* subdirectory of the directory named in the principal name.

This mapping can be expressed as a replacement function:

principal.domain -> [**cname=principal.domain**],**cred.org_dir.domain**

This latter name is an NIS+ name that can be presented to the **nis_list**(3N) interface for resolution. NIS+ principal names are administered using the **nisaddcred**(1M) command.

The *cred* table contains five columns named *cname*, *auth_name*, *auth_type*, *public_data*, and *private_data*. There is one record in this table for each identity mapping for an NIS+ principal. The current service supports two such mappings:

- LOCAL This mapping is used to map from the UID of a given process to the NIS+ principal name associated with that UID. If no mapping exists, the name *nobody* is returned. When the effective UID of the process is 0 (for example, the super-user), the NIS+ name associated with the host is returned. Note that UIDs are sensitive to the context of the machine on which the process is executing.
- DES This mapping is used to map to and from a Secure RPC "netname" into an NIS+ principal name. See **secure_rpc**(3N) for more information on netnames.

Note that since netnames contain the notion of a domain, they span NIS+ directories.

The NIS+ client library function **nis_local_principal**(3N) uses the *cred.org_dir* table to map the UNIX notion of an identity, a process' UID, into an NIS+ principal name. Shell programs can use the program **nisdefaults**(1) with the **-p** switch to return this information.

Mapping from UIDs to an NIS+ principal name is accomplished by constructing a query of the form:

[auth_type=LOCAL, auth_name=uid],cred.org_dir.default-domain.

This query will return a record containing the NIS+ principal name associated with this UID, in the machine's default domain.

The NIS+ service uses the DES mapping to map the names associated with Secure RPC requests into NIS+ principal names. RPC requests that use Secure RPC include the *netname* of the client making the request in the RPC header. This netname has the form:

unix.UID@domain

The service constructs a query using this name of the form:

[auth_type=DES, auth_name=netname],cred.org_dir.domain.

where the domain part is extracted from the netname rather than using the default domain. This query is used to look up the mapping of this netname into an NIS+ principal name in the domain where it was created.

This mechanism of mapping UID and netnames into an NIS+ principal name guarantees that a client of the NIS+ service has only one principal name. This principal name is used as the basis for authorization which is described below. All objects in the NIS+ namespace and all entries in NIS+ tables must have an owner specified for them. This owner field always contains an NIS+ principal name.

Group Names

Like NIS+ principal names, NIS+ group names take the form:

group_name.domain

All objects in the NIS+ namespace and all entries in NIS+ tables may optionally have a *group owner* specified for them. This group owner field, when filled in, always contains the fully qualified NIS+ group name.

The NIS+ client library defines several interfaces (**nis_groups**(3N)) for dealing with NIS+ groups. These interfaces internally map NIS+ group names into an NIS+ simple name which identifies the NIS+ group object associated with that group name. This mapping can be shown as follows:

group.domain -> **group.groups_dir.domain**

This mapping eliminates collisions between NIS+ group names and NIS+ directory names. For example, without this mapping, a directory with the name *engineering.foo.com.*, would make it impossible to have a group named *engineering.foo.com.*. This is due to the restriction that within the NIS+ namespace, a name unambiguously identifies a single object. With this mapping, the NIS+ *group* name *engineering.foo.com.* maps to the NIS+ *object* name *engineering.groups_dir.foo.com.*

NIS+ SECURITY

The contents of a group object is a list of NIS+ principal names, and the names of other NIS+ groups. See **nis_groups(3N)** for a more complete description of their use.

NIS+ defines a security model to control access to information managed by the service. The service defines access rights that are selectively granted to individual clients or groups of clients. Principal names and group names are used to define clients and groups of clients that may be granted or denied access to NIS+ information. These principals and groups are associated with NIS+ domains as defined below.

The security model also uses the notion of a class of principals called *nobody*, which contains all clients, whether or not they have authenticated themselves to the service. The class *world* includes any client who has been authenticated.

Directories and Domains

Some directories within the NIS+ namespace are referred to as NIS+ *Domains*. Domains are those NIS+ directories that contain the subdirectories *groups_dir* and *org_dir*. Further, the subdirectory *org_dir* should contain the table named *cred*. NIS+ Group names and NIS+ Principal names **always** include the NIS+ domain name after their first label.

Authentication

The NIS+ name service uses Secure RPC for the integrity of the NIS+ service. This requires that users of the service and their machines must have a Secure RPC key pair associated with them. This key is initially generated with either the **nisaddcred(1M)** or **nisclient(1M)** commands and modified with the **chkey(1)** or **nispasswd(1)** commands.

The use of Secure RPC allows private information to be stored in the name service that will not be available to untrusted machines or users on the network.

In addition to the Secure RPC key, users need a mapping of their UID into an NIS+ principal name. This mapping is created by the system administrator using the **nisclient(1M)** or **nisaddcred(1M)** command.

Users that will be using machines in several NIS+ domains must insure that they have a *local* credential entry in each of those domains. This credential should be created with the NIS+ principal name of the user in their “home” domain. For the purposes of NIS+ and Secure RPC, the home domain is defined to be the one where your Secure RPC key pair is located.

Authorization

The NIS+ service defines four access rights that can be granted or denied to clients of the service. These rights are *read*, *modify*, *create*, and *destroy*. These rights are specified in the object structure at creation time and may be modified later with the **nischmod(1)** command. In general, the rights granted for an object apply only to that object. However, for purposes of authorization, rights granted to clients reading *directory* and *table* objects are granted to those clients for all of the objects “contained” by the parent object. This notion of containment is abstract. The objects do not actually contain other objects within them. Note that *group* objects do contain the list of principals within their definition.

Access rights are interpreted as follows:

read This right grants read access to an object. For directory and table objects, having read access on the parent object conveys read access to all of the objects that are direct children of a directory, or entries within a table.

- modify** This right grants modification access to an existing object. Read access is not required for modification. However, in many applications, one will need to read an object before modifying it. Such modify operations will fail unless read access is also granted.
- create** This right gives a client permission to create new objects where one had not previously existed. It is only used in conjunction with directory and table objects. Having create access for a table allows a client to add additional entries to the table. Having create access for a directory allows a client to add new objects to an NIS+ directory.
- destroy** This right gives a client permission to destroy or remove an existing object or entry. When a client attempts to destroy an entry or object by removing it, the service first checks to see if the table or directory containing that object grants the client destroy access. If it does, the operation proceeds, if the containing object does not grant this right then the object itself is checked to see if it grants this right to the client. If the object grants the right, then the operation proceeds; otherwise the request is rejected.

Each of these rights may be granted to any one of four different categories.

owner A right may be granted to the *owner* of an object. The owner is the NIS+ principal identified in the owner field. The owner can be changed with the **nischown(1)** command. Note that if the owner does not have modification access rights to the object, the owner cannot change any access rights to the object, unless the owner has modification access rights to its parent object.

group owner

A right may be granted to the *group owner* of an object. This grants the right to any principal that is identified as a member of the group associated with the object. The group owner may be changed with the **nischgrp(1)** command. The object owner need not be a member of this group.

world A right may be granted to everyone in the *world*. This grants the right to all clients who have authenticated themselves with the service.

nobody A right may be granted to the *nobody* principal. This has the effect of granting the right to any client that makes a request of the service, regardless of whether they are authenticated or not.

Note that for bootstrapping reasons, directory objects that are NIS+ domains, the *org_dir* subdirectory and the *cred* table within that subdirectory must have *read* access to the *nobody* principal. This makes navigation of the namespace possible when a client is in the process of locating its credentials. Granting this access does not allow the contents of other tables within *org_dir* to be read (such as the entries in the password table) unless the table itself gives "real" access rights to the *nobody* principal.

**Directory
Authorization**

Additional capabilities are provided for granting access rights to clients for directories. These rights are contained within the *object access rights* (OAR) structure of the directory. This structure allows the NIS+ service to grant rights that are not granted by the directory object to be granted for objects contained by the directory of a specific type.

An example of this capability is a directory object which does not grant create access to all clients, but does grant create access in the OAR structure for *group* type objects to clients who are members of the NIS+ group associated with the directory. In this example the only objects that could be created as children of the directory would have to be of the type *group*.

Another example is a directory object that grants create access only to the owner of the directory, and then additionally grants create access through the OAR structure for objects of type *table*, *link*, *group*, and *private* to any member of the directory's group. This has the effect of giving nearly complete create access to the group with the exception of creating subdirectories. This restricts the creation of new NIS+ domains because creating a domain requires creating both a *groups_dir* and *org_dir* subdirectory.

Note that there is currently no command line interface to set or change the OAR of the directory object.

Table Authorization

As with directories, additional capabilities are provided for granting access to entries within tables. Rights granted to a client by the access rights field in a table object apply to the table object and all of the entry objects "contained" by that table. If an access right is not granted by the table object, it may be granted by an entry within the table. This holds for all rights except *create*.

For example, a table may not grant read access to a client performing a **nis_list(3N)** operation on the table. However, the access rights field of entries within that table may grant read access to the client. Note that access rights in an entry are granted to the owner and group owner of the *entry* and not the owner or group of the table. When the list operation is performed, all entries that the client has read access to are returned. Those entries that do not grant read access are not returned. If none of the entries that match the search criterion grant read access to the client making the request, no entries are returned and the result status contains the NIS_NOTFOUND error code.

Access rights that are granted by the rights field in an entry are granted for the entire entry. However, in the table object an additional set of access rights is maintained for each column in the table. These rights apply to the equivalent column in the entry. The rights are used to grant access when neither the table nor the entry itself grant access. The access rights in a column specification apply to the owner and group owner of the entry rather than the owner and group owner of the table object.

When a read operation is performed, if read access is not granted by the table and is not granted by the entry but *is* granted by the access rights in a column, that entry is returned with the correct values in all columns that are readable and the string ***NP*** (No Permission) in columns where read access is not granted.

As an example, consider a client that has performed a list operation on a table that does not grant read access to that client. Each entry object that satisfied the search criterion specified by the client is examined to see if it grants read access to the client. If it does, it is included in the returned result. If it does not, then each column is checked to see if it grants read access to the client. If any columns grant read access to the client, data in those columns is returned. Columns that do not grant read access have their contents replaced by the string ***NP***. If none of the columns grant read access, then the entry is

not returned.

The following lists all commands and programming functions related to NIS+:

LIST OF COMMANDS

NIS+ User Commands

nisaddent(1)	add /etc files and NIS maps into their corresponding NIS+ tables
niscat(1)	display NIS+ tables and objects
nischgrp(1)	change the group owner of a NIS+ object
nischmod(1)	change access rights on a NIS+ object
nischown(1)	change the owner of a NIS+ object
nischttl(1)	change the time to live value of a NIS+ object
nisdefaults(1)	display NIS+ default values
niserror(1)	display NIS+ error messages
nisgrep(1)	utilities for searching NIS+ tables
nisgrpadm(1)	NIS+ group administration command
nisln(1)	symbolically link NIS+ objects
nisl(1)	list the contents of a NIS+ directory
nismatch(1)	utilities for searching NIS+ tables
nismkdir(1)	create NIS+ directories
nispasswd(1)	change NIS+ password information
nism(1)	remove NIS+ objects from the namespace
nismrmdir(1)	remove NIS+ directories
nisshowcache(1)	NIS+ utility to print out the contents of the shared cache file
nistbladm(1)	NIS+ table administration command
nistest(1)	return the state of the NIS+ namespace using a conditional expression

NIS+ Administrative Commands

aliasadm(1M)	manipulate the NIS+ aliases map
nis_cachemgr(1M)	NIS+ utility to cache location information about NIS+ servers
nisaddcred(1M)	create NIS+ credentials
nisaddent(1M)	create NIS+ tables from corresponding /etc files or NIS maps
nisclient(1M)	initialize NIS+ credentials for NIS+ principals
nisd(1M)	NIS+ service daemon
nisd_resolv(1M)	NIS+ service daemon
nisinit(1M)	NIS+ client and server initialization utility
nislog(1M)	display the contents of the NIS+ transaction log
nisping(1M)	send ping to NIS+ servers
nispopulate(1M)	populate the NIS+ tables in a NIS+ domain
nisserver(1M)	set up NIS+ servers
nissetup(1M)	initialize a NIS+ domain
nisshowcache(1M)	NIS+ utility to print out the contents of the shared cache file

**NIS+ Programming
API**

nisstat(1M)	report NIS+ server statistics
nisupdkeys(1M)	update the public keys in a NIS+ directory object
rpc.nisd(1M)	NIS+ service daemon
rpc.nisd_resolv(1M)	NIS+ service daemon
sysidnis(1M)	system configuration
__nis_map_group(3N)	NIS+ group manipulation functions
db_add_entry(3N)	NIS+ Database access functions
db_checkpoint(3N)	NIS+ Database access functions
db_create_table(3N)	NIS+ Database access functions
db_destroy_table(3N)	NIS+ Database access functions
db_first_entry(3N)	NIS+ Database access functions
db_free_result(3N)	NIS+ Database access functions
db_initialize(3N)	NIS+ Database access functions
db_list_entries(3N)	NIS+ Database access functions
db_next_entry(3N)	NIS+ Database access functions
db_remove_entry(3N)	NIS+ Database access functions
db_reset_next_entry(3N)	NIS+ Database access functions
db_standby(3N)	NIS+ Database access functions
db_table_exists(3N)	NIS+ Database access functions
db_unload_table(3N)	NIS+ Database access functions
nis_add(3N)	NIS+ namespace functions
nis_add_entry(3N)	NIS+ table functions
nis_addmember(3N)	NIS+ group manipulation functions
nis_checkpoint(3N)	misc NIS+ log administration functions
nis_clone_object(3N)	NIS+ subroutines
nis_creategroup(3N)	NIS+ group manipulation functions
nis_db(3N)	NIS+ Database access functions
nis_destroy_object(3N)	NIS+ subroutines
nis_destroygroup(3N)	NIS+ group manipulation functions
nis_dir_cmp(3N)	NIS+ subroutines
nis_domain_of(3N)	NIS+ subroutines
nis_error(3N)	display NIS+ error messages
nis_first_entry(3N)	NIS+ table functions
nis_freenames(3N)	NIS+ subroutines
nis_freeresult(3N)	NIS+ namespace functions
nis_freeservlist(3N)	miscellaneous NIS+ functions
nis_freetags(3N)	miscellaneous NIS+ functions
nis_getnames(3N)	NIS+ subroutines
nis_getservlist(3N)	miscellaneous NIS+ functions
nis_groups(3N)	NIS+ group manipulation functions
nis_ismember(3N)	NIS+ group manipulation functions
nis_leaf_of(3N)	NIS+ subroutines
nis_lerror(3N)	display some NIS+ error messages
nis_list(3N)	NIS+ table functions

	nis_local_directory (3N)	NIS+ local names
	nis_local_group (3N)	NIS+ local names
	nis_local_host (3N)	NIS+ local names
	nis_local_names (3N)	NIS+ local names
	nis_local_principal (3N)	NIS+ local names
	nis_lookup (3N)	NIS+ namespace functions
	nis_map_group (3N)	NIS+ group manipulation functions
	nis_mkdir (3N)	miscellaneous NIS+ functions
	nis_modify (3N)	NIS+ namespace functions
	nis_modify_entry (3N)	NIS+ table functions
	nis_name_of (3N)	NIS+ subroutines
	nis_names (3N)	NIS+ namespace functions
	nis_next_entry (3N)	NIS+ table functions
	nis_objects (3N)	NIS+ object formats
	nis_perror (3N)	display NIS+ error messages
	nis_ping (3N)	misc NIS+ log administration functions
	nis_print_group_entry (3N)	NIS+ group manipulation functions
	nis_print_object (3N)	NIS+ subroutines
	nis_remove (3N)	NIS+ namespace functions
	nis_remove_entry (3N)	NIS+ table functions
	nis_removemember (3N)	NIS+ group manipulation functions
	nis_rmdir (3N)	miscellaneous NIS+ functions
	nis_server (3N)	miscellaneous NIS+ functions
	nis_servstate (3N)	miscellaneous NIS+ functions
	nis_sperrno (3N)	display NIS+ error messages
	nis_sperror (3N)	display NIS+ error messages
	nis_sperror_r (3N)	display NIS+ error messages
	nis_stats (3N)	miscellaneous NIS+ functions
	nis_subr (3N)	NIS+ subroutines
	nis_tables (3N)	NIS+ table functions
	nis_verifygroup (3N)	NIS+ group manipulation functions
NIS+ Files and Directories FILES	nisfiles (4)	NIS+ database files and directory structure
	<rpcsvc/nis_object.x>	protocol description of an NIS+ object
	<rpcsvc/nis.x>	defines the NIS+ protocol using the RPC language as described in the <i>Network Interfaces Programmer's Guide</i>
	<rpcsvc/nis.h>	should be included by all clients of the NIS+ service
SEE ALSO	nischown (1), nisdefaults (1), nismatch (1), nispasswd (1), admintool (1M), newkey (1M), nisaddcred (1M), nisclient (1M), nispopulate (1M), nisserver (1M), nis_add_entry (3N), nis_domain_of (3N), nis_getnames (3N), nis_groups (3N), nis_leaf_of (3N), nis_list (3N), nis_local_directory (3N), nis_lookup (3N), nis_objects (3N)	

Network Interfaces Programmer's Guide

Describes the application programming interfaces for networks including NIS+
Name Service Configuration Guide

Describes how to plan for and configure an NIS+ namespace
Name Services Administration Guide

Describes how to administer a running NIS+ namespace and modify its security
NIS+ Transition Guide

Describes how to make the transition from NIS to NIS+
Administration Application Reference Manual

Describes the **admintool**(1M) window interface for modifying the data in NIS+ tables

NAME	niscat – display NIS+ tables and objects
SYNOPSIS	niscat [-AhLMv] <i>tablename</i> ... niscat [-ALMP] -o <i>name</i> ...
AVAILABILITY	SUNWnisu
DESCRIPTION	In the first synopsis, niscat displays the contents of the NIS+ tables named by <i>tablename</i> . In the second synopsis, it displays the internal representation of the NIS+ objects named by <i>name</i> .
OPTIONS	<p>-A Display the data within the table and all of the data in tables in the initial table's concatenation path.</p> <p>-h Display the header line prior to displaying the table. The header consists of the '#' (hash) character followed by the name of each column. The column names are separated by the table separator character.</p> <p>-L Follow links. When this option is specified, if <i>tablename</i> or <i>name</i> names a LINK type object, the link is followed and the object or table named by the link is displayed.</p> <p>-M Master server only. This option specifies that the request should be sent to the master server of the named data. This guarantees that the most up-to-date information is seen at the possible expense of increasing the load on the master server and increasing the possibility of the NIS+ server being unavailable or busy for updates.</p> <p>-P Follow concatenation path. This option specifies that the request should follow the concatenation path of a table if the initial search is unsuccessful. This option is only useful when using an indexed name for <i>name</i> and the -o option.</p> <p>-v Display binary data directly. This option displays columns containing binary data on the standard output. Without this option binary data is displayed as the string *BINARY*.</p> <p>-o <i>name</i> Display the internal representation of the named NIS+ object(s). If <i>name</i> is an indexed name (see nismatch(1)), then each of the matching entry objects is displayed. This option is used to display access rights and other attributes of individual columns.</p>
EXAMPLES	<p>This example displays the contents of the hosts table.</p> <pre>example% niscat -h host.org_dir # cname name addr comment client1 client1 129.144.201.100 Joe Smith crunchy crunchy 129.144.201.44 Jane Smith crunchy softy 129.144.201.44</pre>

The string ***NP*** is returned in those fields where the user has insufficient access rights.

Display the **passwd.org_dir** on the standard output.

```
example% niscat passwd.org_dir
```

Display the contents of table **frodo** and the contents of all tables in its concatenation path.

```
example% niscat -A frodo
```

Display the entries in the table **groups.org_dir** as NIS+ objects. Note that the brackets are protected from the shell by single quotes.

```
example% niscat -o '[|groups.org_dir'
```

Display the table object of the **passwd.org_dir** table.

```
example% niscat -o passwd.org_dir
```

The previous example displays the passwd table object and not the passwd table. The table object include information such as the number of columns, column type, searchable or not searchable separator, access rights, and other defaults.

Display the directory object for **org_dir**, which includes information such as the access rights and replica information.

```
example% niscat -o org_dir
```

ENVIRONMENT

NIS_PATH If this variable is set, and the NIS+ table name is not fully qualified, each directory specified will be searched until the table is found (see **nisdefaults(1)**).

EXIT CODES

niscat returns **0** on success and **1** on failure.

SEE ALSO

nis+(1), **nismatch(1)**, **nistbladm(1)**, **nisdefaults(1)**, **nis_objects(3N)**, **nis_tables(3N)**

NOTES

Columns without values in the table are displayed by two adjacent separator characters.

NAME	nischgrp – change the group owner of a NIS+ object
SYNOPSIS	nischgrp [-AfLP] <i>group name</i> . . .
AVAILABILITY	SUNWnisu
DESCRIPTION	<p>nischgrp changes the group owner of the NIS+ objects or entries specified by <i>name</i> to the specified NIS+ <i>group</i>. Entries are specified using indexed names (see nismatch(1)). If <i>group</i> is not a fully qualified NIS+ group name, it will be resolved using the directory search path (see nisdefaults(1)).</p> <p>The only restriction on changing an object's group owner is that you must have modify permissions for the object.</p> <p>This command will fail if the master NIS+ server is not running.</p>
OPTIONS	<p>-A Modify all entries in all tables in the concatenation path that match the search criterion specified in <i>name</i>. This option implies the -P switch.</p> <p>-f Force the operation and fail silently if it does not succeed.</p> <p>-L Follow links and change the group owner of the linked object or entries rather than the group owner of the link itself.</p> <p>-P Follow the concatenation path within a named table. This option only makes sense when either <i>name</i> is an indexed name or the -L switch is also specified and the named object is a link pointing to entries.</p>
EXAMPLES	<p>The following two examples show how to change the group owner of an object to a group in a different domain, and how to change it to a group in the local domain, respectively.</p> <pre>example% nischgrp newgroup.remote.domain. object example% nischgrp my-buds object</pre> <p>This example shows how to change the group owner for a password entry.</p> <pre>example% nischgrp admins '[uid=99],passwd.org_dir'</pre> <p>In the previous example, admins is a NIS+ group in the same domain.</p> <p>The next two examples change the group owner of the object or entries pointed to by a link, and the group owner of all entries in the hobbies table.</p> <pre>example% nischgrp -L my-buds linkname example% nischgrp my-buds '[,hobbies]</pre>
ENVIRONMENT	<p>NIS_PATH If this variable is set, and the NIS+ name is not fully qualified, each directory specified will be searched until the object is found (see nisdefaults(1)).</p>

EXIT CODES

nischgrp returns **0** on success and **1** on failure.

SEE ALSO

nis+(1), **nischmod(1)**, **nischown(1)**, **nisdefaults(1)**, **nisgrpadm(1)**, **nis_objects(3N)**

NOTES

The NIS+ server will check the validity of the group name prior to effecting the modification.

NAME	nischmod – change access rights on a NIS+ object																								
SYNOPSIS	nischmod [-AfLP] <i>mode name</i> . . .																								
AVAILABILITY	SUNWnisu																								
DESCRIPTION	<p>nischmod changes the access rights (mode) of the NIS+ objects or entries specified by <i>name</i> to <i>mode</i>. Entries are specified using indexed names (see nismatch(1)). Only principals with modify access to an object may change its mode.</p> <p><i>mode</i> has the following form:</p> <p style="padding-left: 40px;"><i>rights</i> [, <i>rights</i>] . . .</p> <p><i>rights</i> has the form:</p> <p style="padding-left: 40px;">[<i>who</i>] <i>op permission</i> [<i>op permission</i>] . . .</p> <p><i>who</i> is a combination of:</p> <table border="0" style="margin-left: 40px;"> <tr><td>n</td><td>Nobody's permissions.</td></tr> <tr><td>o</td><td>Owner's permissions.</td></tr> <tr><td>g</td><td>Group's permissions.</td></tr> <tr><td>w</td><td>World's permissions.</td></tr> <tr><td>a</td><td>All, or owg.</td></tr> </table> <p style="margin-left: 40px;">If <i>who</i> is omitted, the default is a.</p> <p><i>op</i> is one of:</p> <table border="0" style="margin-left: 40px;"> <tr><td>+</td><td>To grant the <i>permission</i>.</td></tr> <tr><td>-</td><td>To revoke the <i>permission</i>.</td></tr> <tr><td>=</td><td>To set the permissions explicitly.</td></tr> </table> <p><i>permission</i> is any combination of:</p> <table border="0" style="margin-left: 40px;"> <tr><td>r</td><td>Read.</td></tr> <tr><td>m</td><td>Modify.</td></tr> <tr><td>c</td><td>Create.</td></tr> <tr><td>d</td><td>Destroy.</td></tr> </table>	n	Nobody's permissions.	o	Owner's permissions.	g	Group's permissions.	w	World's permissions.	a	All, or owg .	+	To grant the <i>permission</i> .	-	To revoke the <i>permission</i> .	=	To set the permissions explicitly.	r	Read.	m	Modify.	c	Create.	d	Destroy.
n	Nobody's permissions.																								
o	Owner's permissions.																								
g	Group's permissions.																								
w	World's permissions.																								
a	All, or owg .																								
+	To grant the <i>permission</i> .																								
-	To revoke the <i>permission</i> .																								
=	To set the permissions explicitly.																								
r	Read.																								
m	Modify.																								
c	Create.																								
d	Destroy.																								
OPTIONS	<p>-A Modify all entries in all tables in the concatenation path that match the search criteria specified in <i>name</i>. This option implies the -P switch.</p> <p>-f Force the operation and fail silently if it does not succeed.</p> <p>-L Follow links and change the permission of the linked object or entries rather than the permission of the link itself.</p> <p>-P Follow the concatenation path within a named table. This option is only applicable when either <i>name</i> is an indexed name or the -L switch is also specified and the named object is a link pointing to an entry.</p>																								

EXAMPLES	<p>This example gives everyone read access to an object. (i.e., access for owner, group, and all).</p> <p style="padding-left: 2em;">example% nischmod a+r <i>object</i></p> <p>This example denies create and modify privileges to group and unauthenticated clients (nobody).</p> <p style="padding-left: 2em;">example% nischmod gn-cm <i>object</i></p> <p>In this example, a complex set of permissions are set for an object.</p> <p style="padding-left: 2em;">example% nischmod o=rmcd,g=rm,w=rc,n=r <i>object</i></p> <p>This example sets the permissions of an entry in the password table so that the group owner can modify them.</p> <p style="padding-left: 2em;">example% nischmod g+m '[uid=55],passwd.org_dir'</p> <p>The next example changes the permissions of a linked object.</p> <p style="padding-left: 2em;">example% nischmod -L w+mr <i>linkname</i></p>
ENVIRONMENT	<p>NIS_PATH If this variable is set, and the NIS+ name is not fully qualified, each directory specified will be searched until the object is found (see nisdefaults(1)).</p>
EXIT CODES	nischmod returns 0 on success and 1 on failure.
SEE ALSO	chmod(1) , nis+(1) , nischgrp(1) , nischown(1) , nisdefaults(1) , nis_objects(3N)
NOTES	Unlike the system chmod(1) command, this command does not accept an octal notation.

NAME	nischown – change the owner of a NIS+ object
SYNOPSIS	nischown [-AfLP] <i>owner name</i> . . .
AVAILABILITY	SUNWnisu
DESCRIPTION	<p>nischown changes the owner of the NIS+ objects or entries specified by <i>name</i> to <i>owner</i>. Entries are specified using indexed names (see nismatch(1)). If <i>owner</i> is not a fully qualified NIS+ principal name (see nisaddcred(1M)), the default domain (see nisdefaults(1)) will be appended to it.</p> <p>The only restriction on changing an object's owner is that you must have modify permissions for the object. Note: If you are the current owner of an object and you change ownership, you may not be able to regain ownership unless you have modify access to the new object.</p> <p>The command will fail if the master NIS+ server is not running.</p>
OPTIONS	<p>-A Modify all entries in all tables in the concatenation path that match the search criteria specified in <i>name</i>. It implies the -P option.</p> <p>-f Force the operation and fail silently if it does not succeed.</p> <p>-L Follow links and change the owner of the linked object or entries rather than the owner of the link itself.</p> <p>-P Follow the concatenation path within a named table. This option is only meaningful when either <i>name</i> is an indexed name or the -L option is also specified and the named object is a link pointing to entries.</p>
EXAMPLES	<p>The following two examples show how to change the owner of an object to a principal in a different domain, and to change it to a principal in the local domain, respectively.</p> <pre>example% nischown bob.remote.domain. object example% nischown skippy object</pre> <p>The next example shows how to change the owner of an entry in the passwd table.</p> <pre>example% nischown bob.remote.domain. '[uid=99],passwd.org_dir'</pre> <p>This example shows how to change the object or entries pointed to by a link.</p> <pre>example% nischown -L skippy linkname</pre>
ENVIRONMENT	<p>NIS_PATH If this variable is set, and the NIS+ name is not fully qualified, each directory specified will be searched until the object is found (see nisdefaults(1)).</p>

EXIT CODES

nischown returns **0** on success and **1** on failure.

SEE ALSO

nis+(1), **nischgrp(1)**, **nischmod(1)**, **nischttl(1)**, **nisdefaults(1)**, **nisaddcred(1M)**,
nis_objects(3N)

NOTES

The NIS+ server will check the validity of the name before making the modification.

NAME	nischtctl – change the time to live value of a NIS+ object
SYNOPSIS	nischtctl [-A fLP] <i>time name</i> . . .
AVAILABILITY	SUNWnisu
DESCRIPTION	<p>nischtctl changes the time to live value (ttl) of the NIS+ objects or entries specified by <i>name</i> to <i>time</i>. Entries are specified using indexed names (see nismatch(1)).</p> <p>The time to live value is used by object caches to expire objects within their cache. When an object is read into the cache, this value is added to the current time in seconds yielding the time when the cached object would expire. The object may be returned from the cache until the current time is earlier than the calculated expiration time. When the expiration time has been reached, the object will be flushed from the cache.</p> <p>The time to live <i>time</i> may be specified in seconds or in days, hours, minutes, seconds format. The latter format uses a suffix letter of d, h, m, or s to identify the units of time. See the examples below for usage.</p> <p>The command will fail if the master NIS+ server is not running.</p>
OPTIONS	<p>-A Modify all tables in the concatenation path that match the search criterion specified in <i>name</i>. This option implies the -P switch.</p> <p>-f Force the operation and fail silently if it does not succeed.</p> <p>-L Follow links and change the time to live of the linked object or entries rather than the time to live of the link itself.</p> <p>-P Follow the concatenation path within a named table. This option only makes sense when either <i>name</i> is an indexed name or the -L switch is also specified and the named object is a link pointing to entries.</p>
EXAMPLES	<p>The following example shows how to change the ttl of an object using the seconds format and the days, hours, minutes, seconds format. The ttl of the second object is set to 1 day and 12 hours.</p> <pre>example% nischtctl 184000 object example% nischtctl 1d12h object</pre> <p>This example shows how to change the ttl for a password entry.</p> <pre>example% nischtctl 1h30m '[uid=99],passwd.org_dir'</pre> <p>The next two examples change the ttl of the object or entries pointed to by a link, and the ttl of all entries in the hobbies table.</p> <pre>example% nischtctl -L 12h linkname example% nischtctl 3600 '[],hobbies</pre>
ENVIRONMENT	<p>NIS_PATH If this variable is set, and the NIS+ name is not fully qualified, each directory specified will be searched until the object is found (see nisdefaults(1)).</p>

EXIT CODES

nischttl returns **0** on success and **1** on failure.

SEE ALSO

nis+(1), **nischgrp(1)**, **nischmod(1)**, **nischown(1)**, **nisdefaults(1)**, **nis_objects(3N)**

NOTES

Setting a high **tfl** value allows objects to stay persistent in caches for a longer period of time and can improve performance. However, when an object changes, in the worst case, the number of seconds in this attribute must pass before that change is visible to all clients. Setting a **tfl** value of **0** means that the object should not be cached at all.

A high **tfl** value is a week, a low value is less than a minute. Password entries should have **tfl** values of about 12 hours (easily allows one password change per day), entries in the RPC table can have **tfl** values of several weeks (this information is effectively unchanging).

Only directory and group objects are cached in this implementation.

NAME	nisdefaults – display NIS+ default values
SYNOPSIS	nisdefaults [-adghprstv]
AVAILABILITY	SUNWnisu
DESCRIPTION	nisdefaults prints the default values that are returned by calls to the NIS+ local name functions (see nis_local_names(3N)). With no options specified, all defaults will be printed in a verbose format. With options, only that option is displayed in a terse form suitable for shell scripts. See the example below.
OPTIONS	<p>-a Print all defaults in a terse format.</p> <p>-d Print the default domain name.</p> <p>-g Print the default group name.</p> <p>-h Print the default host name.</p> <p>-p Print the default principal name.</p> <p>-r Print the default access rights with which new objects will be created.</p> <p>-s Print the default directory search path.</p> <p>-t Print the default time to live value.</p> <p>-v Print the defaults in a verbose format. This prepends an identifying string to the output.</p>
EXAMPLES	<p>The following prints the NIS+ defaults for a root process on machine example in the foo.bar. domain.</p> <pre>example# nisdefaults Principal Name : example.foo.bar. Domain Name : foo.bar. Host Name : example.foo.bar. Group Name : Access Rights : ----rmcdr---r--- Time to live : 12:00:00 Search Path : foo.bar.</pre> <p>This example sets a variable in a shell script to the default domain.</p> <pre>DOMAIN='nisdefaults -d'</pre> <p>This example prints out the default time to live in a verbose format.</p> <pre>example% nisdefaults -tv Time to live : 12:00:00</pre> <p>This example prints out the time to live in the terse format:</p> <pre>example% nisdefaults -t 43200</pre>

ENVIRONMENT

Several environment variables affect the defaults associated with a process.

NIS_DEFAULTS This variable contains a defaults string that will override the NIS+ standard defaults. The defaults string is a series of tokens separated by colons. These tokens represent the default values to be used for the generic object properties. All of the legal tokens are described below.

ttl=*time*

This token sets the default time to live for objects that are created. The value *time* is specified in the format as defined by the **nischttl**(1) command. The default value is **12** hours.

owner=*ownername*

This token specifies that the NIS+ principal *ownername* should own created objects. The default for this value is the principal who is executing the command.

group=*groupname*

This token specifies that the group *groupname* should be the group owner for created objects. The default is **NULL**.

access=*rights*

This token specifies the set of access rights that are to be granted for created objects. The value *rights* is specified in the format as defined by the **nischmod**(1) command. The default value is
 ----**rmcdr**---**r**---

NIS_GROUP This variable contains the name of the local NIS+ group. If the name is not fully qualified, the default domain will be appended to it.

NIS_PATH This variable overrides the default NIS+ directory search path. It contains an ordered list of directories separated by ':' (colon) characters. The '\$' (dollar sign) character is treated specially. Directory names that end in '\$' have the default domain appended to them, and a '\$' by itself is replaced by the list of directories between the default domain and the global root that are at least two levels deep. The default NIS+ directory search path is '\$'.

Refer to the **Name Expansion** subsection in **nis**+(1) for more details.

SEE ALSO

nis+(1), **nis_local_names**(3N)

NAME	niserror – display NIS+ error messages
SYNOPSIS	niserror <i>error-num</i>
AVAILABILITY	SUNWnisu
DESCRIPTION	niserror prints the NIS+ error associated with status value <i>error-num</i> on the standard output. It is used by shell scripts to translate NIS+ error numbers that are returned into text messages.
EXAMPLES	The following example prints the error associated with the error number 20 : example% niserror 20 Not Found, no such name
SEE ALSO	nis+(1) , nis_error(3N)

NAME	nisgrpadm – NIS+ group administration command
SYNOPSIS	nisgrpadm -a -r -t] [-s] <i>group principal</i> . . nisgrpadm -c -d -l [-M] [-s] <i>group</i>
AVAILABILITY	SUNWnisu
DESCRIPTION	<p>nisgrpadm is used to administer NIS+ groups. This command administers both groups and the groups' membership lists. nisgrpadm can create, destroy, or list NIS+ groups. nisgrpadm can be used to administer a group's membership list. It can add or delete principals to the group, or test principals for membership in the group.</p> <p>The names of NIS+ groups are syntactically similar to names of NIS+ objects but they occupy a separate namespace. A group named "a.b.c.d." is represented by a NIS+ group object named "a.groups_dir.b.c.d."; the functions described here all expect the name of the group, not the name of the corresponding group object.</p> <p>There are three types of group members:</p> <ul style="list-style-type: none"> • An <i>explicit</i> member is just a NIS+ principal-name, for example "wickedwitch.west.oz." • An <i>implicit</i> ("domain") member, written "*.west.oz.", means that all principals in the given domain belong to this member. No other forms of wildcarding are allowed: "wickedwitch.*.oz." is invalid, as is "wickedwitch.west.*.". Note that principals in subdomains of the given domain are <i>not</i> included. • A <i>recursive</i> ("group") member, written "@cowards.oz.", refers to another group; all principals that belong to that group are considered to belong here. <p>Any member may be made <i>negative</i> by prefixing it with a minus sign ('-'). A group may thus contain explicit, implicit, recursive, negative explicit, negative implicit, and negative recursive members.</p> <p>A principal is considered to belong to a group if it belongs to at least one non-negative group member of the group and belongs to no negative group members.</p>
OPTIONS	<p>-a Add the list of NIS+ principals specified to <i>group</i>. The principal name should be fully qualified.</p> <p>-c Create <i>group</i> in the NIS+ namespace. The NIS+ group name should be fully qualified.</p> <p>-d Destroy (remove) <i>group</i> from the namespace.</p> <p>-l List the membership list of the specified <i>group</i>. (See -M.)</p> <p>-M Master server only. Send the lookup to the master server of the named data. This guarantees that the most up to date information is seen at the possible expense that the master server may be busy. Note that the -M flag is applicable only with the -l flag.</p> <p>-r Remove the list of principals specified from <i>group</i>. The principal name should be fully qualified.</p>

- s Work silently. Results are returned using the exit status of the command. This status can be translated into a text string using the **niserror(1)** command.
- t Display whether the principals specified are members in *group*.

EXAMPLES**Administering
Groups**

This example shows how to create a group in the **foo.com.** domain.

```
example% nisgrpadm -c my_buds.foo.com.
```

This example shows how to remove the group from the current domain.

```
example% nisgrpadm -d fred's_group
```

**Administering
Members**

This example shows how one would add two principals, **bob** and **betty** to the group **my_buds.foo.com.**

```
example% nisgrpadm -a my_buds.foo.com. bob.bar.com. betty.foo.com.
```

This example shows how to remove **betty** from **fred's_group**.

```
example% nisgrpadm -r fred's_group betty.foo.com.
```

ENVIRONMENT

NIS_PATH If this variable is set, and the NIS+ group name is not fully qualified, each directory specified will be searched until the group is found (see **nisdefaults(1)**).

SEE ALSO

nis+(1), **nischgrp(1)**, **nisdefaults(1)**, **niserror(1)**, **nis_groups(3N)**

DIAGNOSTICS

NIS_SUCCESS On success, this command returns an exit status of **0**.

NIS_PERMISSION When you do not have the needed access right to change the group, the command returns this error.

NIS_NOTFOUND This is returned when the group does not exist.

NIS_TRYAGAIN This error is returned when the server for the group's domain is currently checkpointing or otherwise in a read-only state. The command should be retried at a later date.

NIS_MODERROR This error is returned when the group was modified by someone else during the execution of the command. Reissue the command and optionally recheck the group's membership list.

NOTES

Principal names *must* be fully qualified, whereas groups can be abbreviated on all operations *except* create.

NAME	nisl - symbolically link NIS+ objects
SYNOPSIS	nisl [-L] [-D defaults] <i>name linkname</i>
AVAILABILITY	SUNWnisu
DESCRIPTION	The nisl command links a NIS+ object named <i>name</i> to a NIS+ name <i>linkname</i> . If <i>name</i> is an indexed name (see nismatch (1)), the link points to entries within a NIS+ table. Clients wishing to look up information in the name service can use the FOLLOW_LINKS flag to force the client library to follow links to the name they point to. Further, all of the NIS+ administration commands accept the -L switch indicating they should follow links (see nis_names (3N) for a description of the FOLLOW_LINKS flag).
OPTIONS	<p>-L When present, this option specifies that this command should follow links. If <i>name</i> is itself a link, then this command will follow it to the linked object that it points to. The new link will point to that linked object rather than to <i>name</i>.</p> <p>-D defaults Specify a different set of defaults to be used for the creation of the link object. The <i>defaults</i> string is a series of tokens separated by colons. These tokens represent the default values to be used for the generic object properties. All of the legal tokens are described below.</p> <p>ttl=time This token sets the default time to live for objects that are created by this command. The value <i>time</i> is specified in the format as defined by the nischttl(1) command. The default is 12 hours.</p> <p>owner=ownername This token specifies that the NIS+ principal <i>ownername</i> should own the created object. The default for this value is the principal who is executing the command.</p> <p>group=groupname This token specifies that the group <i>groupname</i> should be the group owner for the object that is created. The default is NULL.</p> <p>access=rights This token specifies the set of access rights that are to be granted for the given object. The value <i>rights</i> is specified in the format as defined by the nischmod(1) command. The default value is ---rmcdr---r---.</p>

EXAMPLES

In this example we create a link in the domain **foo.com**. named **hosts** that points to the object **hosts.bar.com**.

```
example% nisln hosts.bar.com. hosts.foo.com.
```

In this example we make a link *example.sun.com*. that points to an entry in the hosts table in *eng.sun.com*.

```
example% nisln '[name=example],hosts.eng.sun.com.' example.sun.com.
```

ENVIRONMENT

NIS_PATH If this variable is set, and the NIS+ name is not fully qualified, each directory specified will be searched until the object is found (see **nisdefaults(1)**).

EXIT CODES .LP

nisln returns **0** on success and **1** on failure.

SEE ALSO

nisdefaults(1), **nismatch(1)**, **nismrm(1)**, **nistbladm(1)**, **nis_names(3N)**, **nis_tables(3N)**

NOTES

When creating the link, **nisln** verifies that the linked object exists. Once created, the linked object may be deleted or replaced and the link will not be affected. At that time the link will become invalid and attempts to follow it will return **NIS_LINKNAMEERROR** to the client. When the path attribute in tables specifies a link rather than another table, the link will be followed if the flag **FOLLOW_LINKS** was present in the call to **nis_list()** (see **nis_tables(3N)**) and ignored if the flag is not present. If the flag is present and the link is no longer valid, a warning is sent to the system logger and the link is ignored.

NAME	nisl – list the contents of a NIS+ directory
SYNOPSIS	nisl [-dglLmMR] [<i>name</i> . . .]
AVAILABILITY	SUNWnisu
DESCRIPTION	For each <i>name</i> that is a NIS+ directory, nisl lists the contents of the directory. For each <i>name</i> that is a NIS+ object other than a directory, nisl simply echos the name. If no <i>name</i> is specified, the first directory in the search path (see nisdefaults(1)) is listed.
OPTIONS	<p>-d Treat NIS+ directories like other NIS+ objects, rather than listing their contents.</p> <p>-g Display group owner instead of owner when listing in long format.</p> <p>-l List in long format. This option displays additional information about the objects such as their type, creation time, owner, and access rights. The access rights are listed in the following order in long mode: nobody, owner, group owner, and world.</p> <p>-L This option specifies that links are to be followed. If <i>name</i> actually points to a link, it is followed to the linked object.</p> <p>-m Display modification time instead of creation time when listing in long format.</p> <p>-M Master only. This specifies that information is to be returned from the master server of the named object. This guarantees that the most up to date information is seen at the possible expense that the master server may be busy.</p> <p>-R List directories recursively. This option will reiterate the list for each subdirectory found in the process of listing each <i>name</i> .</p>
ENVIRONMENT	<p>NIS_PATH If this variable is set, and the NIS+ name is not fully qualified, each directory specified will be searched until the object is found (see nisdefaults(1)).</p>
EXIT CODES	nisl returns 0 on success and 1 on failure.
SEE ALSO	nisdefaults(1) , nisrpadm(1) , nismatch(1) , nistbladm(1) , nis_objects(3N)

NAME	nismatch, nisgrep – utilities for searching NIS+ tables
SYNOPSIS	<p>nismatch [-AchMoPv] <i>key tablename</i></p> <p>nismatch [-AchMoPv] <i>colname=key... tablename</i></p> <p>nismatch [-AchMoPv] <i>indexedname</i></p> <p>nisgrep [-AchMov] <i>keypat tablename</i></p> <p>nisgrep [-AchMov] <i>colname=keypat... tablename</i></p>
AVAILABILITY	SUNWnisu
DESCRIPTION	<p>nismatch and nisgrep can be used to search NIS+ tables. The command nisgrep differs from the nismatch command in its ability to accept regular expressions <i>keypat</i> for the search criteria rather than simple text matches.</p> <p>Because nisgrep uses a callback function, it is not constrained to searching only those columns that are specifically made searchable at the time of table creation. This makes it more flexible, but slower, than nismatch.</p> <p>In nismatch, the server does the searching; whereas in nisgrep, the server returns all the readable entries and then the client does the pattern-matching.</p> <p>In both commands, the parameter <i>tablename</i> is the NIS+ name of the table to be searched. If only one key or key pattern is specified without the column name, then it is applied searching the first column. Specific named columns can be searched by using the <i>colname=key</i> syntax. When multiple columns are searched, only entries that match in all columns are returned. This is the equivalent of a logical join operation.</p> <p>nismatch accepts an additional form of search criteria, <i>indexedname</i>, which is a NIS+ indexed name of the form:</p> <p style="padding-left: 40px;"><i>[colname=value, ...],tablename</i></p>
OPTIONS	<p>-A All data. Return the data within the table and all of the data in tables in the initial table's concatenation path.</p> <p>-c Print only a count of the number of entries that matched the search criteria.</p> <p>-h Display a header line before the matching entries that contains the names of the table's columns</p> <p>-M Master server only. Send the lookup to the master server of the named data. This guarantees that the most up to date information is seen at the possible expense that the master server may be busy.</p> <p>-o Display the internal representation of the matching NIS+ object(s).</p> <p>-P Follow concatenation path. Specify that the lookup should follow the concatenation path of a table if the initial search is unsuccessful.</p> <p>-v Verbose. Do not suppress the output of binary data when displaying matching entries. Without this option binary data is displayed as the string *BINARY*.</p>

RETURN VALUES	<p>0 Returns 0 when it successfully matches some entries.</p> <p>1 Returns 1 when it successfully searches the table and no matches are found.</p> <p>2 Returns 2 when an error condition occurs. An error message is also printed.</p>
EXAMPLES	<p>This example searches a table named passwd in the org_dir subdirectory of the zotz.com domain. It returns the entry that has the username of skippy. In this example, all the work is done on the server.</p> <p style="padding-left: 40px;">example% nismatch name=skippy passwd.org_dir.zotz.com.</p> <p>This example is similar to the one above except that it uses nisgrep to find all users in the table named passwd that are using either ksh(1) or cs(1).</p> <p style="padding-left: 40px;">example% nisgrep 'shell=[ck]sh' passwd.org_dir.zotz.com.</p>
ENVIRONMENT	<p>NIS_PATH If this variable is set, and the NIS+ table name is not fully qualified, each directory specified will be searched until the table is found (see nisdefaults(1)).</p>
SEE ALSO	<p>niscat(1), nisdefaults(1), nisls(1), nistbladm(1), nis_objects(3N)</p>
DIAGNOSTICS	<p>No memory An attempt to allocate some memory for the search failed.</p> <p>tablename is not a table The object with the name <i>tablename</i> was not a table object.</p> <p>Can't compile regular expression The regular expression in <i>keypat</i> was malformed.</p> <p>column not found: colname The column named <i>colname</i> does not exist in the table named <i>tablename</i>.</p>

NAME	nismkdir – create NIS+ directories
SYNOPSIS	nismkdir [-D defaults] [-m hostname -s hostname] <i>dirname</i>
AVAILABILITY	SUNWnisu
DESCRIPTION	<p>The nismkdir command creates new NIS+ subdirectories within an existing domain. It can also be used to create replicated directories. Without options, this command will create a subdirectory with the same master and the replicas as its parent directory.</p> <p>It is advisable to use nisserver(1M) to create an NIS+ domain which consists of the specified directory along with the org_dir and group_dir subdirectories.</p> <p>The two primary aspects that are controlled when making a directory are its access rights, and its degree of replication.</p> <p><i>dirname</i> is the fully qualified NIS+ name of the directory that has to be created.</p>
OPTIONS	<p>-D defaults Specify a different set of defaults to be used when creating new directories. The <i>defaults</i> string is a series of tokens separated by colons. These tokens represent the default values to be used for the generic object properties. All of the legal tokens are described below.</p> <p>ttl=time This token sets the default time to live for objects that are created by this command. The value <i>time</i> is specified in the format as defined by the nischttl(1) command. The default value is 12h (12 hours).</p> <p>owner=ownername This token specifies that the NIS+ principal <i>ownername</i> should own the created object. The default for this value is the principal who is executing the command.</p> <p>group=groupname This token specifies that the group <i>groupname</i> should be the group owner for the object that is created. The default value is NULL.</p> <p>access=rights This token specifies the set of access rights that are to be granted for the given object. The value <i>rights</i> is specified in the format as defined by the nischmod(1) command. The default value is ----rmcdr---r---.</p> <p>-m hostname If the directory named by <i>dirname</i> does not exist, then a new directory that is <i>not</i> replicated is created with host <i>hostname</i> as its master server. If the directory name by <i>dirname</i> does exist, then the host named by <i>hostname</i> is made its master server.</p>

-s *hostname* Specify that the host *hostname* will be a replica for an existing directory named *dirname*.

RETURN VALUES **0** This command returns 0 if successful and 1 otherwise.

EXAMPLES To create a new directory **bar** under the **foo.com.** domain that shares the same master and replicas as the **foo.com.** directory one would use the command:

example% nismkdir bar.foo.com.

To create a new directory *bar.foo.com.* that is not replicated under the **foo.com.** domain one would use the command:

example% nismkdir -m myhost.foo.com. bar.foo.com.

To add a replica server of the *bar.foo.com.* directory, one would use the command:

example% nismkdir -s replica.foo.com. bar.foo.com.

ENVIRONMENT **NIS_DEFAULTS** This variable contains a defaults string that will override the NIS+ standard defaults. If the **-D** switch is used those values will then override both the **NIS_DEFAULTS** variable and the standard defaults.

NIS_PATH If this variable is set, and the NIS+ directory name is not fully qualified, each directory specified will be searched until the directory is found (see **nisdefaults(1)**).

SEE ALSO **nis+(1), nischmod(1), nisdefaults(1), nisl(1), nisrmdir(1), nisserver(1M)**

NOTES A host that serves a NIS+ directory *must be* a NIS+ client in a directory above the one it is serving. The exceptions to this rule are the root NIS+ servers which are both clients and servers of the same NIS+ directory.

When the host's default domain is different than the default domain on the client where the command is executed, the hostname supplied as an argument to the **-s** or **-m** options must be fully qualified.

NAME	nispasswd – change NIS+ password information
SYNOPSIS	nispasswd [-ghs] [-D <i>domainname</i>] [<i>username</i>] nispasswd -a nispasswd -D <i>domainname</i>] [-d [<i>username</i>]] nispasswd [-l] [-f] [-n <i>min</i>] [-x <i>max</i>] [-w <i>warn</i>] [-D <i>domainname</i>] <i>username</i>
AVAILABILITY	SUNWnisu
DESCRIPTION	<p>nispasswd changes a password, gecost (finger) field (-g option), home directory (-h option), or login shell (-s option) associated with the <i>username</i> (invoker by default) in the NIS+ passwd table.</p> <p>Additionally, the command can be used to view or modify aging information associated with the user specified if the invoker has the right NIS+ privileges.</p> <p>nispasswd uses secure RPC to communicate with the NIS+ server, and therefore, never sends unencrypted passwords over the communication medium.</p> <p>nispasswd does not read or modify the local password information stored in the <i>/etc/passwd</i> and <i>/etc/shadow</i> files.</p> <p>When used to change a password, nispasswd prompts non-privileged users for their old password. It then prompts for the new password twice to forestall typing mistakes. When the old password is entered, nispasswd checks to see if it has “aged” sufficiently. If “aging” is insufficient, nispasswd terminates; see getspnam(3C).</p> <p>The old password is used to decrypt the username’s secret key. If the password does not decrypt the secret key, nispasswd prompts for the old secure-RPC password. It uses this password to decrypt the secret key. If this fails, it gives the user one more chance. The old password is also used to ensure that the new password differs from the old by at least three characters. Assuming aging is sufficient, a check is made to ensure that the new password meets construction requirements described below. When the new password is entered a second time, the two copies of the new password are compared. If the two copies are not identical, the cycle of prompting for the new password is repeated twice. The new password is used to re-encrypt the user’s secret key. Hence, it also becomes their secure-RPC password.</p> <p>Passwords must be constructed to meet the following requirements:</p> <ul style="list-style-type: none"> • Each password must have at least six characters. Only the first eight characters are significant. • Each password must contain at least two alphabetic characters and at least one numeric or special character. In this case, “alphabetic” refers to all upper or lower case letters. • Each password must differ from the user’s login <i>username</i> and any reverse or circular shift of that login <i>username</i>. For comparison purposes, an upper case letter and its corresponding lower case letter are equivalent.

- New passwords must differ from the old by at least three characters. For comparison purposes, an upper case letter and its corresponding lower case letter are equivalent.

Network administrators, who own the NIS+ password table, may change any password attributes if they establish their credentials (see **keylogin(1)**) before invoking **nispasswd**. Hence, **nispasswd** does not prompt these privileged-users for the old password and they are not forced to comply with password aging and password construction requirements.

Any user may use the **-d** option to display password attributes for his or her own login name. The format of the display will be:

```
username status mm/dd/yy min max warn
```

or, if password aging information is not present,

```
username status
```

where

<i>username</i>	The login ID of the user.
<i>status</i>	The password status of <i>username</i> : "PS" stands for password exists or locked, "LK" stands for locked, and "NP" stands for no password.
<i>mm/dd/yy</i>	The date password was last changed for <i>username</i> . (Note that all password aging dates are determined using Greenwich Mean Time and, therefore, may differ by as much as a day in other time zones.)
<i>min</i>	The minimum number of days required between password changes for <i>username</i> .
<i>max</i>	The maximum number of days the password is valid for <i>username</i> .
<i>warn</i>	The number of days relative to <i>max</i> before the password expires that the <i>username</i> will be warned.

OPTIONS

-g	Change the <i>gecos</i> (finger) information.
-h	Change the home directory.
-s	Change the login shell. By default, only the NIS+ administrator can change the login shell. User will be prompted for the new login shell.
-a	Show the password attributes for all entries. This will show only the entries in the NIS+ passwd table in the local domain that the invoker is authorized to "read".
-d [username]	Display password attributes for the caller or the user specified if the invoker has the right privileges.
-l	Locks the password entry for <i>username</i> . Subsequently, login(1) would disallow logins with this NIS+ password entry.
-f	Force the user to change password at the next login by expiring the password for <i>username</i> .

- n *min*** Set minimum field for *username*. The *min* field contains the minimum number of days between password changes for *username*. If *min* is greater than *max*, the user may not change the password. Always use this option with the **-x** option, unless *max* is set to -1 (aging turned off). In that case, *min* need not be set.
- x *max*** Set maximum field for *username*. The *max* field contains the number of days that the password is valid for *username*. The aging for *username* will be turned off immediately if *max* is set to -1. If it is set to 0, then the user is forced to change the password at the next login session and aging is turned off.
- w *warn*** Set *warn* field for *username*. The *warn* field contains the number of days before the password expires that the user will be warned whenever he or she attempts to login.
- D *domainname*** Consult the **passwd.org_dir** table in *domainname*. If this option is not specified, the default domainname returned by **nis_local_directory()** will be used. This domainname is the same as that returned by **domainname(1M)**.

SEE ALSO

keylogin(1), **login(1)**, **nis+(1)**, **nistbladm(1)**, **passwd(1)**, **domainname(1M)**, **getpwnam(3C)**, **getspnam(3C)**, **nsswitch.conf(4)**, **passwd(4)**, **shadow(4)**

DIAGNOSTICS

The **nispasswd** command exits with one of the following values:

- 0** SUCCESS.
- 1** Permission denied.
- 2** Invalid combination of options.
- 3** Unexpected failure. NIS+ passwd table unchanged.
- 4** NIS+ passwd table missing.
- 5** NIS+ is busy. Try again later.
- 6** Invalid argument to option.
- 7** Aging is disabled.

NOTES

The login program, file access display programs (for example, **'ls -l'**) and network programs that require user passwords (for example, **rlogin(1)**, **ftp(1)**, etc.) use the standard **getpwnam(3C)** and **getspnam(3C)** interfaces to get password information. These programs will get the NIS+ password information, that is modified by **nispasswd**, only if the **passwd:** entry in the **/etc/nsswitch.conf** file includes **to nisplus**. See **nsswitch.conf(4)** for more details.

NAME	nism – remove NIS+ objects from the namespace
SYNOPSIS	nism [-if] <i>name</i> ...
AVAILABILITY	SUNWnisu
DESCRIPTION	The nism command removes NIS+ objects named <i>name</i> from the NIS+ namespace. This command will fail if the NIS+ master server is not running.
OPTIONS	-i Interactive mode. Like the system rm (1) command the nism command will ask for confirmation prior to removing an object. If the name specified by <i>name</i> is a non-fully qualified name this option is forced on. This prevents the removal of unexpected objects. -f Force. The removal is attempted, and if it fails for permission reasons, a nischmod (1) is attempted and the removal retried. If the command fails, it fails silently.
EXAMPLES	Remove the objects <i>foo</i> , <i>bar</i> , and <i>baz</i> from the namespace. example% nism foo bar baz
ENVIRONMENT	NIS_PATH If this variable is set, and the NIS+ name is not fully qualified, each directory specified will be searched until the object is found (see nisdefaults (1)).
EXIT CODES	nism returns 0 on success and 1 on failure.
SEE ALSO	nis+ (1), nischmod (1), nisdefaults (1), nismdir (1), nistbladm (1), rm (1)
NOTES	This command will not remove directories (see nismdir (1)) nor will it remove non-empty tables (see nistbladm (1)).

NAME	nismkdir – remove NIS+ directories
SYNOPSIS	nismkdir [-if] [-s <i>hostname</i>] <i>dirname</i>
AVAILABILITY	SUNWnisu
DESCRIPTION	<p>nismkdir deletes existing NIS+ subdirectories. It can remove a directory outright, or simply remove replicas from serving a directory.</p> <p>This command modifies the object that describes the directory <i>dirname</i>, and then notifies each replica to remove the directory named <i>dirname</i>. If the notification of any of the affected replicas fails, the directory object is returned to its original state unless the -f option is present.</p> <p>This command will fail if the NIS+ master server is not running.</p>
OPTIONS	<p>-i Interactive mode. Like the system rm(1) command the nismkdir command will ask for confirmation prior to removing a directory. If the name specified by <i>dirname</i> is a non-fully qualified name this option is forced on. This prevents the removal of unexpected directories.</p> <p>-f Force the command to succeed even though it may not be able to contact the affected replicas. This option should be used when a replica is known to be down and will not be able to respond to the removal notification. When the replica is finally rebooted it will read the updated directory object, note that it is no longer a replica for that directory, and stop responding to lookups on that directory. Cleanup of the files that held the now removed directory can be accomplished manually by removing the appropriate files in the <i>/var/nis</i> directory (see nisfiles(4) for more information).</p> <p>-s <i>hostname</i> Specify that the host <i>hostname</i> should be removed as a replica for the directory named <i>dirname</i>. If this option is not present <i>all</i> replicas and the master server for a directory are removed and the directory is removed from the namespace.</p>
RETURN VALUES	This command returns 0 if it is successful, and 1 otherwise.
EXAMPLES	<p>To remove a directory bar under the foo.com. domain, one would use the command:</p> <p style="padding-left: 40px;">example% nismkdir bar.foo.com.</p> <p>To remove a replica that is serving directory bar.foo.com. one would use the command:</p> <p style="padding-left: 40px;">example% nismkdir -s replica.foo.com. bar.foo.com.</p> <p>To force the removal of directory bar.foo.com. from the namespace, one would use the command:</p>

example% nismdir -f bar.foo.com.

ENVIRONMENT

NIS_PATH If this variable is set, and the NIS+ directory name is not fully qualified, each directory specified will be searched until the directory is found (see **nisdefaults(1)**).

SEE ALSO

nis+(1), nisdefaults(1), nism(1), nisfiles(4)

NAME	nistbladm – NIS+ table administration command
SYNOPSIS	<p>nistbladm -a -A [-D defaults] <i>colname=value ... tablename</i></p> <p>nistbladm -a -A [-D defaults] <i>indexedname</i></p> <p>nistbladm -c [-D defaults] [-p path] [-s sep] <i>type colname=[flags][,access] ... tablename</i></p> <p>nistbladm -d <i>tablename</i></p> <p>nistbladm -m <i>colname=value ... indexedname</i></p> <p>nistbladm -r -R [<i>colname=value ...</i>] <i>tablename</i></p> <p>nistbladm -r -R <i>indexedname</i></p> <p>nistbladm -u [-p path] [-s sep] [-t type] [<i>colname=access ...</i>] <i>tablename</i></p>
AVAILABILITY	SUNWnisu
DESCRIPTION	<p>The nistbladm command is used to administer NIS+ tables. There are five primary operations that it performs: creating and deleting tables, adding entries to, modifying entries within, and removing entries from tables.</p> <p>Though NIS+ does not place restrictions on the size of tables or entries, the size of data has an impact on the performance and the disk space requirements of the NIS+ server. NIS+ is not designed to store huge pieces of data, such as files; instead pointer to files should be stored in NIS+.</p> <p>NIS+ design is optimized to support 10,000 objects with a total size of 10M bytes. If the requirements exceed the above, it is suggested that the domain hierarchy be created, or the data stored in the tables be pointers to the actual data, instead of the data itself.</p> <p>When creating tables, a table type, <i>type</i>, and a list of column definitions must be provided.</p> <p><i>type</i> is a string that is stored in the table and later used by the service to verify that entries being added to it are of the correct type.</p> <p>Syntax for column definitions is:</p> <p><i>colname=[flags][,access]</i></p> <p><i>flags</i> is a combination of:</p> <ul style="list-style-type: none"> S Searchable. Specifies that searches can be done on the column's values (see nismatch(1)). I Case-insensitive (only makes sense in combination with S). Specifies that searches should ignore case. C Crypt. Specifies that the column's values should be encrypted. B Binary data (does not make sense in combination with S). If not set, the column's values are expected to be null terminated ASCII strings.

X XDR encoded data (only makes sense in combination with **B**).

access is specified in the format as defined by the **nischmod**(1) command.

When manipulating entries, this command takes two forms of entry name. The first uses a series of space separated *colname=value* pairs that specify column values in the entry. The second is a NIS+ indexed name, *indexedname*, of the form:

[*colname=value, ...*],*tablename*

OPTIONS

- a | A** Add entries to a NIS+ table. The difference between the lowercase 'a' and the uppercase 'A' is in the treatment of preexisting entries. The entry's contents are specified by the *column=value* pairs on the command line. Note: Values for *all* columns must be specified when adding entries to a table.

Normally, NIS+ reports an error if an attempt is made to add an entry to a table that would overwrite an entry that already exists. This prevents multiple parties from adding duplicate entries and having one of them get overwritten. If you wish to force the add, the uppercase 'A' specifies that the entry is to be added, even if it already exists. This is analogous to a modify operation on the entry.
- c** Create a table named *tablename* in the namespace. The table that is created must have at least one column and at least one column must be searchable.
- d *tablename***
Destroy the table named *tablename*. The table that is being destroyed must be empty. The table's contents can be deleted with the **-R** option below.
- m** Modify an entry in the table that is specified by *indexedname*. Note: Since it is possible to modify the value in a column that would change the indexed name for an entry both the column value pair and the indexed name are required. It uses the indexed name to look up the entry, modify it, and write it back with the new value. The indexed name must uniquely identify a single entry.
- r | R** Remove entries from a table. The entry is specified by either a series of *column=value* pairs on the command line, or an indexed name that is specified as *entryname*. The difference between the interpretation of the lowercase 'r' versus the uppercase 'R' is in the treatment of non-unique entry specifications. Normally the NIS+ server will disallow an attempt to remove an entry when the search criterion specified for that entry resolves to more than one entry in the table. However, it is sometimes desirable to remove more than one entry, as when you are attempting to remove all of the entries from a table. In this case, using the uppercase 'R' will force the NIS+ server to remove all entries matching the passed search criterion. If that criterion is null and no column values specified, then all entries in the table will be removed.
- u** Update attributes of a table. This allows the concatenation path (**-p**), separation character (specified with the **(-s)**), column access rights, and table type string (**(-t)**) of a table to be changed. Neither the number of columns, nor the columns that are searchable may be changed.
- D *defaults***

When creating objects, this option specifies a different set of defaults to be used during this operation. The *defaults* string is a series of tokens separated by colons. These tokens represent the default values to be used for the generic object properties. All of the legal tokens are described below.

ttl=*time* This token sets the default time to live for objects that are created by this command. The value *time* is specified in the format as defined by the **nischttl(1)** command. The default value is 12 hours.

owner=*ownername*

This token specifies that the NIS+ principal *ownername* should own the created object. Normally this value is the same as the principal who is executing the command.

group=*groupname*

This token specifies that the group *groupname* should be the group owner for the object that is created. The default value is NULL.

access=*rights*

This token specifies the set of access rights that are to be granted for the given object. The value *rights* is specified in the format as defined by the **nischmod(1)** command. The default value is

----rmcdr---r---

- p *path*** When creating or updating a table, this option specifies the table's search path. When a **nisl_list()** function is invoked, the user can specify the flag FOLLOW_PATH to tell the client library to continue searching tables in the table's path if the search criteria used does not yield any entries. The path consists of an ordered list of table names, separated by colons. The names in the path must be fully qualified.
- s *sep*** When creating or updating a table, this option specifies the table's separator character. The separator character is used by **niscat(1)** when displaying tables on the standard output. Its purpose is to separate column data when the table is in ASCII form. The default value is a space.
- t *type*** When updating a table, this option specifies the table's type string.

RETURN VALUES

This command returns **0** on success and **1** on failure.

EXAMPLES

This example creates a table named **hobbies** in the directory **foo.com.** of the type **hobby_tbl** with two searchable columns, **name** and **hobby**.

example% nistbladm -c hobby_tbl name=S,a+r,o+m hobby=S,a+r hobbies.foo.com.

The column **name** has read access for all (that is, **owner**, **group**, and **world**) and modify access for only the owner. The column **hobby** is readable by all, but not modifiable by anyone.

In this example, if the access rights had not been specified, the tables access rights would have come from either the standard defaults or the **NIS_DEFAULTS** variable (see below).

To add entries to this table:

```
example% nistbladm -a name=bob hobby=skiing hobbies.foo.com.
example% nistbladm -a name=sue hobby=skiing hobbies.foo.com.
example% nistbladm -a name=ted hobby=swimming hobbies.foo.com.
```

To add the concatenation path:

```
example% nistbladm -u -p hobbies.bar.com.:hobbies.baz.com. hobbies
```

To delete the skiers from our list:

```
example% nistbladm -R hobby=skiing hobbies.foo.com.
```

Note: The use of the `-r` option would fail because there are two entries with the value of **skiing**.

To create a table with a column that is named with no flags set, you supply only the name and the equals (=) sign as follows.

```
example% nistbladm -c notes_tbl name=S,a+r,o+m note= notes.foo.com.
```

This example created a table, named *notes.foo.com.*, of type *notes_tbl* with two columns **name** and **note**. The **note** column is not searchable.

When entering data for columns in the form of a *value* string, it is essential that terminal characters be protected by single or double quotes. These are the characters equals (=), comma (,), left bracket ([), right bracket (]), and space (). These characters are parsed by NIS+ within an indexed name. These characters are protected by enclosing the entire value in double quote (") characters as follows.

```
example% nistbladm -a fullname="Joe User" nickname=Joe nicknames
```

If there is any doubt about how the string will be parsed, it is better to enclose it in quotes.

ENVIRONMENT

NIS_DEFAULTS This variable contains a defaults string that will be override the NIS+ standard defaults. If the `-D` switch is used those values will then override both the **NIS_DEFAULTS** variable and the standard defaults.

NIS_PATH If this variable is set, and the NIS+ table name is not fully qualified, each directory specified will be searched until the table is found (see **nisdefaults(1)**).

SEE ALSO

nis+(1), **niscat(1)**, **nischmod(1)**, **nischown(1)**, **nisdefaults(1)**, **nismatch(1)**, **nissetup(1M)**

WARNINGS

To modify one of the entries, say, for example, from "bob" to "robert":

```
example% nistbladm -m name=robert [name=bob],hobbies
```

Note that "[**name=bob**],hobbies" is an indexed name, and that the characters '[' (open bracket) and ']' (close bracket) are interpreted by the shell. When typing entry names in the form of NIS+ indexed names, the name must be protected by using single quotes.

It is possible to specify a set of defaults such that you cannot read or modify the table object later.

NAME	nistest – return the state of the NIS+ namespace using a conditional expression
SYNOPSIS	nistest [-ALMP] [-a rights -t type] <i>object</i> nistest [-ALMP] [-a rights] <i>indexedname</i>
AVAILABILITY	SUNWnisu
DESCRIPTION	nistest provides a way for shell scripts and other programs to test for the existence, type, and access rights of objects and entries. Entries are named using indexed names (see nismatch(1)).
OPTIONS	<p>-A All data. This option specifies that the data within the table and all of the data in tables in the initial table's concatenation path be returned. This option is only valid when using indexed names or following links.</p> <p>-L Follow links. If the object named by <i>object</i> or the tablename component of <i>indexedname</i> names a LINK type object, the link is followed when this switch is present.</p> <p>-M Master server only. This option specifies that the lookup should be sent to the master server of the named data. This guarantees that the most up to date information is seen at the possible expense that the master server may be busy.</p> <p>-P Follow concatenation path. This option specifies that the lookup should follow the concatenation path of a table if the initial search is unsuccessful. This option is only valid when using indexed names or following links.</p> <p>-a rights This option is used to verify that the current process has the desired or required access rights on the named object or entries. The access rights are specified in the same way as the nischmod(1) command.</p> <p>-t type This option tests the type of <i>object</i>. The value of <i>type</i> can be one of the following:</p> <p>G Return true if the object is a group object.</p> <p>D Return true if the object is a directory object.</p> <p>T Return true if the object is a table object.</p> <p>L Return true if the object is a link object.</p> <p>P Return true if the object is a private object.</p>
RETURN VALUES	<p>0 Success.</p> <p>1 Failure due to object not present, not of specified type and/or no such access.</p> <p>2 Failure due to illegal usage.</p>

EXAMPLES

When testing for access rights, **nistest** returns success (0) if the specified rights are granted to the current user. Thus testing for access rights

```
example% nistest -a w=mr skippy.domain
```

Tests that all authenticated NIS+ clients have read and modify access to the object named *skippy.domain*.

Testing for access on a particular entry in a table can be accomplished using the indexed name syntax. The following example tests to see if an entry in the password table can be modified.

```
example% nistest -a o=m '[uid=99],passwd.org_dir'
```

ENVIRONMENT

NIS_PATH If this variable is set, and the NIS+ name is not fully qualified, each directory specified will be searched until the object is found (see **nisdefaults(1)**).

SEE ALSO

nis+(1), **nischmod(1)**, **nisdefaults(1)**

NAME	nl – line numbering filter								
SYNOPSIS	nl [- <i>btype</i>] [- <i>ftype</i>] [- <i>htype</i>] [- <i>vstart#</i>] [- <i>iincr</i>] [- <i>p</i>] [- <i>lnum</i>] [- <i>ssep</i>] [- <i>wwidth</i>] [- <i>nformat</i>] [- <i>ddelim</i>] [<i>filename</i>]								
AVAILABILITY	SUNWesu								
DESCRIPTION	<p>nl reads lines from the named <i>filename</i>, or the standard input if no <i>filename</i> is named, and reproduces the lines on the standard output. Lines are numbered on the left in accordance with the command options in effect.</p> <p>nl views the text it reads in terms of logical pages. Line numbering is reset at the start of each logical page. A logical page consists of a header, a body, and a footer section. Empty sections are valid. Different line numbering options are independently available for header, body, and footer. For example, -bt (the default) numbers non-blank lines in the body section and does not number any lines in the header and footer sections.</p> <p>The start of logical page sections are signaled by input lines containing nothing but the following delimiter character(s):</p> <table border="0" style="margin-left: 40px;"> <tr> <td style="padding-right: 20px;"><i>Line contents</i></td> <td><i>Start of</i></td> </tr> <tr> <td style="padding-right: 20px;">\:\:\:</td> <td>header</td> </tr> <tr> <td style="padding-right: 20px;">\:\:</td> <td>body</td> </tr> <tr> <td style="padding-right: 20px;">\:</td> <td>footer</td> </tr> </table> <p>Unless optioned otherwise, nl assumes the text being read is in a single logical page body.</p>	<i>Line contents</i>	<i>Start of</i>	\:\:\:	header	\:\:	body	\:	footer
<i>Line contents</i>	<i>Start of</i>								
\:\:\:	header								
\:\:	body								
\:	footer								
OPTIONS	<p>Command options may appear in any order and may be intermingled with an optional file name. Only one file may be named. The options are:</p> <p>-btype Specifies which logical page body lines are to be numbered. Recognized <i>types</i> and their meanings are:</p> <table border="0" style="margin-left: 40px;"> <tr> <td style="padding-right: 20px;">a</td> <td>number all lines</td> </tr> <tr> <td style="padding-right: 20px;">t</td> <td>number lines with printable text only</td> </tr> <tr> <td style="padding-right: 20px;">n</td> <td>no line numbering</td> </tr> </table> <p>pexp number only lines that contain the regular expression specified in <i>exp</i> (see ed(1))</p> <p>Default <i>type</i> for logical page body is t (text lines numbered).</p> <p>-ftype Same as -btype except for footer. Default <i>type</i> for logical page footer is n (no lines numbered).</p> <p>-htype Same as -btype except for header. Default <i>type</i> for logical page header is n (no lines numbered).</p>	a	number all lines	t	number lines with printable text only	n	no line numbering		
a	number all lines								
t	number lines with printable text only								
n	no line numbering								

- vstart#** *start#* is the initial value used to number logical page lines. Default *start#* is **1**.
- iincr** *incr* is the increment value used to number logical page lines. Default *incr* is **1**.
- p** Do not restart numbering at logical page delimiters.
- lnum** *num* is the number of blank lines to be considered as one. For example, **-l2** results in only the second adjacent blank being numbered (if the appropriate **-ha**, **-ba**, and/or **-fa** option is set). Default *num* is **1**.
- sssep** *sep* is the character(s) used in separating the line number and the corresponding text line. Default *sep* is a tab.
- wwidth** *width* is the number of characters to be used for the line number. Default *width* is **6**.
- nformat** *format* is the line numbering format. Recognized values are: **ln**, left justified, leading zeroes suppressed; **rn**, right justified, leading zeroes suppressed; **rz**, right justified, leading zeroes kept. Default *format* is **rn** (right justified).
- ddelim** The two delimiter characters specifying the start of a logical page section may be changed from the default characters (\ :) to two user-specified characters. If only one character is entered, the second character remains the default character (:). No space should appear between the **-d** and the delimiter characters. To enter a backslash, use two backslashes.

EXAMPLES

The command:

```
example% nl -v10 -i10 -d!+ filename1
```

will cause the first line of the page body to be numbered **10**, the second line of the page body to be numbered **20**, the third **30**, and so forth. The logical page delimiters are **!+**.

SEE ALSO

ed(1), **pr(1)**

NAME	nm – print name list of an object file
SYNOPSIS	nm [-CefhlnoprRsTuvVx] <i>filename</i> . . .
DESCRIPTION	<p>The nm command displays the symbol table of each ELF object file that is specified by <i>filename(s)</i>. For each symbol, the following information will be printed:</p> <p>Index The index of the symbol. (The index appears in brackets.)</p> <p>Value The value of the symbol is one of the following: a section offset for defined symbols in a relocatable file; alignment constraints for symbols whose section index is SHN_COMMON; a virtual address in executable and dynamic library files.</p> <p>Size The size in bytes of the associated object.</p> <p>Type A symbol is of one of the following types: NOTYPE (no type was specified), OBJECT (a data object such as an array or variable), FUNC (a function or other executable code), SECTION (a section symbol), or FILE (name of the source file).</p> <p>Bind The symbol's binding attributes. LOCAL symbols have a scope limited to the object file containing their definition; GLOBAL symbols are visible to all object files being combined; and WEAK symbols are essentially global symbols with a lower precedence than GLOBAL.</p> <p>Other A field reserved for future use, currently containing 0.</p> <p>Shndx Except for three special values, this is the section header table index in relation to which the symbol is defined. The following special values exist: ABS indicates the symbol's value will not change through relocation; COMMON indicates an unallocated block and the value provides alignment constraints; and UNDEF indicates an undefined symbol.</p> <p>Name The name of the symbol.</p>
OPTIONS	<p>The output of nm may be controlled using the following options:</p> <p>-C Demangle C++ symbol names before printing them out.</p> <p>-e See NOTES below.</p> <p>-f See NOTES below.</p> <p>-h Do not display the output heading data.</p> <p>-l Distinguish between WEAK and GLOBAL symbols by appending a * to the key letter for WEAK symbols.</p> <p>-n Sort external symbols by name before they are printed.</p> <p>-o Print the value and size of a symbol in octal instead of decimal.</p> <p>-p Produce easily parsable, terse output. Each symbol name is preceded by its value (blanks if undefined) and one of the letters U (undefined), N (symbol</p>

has no type), **D** (data object symbol), **B** (bss symbol), **T** (text symbol), **S** (section symbol), or **F** (file symbol). If the symbol's binding attribute is **LOCAL**, the key letter is lower case; if the symbol's binding attribute is **WEAK**, the key letter is upper case; if the **-l** modifier is specified, the upper case key letter is followed by a *; if the symbol's binding attribute is **GLOBAL**, the key letter is upper case.

- r** Prepend the name of the object file or archive to each output line.
- R** Print the archive name (if present), followed by the object file and symbol name. If the **-r** option is also specified, this option is ignored.
- s** Print section name instead of section index.
- T** See NOTES below.
- u** Print undefined symbols only.
- v** Sort external symbols by value before they are printed.
- V** Print the version of the **nm** command executing on the standard error output.
- x** Print the value and size of a symbol in hexadecimal instead of decimal.

Options may be used in any order, either singly or in combination, and may appear anywhere in the command line. When conflicting options are specified (such as **-v** and **-n**; and **-o** and **-x**) the first is taken and the second ignored with a warning message to the user. (See **-R** for exception.)

SEE ALSO **as(1)**, **dump(1)**, **ld(1)**, **a.out(4)**, **ar(4)**

NOTES

The following options are obsolete because of changes to the object file format and will be deleted in a future release.

- e** Print only external and static symbols. The symbol table now contains only static and external symbols. Automatic symbols no longer appear in the symbol table. They do appear in the debugging information produced by **cc -g**, which may be examined using **dump(1)**.
- f** Produce full output. Redundant symbols (such as .text, .data, and so forth), which existed previously do not exist and producing full output will be identical to the default output.
- T** By default, **nm** prints the entire name of the symbols listed. Since symbol names have been moved to the last column, the problem of overflow is removed and it is no longer necessary to truncate the symbol name.

NAME	nohup – run a command immune to hangups and quits
SYNOPSIS	<code>/usr/bin/nohup <i>command</i> [<i>arguments</i>]</code>
AVAILABILITY	SUNWcsu
DESCRIPTION	Processes run by nohup are immune to HUP (hangup) and QUIT (quit) signals; nohup does not arrange to make them immune to a TERM (terminate) signal, so unless they arrange to be immune to a TERM signal, or the shell makes them immune to a TERM signal, they will receive that signal. If nohup.out is not writable in the current directory, output is redirected to \$HOME/nohup.out . If the standard error is a terminal, it is redirected to the standard output, otherwise it is not redirected. The priority of the process run by nohup is not altered.
EXAMPLES	<p>It is frequently desirable to apply nohup to pipelines or lists of commands. This can be done only by placing pipelines and command lists in a single file, called a shell procedure. One can then issue:</p> <pre style="margin-left: 40px;">\$ nohup sh <i>filename</i></pre> <p>and the nohup applies to everything in <i>filename</i>. If the shell procedure <i>filename</i> is to be executed often, then the need to type sh can be eliminated by giving <i>filename</i> execute permission. Add an ampersand and the contents of <i>filename</i> are run in the background with interrupts also ignored (see sh(1)):</p> <pre style="margin-left: 40px;">\$ nohup <i>filename</i> &</pre> <p>An example of what the contents of <i>filename</i> could be is:</p> <pre style="margin-left: 40px;">sort <i>ofilename</i> > <i>nfilename</i></pre>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of nohup for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how nohup behaves.</p> <p>LC_CTYPE</p> <p>Determines how nohup handles characters. When LC_CTYPE is set to a valid value, nohup can display and handle text and filenames containing valid characters for that locale. nohup can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. nohup can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p>

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

chmod(1), csh(1), nice(1), sh(1), shell_builtins(1), signal(3C), environ(5)

NOTES

If you use **csh(1)**, then commands executed with '&' are automatically immune to HUP signals while in the background.

There is a C shell built-in command **nohup** that provides immunity from terminate, but does not redirect output to **nohup.out**.

nohup does not recognize command sequences. In the case of the following command

```
$ nohup command1; command2
```

nohup applies only to **command1**. The command

```
$ nohup (command1; command2)
```

is syntactically incorrect.

NAME	nroff – format documents for display or line-printer																				
SYNOPSIS	nroff [-ehiq] [-mname] [-nN] [-opagelist] [-raN] [-sN] [-Tname]																				
AVAILABILITY	SUNWdoc																				
DESCRIPTION	nroff formats text in the named <i>files</i> for typewriter-like devices. See also troff (1). If no <i>file</i> argument is present, nroff reads the standard input. An argument consisting of a '-' is taken to be a file name corresponding to the standard input.																				
OPTIONS	Options may appear in any order so long as they appear <i>before</i> the files. -e Produce equally-spaced words in adjusted lines, using full terminal resolution. -h Use output TAB characters during horizontal spacing to speed output and reduce output character count. TAB settings are assumed to be every 8 nominal character widths. -i Read the standard input after the input files are exhausted. -q Invoke the simultaneous input-output mode of the rd request. -mname Prepend the macro file /usr/share/lib/tmac/tmac.name to the input files. -nN Number first generated page <i>N</i> . -opagelist Print only pages whose page numbers appear in the comma-separated <i>list</i> of numbers and ranges. A range <i>N–M</i> means pages <i>N</i> through <i>M</i> ; an initial <i>-N</i> means from the beginning to page <i>N</i> ; and a final <i>N–</i> means from <i>N</i> to the end. -raN Set register <i>a</i> (one-character) to <i>N</i> . -sN Stop every <i>N</i> pages. nroff will halt prior to every <i>N</i> pages (default <i>N</i> =1) to allow paper loading or changing, and will resume upon receipt of a NEWLINE. -Tname Prepare output for a device of the specified <i>name</i> . Known <i>names</i> are: <table border="0" style="margin-left: 2em;"> <tr> <td>37</td> <td>Teletype Corporation Model 37 terminal — this is the default.</td> </tr> <tr> <td>lp tn300</td> <td>GE Any line printer or terminal without half-line capability.</td> </tr> <tr> <td>300</td> <td>DASI-300.</td> </tr> <tr> <td>300-12</td> <td>DASI-300 — 12-pitch.</td> </tr> <tr> <td>300S</td> <td>DASI-300S.</td> </tr> <tr> <td>300S-12</td> <td>DASI-300S.</td> </tr> <tr> <td>382</td> <td>DASI-382 (fancy DTC 382).</td> </tr> <tr> <td>450</td> <td>DASI-450 (Diablo Hyterm).</td> </tr> <tr> <td>450-12</td> <td>DASI-450 (Diablo Hyterm) — 12-pitch.</td> </tr> <tr> <td>832</td> <td>AJ 832.</td> </tr> </table>	37	Teletype Corporation Model 37 terminal — this is the default.	lp tn300	GE Any line printer or terminal without half-line capability.	300	DASI-300.	300-12	DASI-300 — 12-pitch.	300S	DASI-300S.	300S-12	DASI-300S.	382	DASI-382 (fancy DTC 382).	450	DASI-450 (Diablo Hyterm).	450-12	DASI-450 (Diablo Hyterm) — 12-pitch.	832	AJ 832.
37	Teletype Corporation Model 37 terminal — this is the default.																				
lp tn300	GE Any line printer or terminal without half-line capability.																				
300	DASI-300.																				
300-12	DASI-300 — 12-pitch.																				
300S	DASI-300S.																				
300S-12	DASI-300S.																				
382	DASI-382 (fancy DTC 382).																				
450	DASI-450 (Diablo Hyterm).																				
450-12	DASI-450 (Diablo Hyterm) — 12-pitch.																				
832	AJ 832.																				

- EXAMPLE** The following command:
example% nroff -s4 -me users.guide
formats **users.guide** using the **-me** macro package, and stopping every 4 pages.
- FILES**
- | | |
|------------------------------------|--|
| /var/tmp/trtmp* | temporary file |
| /usr/share/lib/tmac/tmac.* | standard macro files |
| /usr/share/lib/nterm/* | terminal driving tables for nroff |
| /usr/share/lib/nterm/README | index to terminal description files |
- SEE ALSO** **checknr(1)**, **col(1)**, **eqn(1)**, **man(1)**, **tbl(1)**, **troff(1)**, **term(5)**, **me(5)**, **ms(5)**
- NOTES** **nroff** is not 8-bit clean because making **nroff** 8-bit clean would require rewriting the **nroff** internals and filters. Also, some **nroff** syntax is based on ASCII only and does not lend itself to 8-bit character sequences.

NAME	od – octal dump
SYNOPSIS	od [-bcCDdFfOoSsvXx] [<i>filename</i>] [[+] <i>offset</i> [.] [b]]
DESCRIPTION	od displays <i>filename</i> in one or more formats, as selected by the first argument. If the first argument is missing, -o is default. If no <i>filename</i> is specified, the standard input is used. For the purposes of this description, "word" refers to a 16-bit unit, independent of the word size of the machine; "long word" refers to a 32-bit unit, and "double long word" refers to a 64-bit unit.
OPTIONS	<p>-b Interpret bytes in octal.</p> <p>-c Display single-byte characters. Certain non-graphic characters appear as C-language escapes: null=\0, backspace=\b, form-feed=\f, new-line=\n, return=\r, tab=\t; others appear as 3-digit octal numbers. For example:</p> <pre>echo "hello world" od -c 0000000 h e l l o w o r l d \n 0000014</pre> <p>-C Interpret bytes as single-byte or multibyte characters according to the current setting of the LC_CTYPE locale category. Printable multibyte characters are written in the area corresponding to the first byte of the character; two character sequence ** are written in the area corresponding to each remaining byte in the character, as an indication that the character is continued. Non graphic characters appear as same as they would using the -c option.</p> <p>-D Interpret long words in unsigned decimal.</p> <p>-d Interpret words in unsigned decimal.</p> <p>-F Interpret double long words in extended precision.</p> <p>-f Interpret long words in floating point.</p> <p>-O Interpret long words in unsigned octal.</p> <p>-o Interpret words in octal.</p> <p>-S Interpret long words in signed decimal.</p> <p>-s Interpret words in signed decimal.</p> <p>-v Show all data (verbose).</p> <p>-X Interpret long words in hex.</p> <p>-x Interpret words in hex.</p> <p><i>offset</i> specifies an offset from the beginning of <i>filename</i> where the display will begin. <i>offset</i> is normally interpreted as octal bytes. If . is appended, <i>offset</i> is interpreted in decimal. If b is appended, <i>offset</i> is interpreted in blocks of 512 bytes. If <i>filename</i> is omitted, <i>offset</i> must be preceded by +.</p>

The display continues until an end-of-file is reached.

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

environ(5)

NAME	on – execute a command on a remote system, but with the local environment
SYNOPSIS	on [-i] [-d] [-n] <i>host command</i> [<i>argument</i>] ...
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The on program is used to execute commands on another system, in an environment similar to that invoking the program. All environment variables are passed, and the current working directory is preserved. To preserve the working directory, the working file system must be either already mounted on the host or be exported to it. Relative path names will only work if they are within the current file system; absolute path names may cause problems.</p> <p>The standard input is connected to the standard input of the remote command, and the standard output and the standard error from the remote command are sent to the corresponding files for the on command.</p>
OPTIONS	<p>-i Interactive mode. Use remote echoing and special character processing. This option is needed for programs that expect to be talking to a terminal. All terminal modes and window size changes are propagated.</p> <p>-d Debug mode. Print out some messages as work is being done.</p> <p>-n No Input. This option causes the remote program to get EOF when it reads from the standard input, instead of passing the standard input from the standard input of the on program. For example, -n is necessary when running commands in the background with job control.</p>
SEE ALSO	chkey(1)
DIAGNOSTICS	<p>unknown host Host name not found.</p> <p>cannot connect to server Host down or not running the server.</p> <p>can't find Problem finding the working directory.</p> <p>can't locate mount point Problem finding current file system.</p> <p>RPC: Authentication error The server requires des authentication and you do not have a secret key registered with keyserv. Perhaps you logged in without a password. Try to keylogin. If that fails try to set your publickey with chkey.</p> <p>Other error messages may be passed back from the server.</p>

BUGS When the working directory is remote mounted over NFS, a **^Z** hangs the window.
Root cannot use **on**.

NAME	pack, pcat, unpack – compress and expand files
SYNOPSIS	pack [-] [-f] <i>name</i> ... pcat <i>name</i> ... unpack <i>name</i> ...
AVAILABILITY	SUNWesu
DESCRIPTION	<p>pack attempts to store the specified files in a compressed form. Wherever possible (and useful), each input file <i>name</i> is replaced by a packed file <i>name.z</i> with the same access modes, access and modified dates, and owner as those of <i>name</i>.</p> <p>The amount of compression obtained depends on the size of the input file and the character frequency distribution. Because a decoding tree forms the first part of each <i>.z</i> file, it is usually not worthwhile to pack files smaller than three blocks, unless the character frequency distribution is very skewed, which may occur with printer plots or pictures.</p> <p>Typically, text files are reduced to 60-75% of their original size. Load modules, which use a larger character set and have a more uniform distribution of characters, show little compression, the packed versions being about 90% of the original size.</p> <p>pack returns a value that is the number of files that it failed to compress.</p> <p>No packing will occur if:</p> <ul style="list-style-type: none"> the file appears to be already packed; the file name has more than 12 characters; the file has links; the file is a directory; the file cannot be opened; no disk storage blocks will be saved by packing; a file called <i>name.z</i> already exists; the <i>.z</i> file cannot be created; an I/O error occurred during processing. <p>The last segment of the file name must contain no more than 12 characters to allow space for the appended <i>.z</i> extension. Directories cannot be compressed.</p> <p>pcat does for packed files what cat(1) does for ordinary files, except that pcat cannot be used as a filter. The specified files are unpacked and written to the standard output. Thus to view a packed file named name.z use:</p> <p style="padding-left: 2em;">pcat name.z</p> <p>or just:</p> <p style="padding-left: 2em;">pcat name</p>

To make an unpacked copy, say **nnn**, of a packed file named **name.z** (without destroying **name.z**) use the command:

```
pcat name >nnn
```

pcat returns the number of files it was unable to unpack. Failure may occur if:

- the file name (exclusive of the **.z**) has more than 12 characters;
- the file cannot be opened;
- the file does not appear to be the output of **pack**.

unpack expands files created by **pack**. For each file *name* specified in the command, a search is made for a file called *name.z* (or just *name*, if *name* ends in **.z**). If this file appears to be a packed file, it is replaced by its expanded version. The new file has the **.z** suffix stripped from its name, and has the same access modes, access and modification dates, and owner as those of the packed file.

Unpack returns a value that is the number of files it was unable to unpack. Failure may occur for the same reasons that it may in **pcat**, as well as for the following:

- a file with the "unpacked" name already exists;
- if the unpacked file cannot be created.

OPTIONS

- f** Forces packing of *name*. This is useful for causing an entire directory to be packed even if some of the files will not benefit. If **pack** is successful, *name* will be removed. Packed files can be restored to their original form using **unpack** or **pcat**.
- **pack** uses Huffman (minimum redundancy) codes on a byte-by-byte basis. If the **-** argument is used, an internal flag is set that causes the number of times each byte is used, its relative frequency, and the code for the byte to be printed on the standard output. Additional occurrences of **-** in place of *name* will cause the internal flag to be set and reset.

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

cat(1), compress(1), environ(5)

NAME	pagesize – display the size of a page of memory
SYNOPSIS	/usr/ucb/pagsize
AVAILABILITY	SUNWscpu
DESCRIPTION	pagesize prints the size of a page of memory in bytes, as returned by getpagesize . This program is useful in constructing portable shell scripts.
SEE ALSO	getpagesize(3B)

NAME	passwd – change login password and password attributes
SYNOPSIS	<pre>passwd [name] passwd [-d -l] [-f] [-n min] [-w warn] [-x max] name passwd -s [-a] passwd -s [name]</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The passwd command changes the password or lists password attributes associated with the user's login <i>name</i>. Additionally, privileged-users may use passwd to install or change passwords and attributes associated with any login <i>name</i>.</p> <p>When used to change a password, passwd prompts ordinary users for their old password, if any. It then prompts for the new password twice. When the old password is entered, passwd checks to see if it has "aged" sufficiently. If "aging" is insufficient, passwd terminates; see pwconv(1M) and shadow(4) for additional information. The pwconv command creates and updates /etc/shadow with information from /etc/passwd. pwconv relies on a special value of 'x' in the password field of /etc/passwd. This value of 'x' indicates that the password for the user is already in /etc/shadow and should not be modified.</p> <p>Assuming aging is sufficient, a check is made to ensure that the new password meets construction requirements. When the new password is entered a second time, the two copies of the new password are compared. If the two copies are not identical the cycle of prompting for the new password is repeated for at most two more times.</p> <p>Passwords must be constructed to meet the following requirements:</p> <ul style="list-style-type: none"> • Each password must have at least six characters. Only the first eight characters are significant. PASSLENGTH is found in /etc/default/passwd and is set to 6. • Each password must contain at least two alphabetic characters and at least one numeric or special character. In this case, "alphabetic" refers to all upper or lower case letters. • Each password must differ from the user's login <i>name</i> and any reverse or circular shift of that login <i>name</i>. For comparison purposes, an upper case letter and its corresponding lower case letter are equivalent. • New passwords must differ from the old by at least three characters. For comparison purposes, an upper case letter and its corresponding lower case letter are equivalent.

Super-users (for instance, real and effective uid equal to zero, see **id(1M)** and **su(1M)**) may change any password; hence, **passwd** does not prompt privileged-users for the old password. Privileged-users are not forced to comply with password aging and password construction requirements. A privileged-user can create a null password by entering a carriage return in response to the prompt for a new password. (This differs from **passwd -d** because the "password" prompt will still be displayed.)

Any user may use the **-s** option to show password attributes for his or her own login *name*.

The format of the display will be:

name status mm/dd/yy min max warn

or, if password aging information is not present,

name status

where

name The login ID of the user.

status The password status of *name*: "PS" stands for passworded or locked, "LK" stands for locked, and "NP" stands for no password.

mm/dd/yy The date password was last changed for *name*. (Note: All password aging dates are determined using Greenwich Mean Time and, therefore, may differ by as much as a day in other time zones.)

min The minimum number of days required between password changes for *name*. MINWEEKS is found in **/etc/default/passwd** and is set to NULL.

max The maximum number of days the password is valid for *name*. MAXWEEKS is found in **/etc/default/passwd** and is set to NULL.

warn The number of days relative to *max* before the password expires that the *name* will be warned.

OPTIONS

Only a privileged-user can use the following options:

- a** Show password attributes for all entries. Use only with **-s** option; *name* must not be provided.
- d** Deletes password for *name*. The login *name* will not be prompted for password.
- f** Force the user to change password at the next login by expiring the password for *name*.
- l** Locks password entry for *name*.
- s** Show password attributes for the login *name*.

- n *min*** Set minimum field for *name*. The *min* field contains the minimum number of days between password changes for *name*. If *min* is greater than *max*, the user may not change the password. Always use this option with the **-x** option, unless *max* is set to -1 (aging turned off). In that case, *min* need not be set.
- w *warn*** Set warn field for *name*. The *warn* field contains the number of days before the password expires that the user will be warned.
- x *max*** Set maximum field for *name*. The *max* field contains the number of days that the password is valid for *name*. The aging for *name* will be turned off immediately if *max* is set to -1. If it is set to 0, then the user is forced to change the password at the next login session and aging is turned off.

ENVIRONMENT

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

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

LC_MESSAGES Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

FILES

/etc/oshadow

/etc/passwd

/etc/shadow

/etc/default/passwd Default values can be set for the following flags in **/etc/default/passwd**. For example: MAXWEEKS=26

MAXWEEKS Maximum time period that password is valid.

MINWEEKS Minimum time period before the password can be changed.

PASSLENGTH Minimum length of password, in characters.

WARNWEEKS Time period until warning of date of password's ensuing expiration.

SEE ALSO

finger(1), login(1), nispasswd(1), yppasswd(1), domainname(1M), eeprom(1M), id(1M), passmgmt(1M), pwconv(1M), su(1M), useradd(1M), userdel(1M), usermod(1M), crypt(3C), getpwnam(3C), getspnam(3C), loginlog(4), passwd(4), shadow(4), environ(5)

DIAGNOSTICS

The **passwd** command exits with one of the following values:

- 0 SUCCESS.
- 1 Permission denied.
- 2 Invalid combination of options.
- 3 Unexpected failure. Password file unchanged.
- 4 Unexpected failure. Password file(s) missing.
- 5 Password file(s) busy. Try again later.
- 6 Invalid argument to option.

NAME	paste – merge same lines of several files or subsequent lines of one file
SYNOPSIS	<pre>paste filename1 filename2... paste -d list filename1 filename2... paste -s [-d list] filename1 filename2...</pre>
AVAILABILITY	SUNWesu
DESCRIPTION	<p>In the first two forms, paste concatenates corresponding lines of the given input files <i>filename1</i>, <i>filename2</i>, and so forth. It treats each file as a column or columns of a table and pastes them together horizontally (parallel merging). If you will, it is the counterpart of cat(1) which concatenates vertically, that is, one file after the other. In the last form above, paste replaces the function of an older command with the same name by combining subsequent lines of the input file (serial merging). In all cases, lines are glued together with the <i>tab</i> character, or with characters from an optionally specified <i>list</i>. Output is to the standard output, so it can be used as the start of a pipe, or as a filter, if - is used in place of a file name.</p>
OPTIONS	<p>-d list Without this option, the new-line characters of each but the last file (or last line in case of the -s option) are replaced by a <i>tab</i> character. This option allows replacing the <i>tab</i> character by one or more alternate characters (see below).</p> <p>One or more characters immediately following -d replace the default <i>tab</i> as the line concatenation character. The <i>list</i> is used circularly, that is, when exhausted, it is reused. In parallel merging (that is, no -s option), the lines from the last file are always terminated with a new-line character, not from the <i>list</i>. The <i>list</i> may contain the special escape sequences: \n (new-line), \t (tab), \\ (backslash), and \0 (empty string, not a null character). Quoting may be necessary, if characters have special meaning to the shell (for instance, to get one backslash, use <code>" -d "\\ \\"</code>).</p> <p>-s Merge subsequent lines rather than one from each input file. Use <i>tab</i> for concatenation, unless a <i>list</i> is specified with -d option. Regardless of the <i>list</i>, the very last character of the file is forced to be a new-line.</p> <p>- May be used in place of any file name, to read a line from the standard input. (There is no prompting).</p>
EXAMPLES	<p>The first example will list a directory in one column.</p> <pre>ls paste -d" " -</pre> <p>The second example will list a directory in four columns.</p> <pre>ls paste - - - -</pre>

The next example will combine pairs of lines into lines.

```
paste -s -d"\t\n" file
```

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

`cut(1)`, `grep(1)`, `pr(1)`, `environ(5)`

DIAGNOSTICS

"line too long" Output lines are restricted to 511 characters.

"too many files" Except for `-s` option, no more than 12 input files may be specified.

NAME	pathconv – search FMLI criteria for filename
SYNOPSIS	pathconv [-f] [-v alias] pathconv [-t] [-l] [-nnum] [-v string]
DESCRIPTION	The pathconv function converts an alias to its pathname. By default, it takes the alias as a string from the standard input.
OPTIONS	<p>-f If -f is specified, the full path will be returned (this is the default).</p> <p>-t If -t is specified, pathconv will truncate a pathname specified in <i>string</i> in a format suitable for display as a frame title. This format is a shortened version of the full pathname, created by deleting components of the path from the middle of the string until it is under DISPLAYW - 6 characters in length, and then inserting ellipses (...) between the remaining pieces. Ellipses are also used to show truncation at the ends of the strings if necessary, unless the -l option is given.</p> <p>-l If -l is specified, < and > will be used instead of ellipses (...) to indicate truncation at the ends of the string generated by the -t option. Using -l allows display of the longest possible string while still notifying users it has been truncated.</p> <p>-nnum If -n is specified, <i>num</i> is the maximum length of the string (in characters) generated by the -t option. The argument <i>num</i> can be any integer from 1 to 255.</p> <p>-v alias string If the -v option is used, then <i>alias</i> or <i>string</i> can be specified when pathconv is called. The argument <i>alias</i> must be an alias defined in the <i>alias_file</i> named when fmli was invoked. The argument <i>string</i> can only be used with the -t option and must be a pathname.</p>
EXAMPLES	<p>Here is a menu descriptor that uses pathconv to construct the menu title. It searches for MYPATH in the <i>alias_file</i> named when fmli was invoked:</p> <pre>menu=`pathconv -v MYPATH/ls` . . .</pre> <p>where there is a line in <i>alias_file</i> that defines MYPATH. For example, MYPATH=\$HOME/bin:/usr/bin.</p> <p>Here is a menu descriptor that takes <i>alias</i> from the standard input.</p> <pre>menu=`echo MYPATH/ls pathconv` . . .</pre>
SEE ALSO	fml i(1)

NAME	pcmapkeys – set keyboard extended map and scancode translation for the PC console in text mode
SYNOPSIS	pcmapkeys [<i>-f mapfile</i> <i>-n</i> <i>-g</i> <i>-d</i> <i>-e</i>]
AVAILABILITY	x86 SUNWcsu
DESCRIPTION	pcmapkeys is a utility that permits a user to activate character mapping on input and output and keyboard extended mapping on the PC console in text mode. The keyboard extended mapping consists of the support for the deadkey and compose key sequences.
Options	<p><i>-f mapfile</i> Installs the contents of the file <i>mapfile</i> and sets the corresponding mapping as supported by the console driver. The layout of the <i>mapfile</i> and the supported functionality are described below.</p> <p><i>-n</i> Disables and dismantles the current keyboard extended mapping. The <i>-f</i> option must be used to re-install the keyboard extended mapping.</p> <p><i>-g</i> Displays the current mappings and keyboard extended mapping (if one is installed) in hex values (see <i>/usr/include/sys/emap.h</i>). This option is mainly used for debugging purposes.</p> <p><i>-d</i> and <i>-e</i> <i>-d</i> temporarily disables the compose key and deadkey sequences if the keyboard extended mapping is installed. The keyboard extended mapping can be enabled again by using the <i>-e</i> option (or it can be re-installed by using the <i>-f</i> option).</p>
Consistent Keyboard-Display Mapping	<p>The original UNIX operating system was written to support the ASCII codeset. ASCII is one of many standards to represent a number of characters internally as certain numbers. Typical for ASCII is that it supports 128 different characters, each represented by a single byte of which the 8th bit is not used. Many UNIX system applications, including the shell, took advantage of this. Starting with UNIX System V Release 3.1, most of these applications have been modified to properly support characters represented as a byte with the 8th bit set as well. This means that now 256 characters can be supported at the same time. However, a consistent coding convention needs to be applied. In the IBM PC world, an 8-bit coding referred to as IBM extended ASCII has been used for several years; MS-DOS users are quite familiar with that. In heterogeneous UNIX system environments, a different codeset, called ISO 8859, has been promoted. In both codesets, characters found in the ASCII codeset are represented in the same way. The other 128 characters are encoded differently, however, and some characters found in one codeset will be missing in the other. The Solaris for x86 system supports both codesets; actually, it supports any 8-bit one byte codeset.</p> <p>To be able to use characters from the French, German, Finnish, and other alphabets, there are systems available on the market that generate 7-bit codes but display the above-mentioned characters on the screen instead of the ones found on a U.S. console. On the</p>

keyboard there are an equal number of keys, but there are different characters on the key caps. Others may support 256 different characters at a time but use their own proprietary codesets.

For example, if you are using the Solaris for x86 system with a console and a French keyboard and you do not use **pcmapkeys** to map the French keyboard tables, then if you edit a file and use the French character *é* in text, the actual code generated is ASCII 123, which is the code normally used for the left curly brace. If you look at the edited file on the console, the letter will actually appear to be a curly brace. Using **pcmapkeys** to map in the French keyboard allows consistent input and output mappings.

Input mapping

On input, any byte can be mapped to any byte. Using the example above, you could map 123 to 130, the code used for *é* in the IBM extended ASCII codeset.

Output mapping

On output, any byte can be mapped to either a byte or a string. In the above example, 130 would be mapped back to 123 to properly display the character on the screen. If the connected device is a printer that does not support the *é* character, it could be mapped to the string:

```
e BACKSPACE '
```

Deadkeys

On typewriters, keys can be found that behave slightly differently than all the others, because when you press them, the printing wheel of the typewriter does not move. Ctrl (`^`) and the grave accent (```) are such characters. When ``` is followed by an *e*, the letter *è* is generated. This is called a deadkey or a non-spacing character. Solaris for x86 supports the use of deadkeys. Typically, the `^` character, the ``` character, and the umlaut character are used as deadkeys.

Compose sequences

Characters can also be generated using a compose sequence. A dedicated character called the “compose character” followed by two other keystrokes will generate a single character. As an example, COMPOSE followed by the plus and the minus sign could generate the plus/minus sign (\pm). Compose sequences can also be used as an alternative for deadkeys, e.g., “COMPOSE `^` *e*” instead of “`^` *e*.”

Numeric compose sequences

Compose sequence characters that are not present on the keyboard and cannot be intuitively composed by some key sequence, for example, graphics characters, can be generated by pressing the compose key followed by three digits.

Toggle key

An optional toggle key can be defined to temporarily disable the current mapping from within an application. This can be useful when, for example, a German programmer wants easy access to the curly braces and the brackets. Use of the toggle key is analogous to the use of the `-d` and `-e` command line options.

Scancode Mapping

The keyboards of the console and some other peripherals such as SunRiver workstations behave differently than those of regular terminals. They generate what are called *scancodes* and you will also find a number of keys on these keyboards, such as the Alt key, that are not found on regular terminals. Scancodes generated by PC keyboards typically represent the location of the key on the keyboard. The keyboard driver has to properly translate these scancodes. The different national variants of a PC keyboard not only have non-English characters printed on some of the keycaps, but the order of some of the keys is different as well. Without changing the scancode translation, a French user would type **A** and see **Q** on his screen. Several status keys can influence the translated code as well. The keyboard driver, and thus the **pcmapkeys** program, makes a distinction between two sets of key combinations that can be translated.

Function keys

Up to 60 key combinations are recognized as function keys. The first 12 are the 12 function keys of a 101-key PC-keyboard (the first 10 on an 84-key keyboard).

If you do not know whether you have an 84- or 101-key keyboard, you can use the following scheme to determine which type you have:

If your keyboard has arrow keys that are separate from the ones on the numeric keypad, then you have a 101-key keyboard.

If the arrow keys on your keyboard are located on the numeric keypad only, then you have an 84-key keyboard.

F13 to F24 are the same keys used in combination with Shift, F25 to F36 when used with Ctrl, and F37 to F48 when used with Ctrl and Shift together. F49 to F60 are the keys on the numeric keypad, in the following order:

7
8
9
-
4
5
6
+
1
2
3
INS

Each of these function keys can be given a string as a value. The total length of all strings should not exceed 512 characters.

Regular keys

Scancodes generated by all keys on the PC keyboard can be translated in a different way as well. For each key, a different translation can be specified for each of the following four cases:

1. The key is pressed.
2. The key and the Shift key are pressed simultaneously.
3. The key and the Alt key are pressed simultaneously.
4. The key, the Shift, and the Alt keys are pressed simultaneously.

For each of these cases, the scancode can be translated into one of the following:

- a single byte
- a single byte preceded by ESC N
- a single byte preceded by ESC O
- a single byte preceded by ESC I

Internally, special bits are set to indicate that an escape sequence needs to be generated. Other bits are used to indicate whether the translated code should be influenced by some special keys.

Num Lock

If the Num Lock bit is set, the regular and Shift values are swapped, as are the Alt and Shift Alt values, whenever the Num Lock LED is on. By default, only the keys on the numeric keypad have this bit set. That is why these keys generate 7, 8, 9, etc. when the Num Lock LED is on, which is the same value that would be produced if Shift were used with these keys.

Caps Lock

This has the same effect as the Num Lock key. By default, this bit is set for all letters and not set for punctuation signs.

Ctrl When a key is translated into a single byte (no escape sequence) and this bit is set, the corresponding control character will be generated when the Ctrl key is pressed simultaneously. This is equally valid for the Shift, Alt, and Shift Alt combination. When this bit is not used, the Ctrl key combination will not generate anything.

mapfiles

This section describes the layout of a *mapfile* that is read by the **pcmapkeys** program.

A *mapfile* is a text file that consists of several **sections**. A sharp sign (#) can be used to include comments. Everything following the # until the end of the line will be ignored by the **pcmapkeys** program. Inside a line, C-style comments can be used as well. The beginning of each section is indicated by a *keyword*. Spaces and tabs are silently ignored and can be used at all times to improve readability. All but one section, the one that defines the *compose character*, can be left out. The order in which the different sections should appear is predefined. Here is the list of keywords in the order they should appear:

- input:**
- toggle:**
- dead:**
- compose:**
- output:**
- scancodes:**

Characters can be described in several different ways. ASCII characters can be described by putting them between single quotes. For example:

```
'a' '{'
```

Between single quotes, control characters can be listed by using a circumflex sign before the character that needs to be quoted. For example:

```
'^x'
```

When a backslash (\) is used, what follows will be interpreted as a decimal, octal (leading zero), or hexadecimal (leading x or X) representation of the character, although in this case the use of single quotes is not mandatory. For example:

```
'\x88'
```

is the same as:

```
0x88 (zero needed when not quoted)
```

and:

```
'\007'
```

is the same as:

```
007
```

When strings are needed, a list of character representations should be used. Quoted strings will be supported in the future.

The following paragraphs describe what goes in each section.

Input section

The input section describes which input characters should be mapped into a single byte. A very small sample input section could be:

input:

```
'A' 'B' # map A into B on input
```

```
'#' 0x9c # map sharp sign into pound sign
```

Toggle section

The toggle section is a one-line section that defines which key is to toggle between mapping and no mapping. For example:

toggle:

```
'^y' # ctrl y is the toggle key
```

Deadkey section

The deadkey section defines which keys should be treated as deadkeys. A **dead:** keyword followed by the specification of the character appears in this section for each deadkey. The subsequent lines describe what key should be generated for each key following the deadkey. A deadkey followed by a key not described in this part of the *mapfile* will not generate any key and a beep tone will be produced on the terminal. For example:

```
dead: '^' # circumflex is a deadkey
```

```
' ' '^' # circumflex followed by space generates circumflex
```

```
'e' 0x88 # circumflex followed by e generates e circumflex
```

```
dead: ""          # double quote used as a deadkey
'' ''           # double quote space generates double quote
'a' 0x84        # double quote a generates an umlaut
```

Compose section

The first line of this section describes what the compose character is. That line should always be present in the *mapfile*. Subsequent lines consist of three character representations indicating each time that the third character needs to be generated on input when the compose character is followed by the first two. Compose sequences with the same first character should be grouped together. For example:

```
compose: '^x'
"" 'e' 0x89      # e with umlaut is generated when typing ^x " e
"" 'a' 0x84      # a with umlaut
'e' "" 0x89      # e with umlaut is generated when typing ^x e "
'a' "" 0x84      # a with umlaut
```

The following example would give the wrong result. All lines starting with the same character specification should be grouped together.

```
compose: '^x'
"" 'e' 0x89      # e with umlaut is generated when typing ^x " e
'e' "" 0x89      # e with umlaut is generated when typing ^x e "
"" 'a' 0x84      # a with umlaut
'a' "" 0x84# a with umlaut
```

Output section

This section describes the mapping on output, either single byte to single byte, or single byte to string. A string is specified as a series of character specifications. For example:

```
output:
0x82 '{'          # map e with accent to { to display e with accent
'^u' ('K''I''L''L') # print (KILL) when kill character is used
```

Scancodes section

This section will only have an effect when your terminal is a scancode device. No error message will be produced if this section is in your *mapfile* when not needed, because the **pcmapkeys** program will find out whether the terminal is a scancode device or not. The lines in this section can have two different formats. One format will be used to describe what the values of the function keys must be. The other format describes the translation of scancodes into a byte or an escape sequence. No specific order is required.

Function keys

Here is an example of a line defining a string for a function key:

```
F13 'd''a''t''e''\n' # Shift F1 is the date command
```

The numbering convention of the function keys is described in a previous section. Currently, the use of quoted strings such as *"date\n"* is not supported.

Scancodes

Specifying how to translate a scancode is a more complex task. The general format of such a line is:

scancode normal shift alt shiftalt flags

scancode should list the hexadecimal representation of a scancode generated by a key (unquoted). How keys correspond with scancodes can be found in *keyboard(7)*.

normal, **shift**, **alt** and **shiftalt** are character representations in one of the formats described throughout this document, optionally followed by one of the following special keywords:

|C This indicates that the key is influenced by the Ctrl key.

|N This indicates that Esc N should precede the specified character.

|O This indicates that Esc O should precede the specified character.

|[This indicates that Esc [should precede the specified character.

The **normal** field defines how the scancode is translated when no other key is pressed, the **shift** field defines the translation for when the Shift key is used simultaneously, the **alt** field specifies what to do when the Alt key is pressed together with this and the **shiftalt** field contains the information on what to generate when both the Shift and Alt keys are pressed.

All five fields must be filled in. When no translation is requested (that is, the current active translation does not need to be changed) a dash (-) can be used. The sixth field is optional. This field can contain the special keyword CAPS or NUM or both, to indicate whether or not the Caps Lock key or Num Lock key status have any effect. Here is a sample line that describes the default translation for the 'Q' key:

```
0x10 'q'|C 'Q'|C 'q'|N 'Q'|N CAPS
```

If the normal or shift field is filled out for a scancode that represents a function key, a self-explanatory message will be produced and that translation information will be ignored.

A more detailed example of a **scancodes** section is:

scancodes:

the w key

```
0x11 'w'|C 'W'|C 'w'|N 'W'|N CAPS
```

left square bracket and curly brace key

control shift [does not generate anything (no C flag)

```
0x1a '['|C '{' '['|N '{'|N
```

9 on numeric keypad

```
0x49 'V'|[ '9' '9'|N '9'|N NUM
```

```
F13 'd''a''t''e''0 # SHIFT F1
```

More complete examples of *mapfiles* can be found in the `/usr/share/lib/keyboards` directory.

FILES

`/usr/share/lib/keyboards/8859/*`

sample mapfiles to be used in conjunction with ISO-8859-1 fonts (see `loadfont(1)`)

`/usr/share/lib/keyboards/437/*`

sample mapfiles to be used in conjunction with IBM 437 fonts (see **loadfont(1)**)

SEE ALSO **loadfont(1)**

NOTE The default keyboard mappings on the system are those of the ISO 8859-1 codeset. The optional IBM DOS 437 codeset is supported *only* at internationalization level 1. That is, if you choose to download keyboard mappings of the optional IBM DOS 437 codeset, there will be no support for non-standard U.S. date, time, currency, numbers, unit, and collation. There will be no support for non-English message and text presentation, and no multi-byte character support. Therefore, non-Windows users should only use IBM DOS 437 codeset in the default C locale.

NAME	pg – files perusal filter for CRTs
SYNOPSIS	pg [<i>-number</i>] [<i>-p string</i>] [<i>-cefnrs</i>] [<i>+linenumber</i>] [<i>+/pattern/</i>] [<i>filename ...</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The pg command is a filter that allows the examination of <i>filenames</i> one screenful at a time on a CRT. If no <i>filename</i> is specified or if it encounters the file name <i>-</i>, pg reads from standard input. Each screenful is followed by a prompt. If the user types a RETURN, another page is displayed; other possibilities are listed below.</p> <p>This command is different from previous paginators in that it allows you to back up and review something that has already passed. The method for doing this is explained below. To determine terminal attributes, pg scans the terminfo(4) data base for the terminal type specified by the environment variable TERM. If TERM is not defined, the terminal type dumb is assumed.</p>
OPTIONS	<p><i>-number</i> An integer specifying the size (in lines) of the window that pg is to use instead of the default. (On a terminal containing 24 lines, the default window size is 23).</p> <p><i>-p string</i> pg uses <i>string</i> as the prompt. If the prompt string contains a %d, the first occurrence of %d in the prompt will be replaced by the current page number when the prompt is issued. The default prompt string is “:”.</p> <p><i>-c</i> Home the cursor and clear the screen before displaying each page. This option is ignored if clear_screen is not defined for this terminal type in the terminfo(4) data base.</p> <p><i>-e</i> pg does <i>not</i> pause at the end of each file.</p> <p><i>-f</i> Normally, pg splits lines longer than the screen width, but some sequences of characters in the text being displayed (for instance, escape sequences for underlining) generate undesirable results. The <i>-f</i> option inhibits pg from splitting lines.</p> <p><i>-n</i> Normally, commands must be terminated by a <i><newline></i> character. This option causes an automatic end of command as soon as a command letter is entered.</p> <p><i>-r</i> Restricted mode. The shell escape is disallowed. pg prints an error message but does not exit.</p> <p><i>-s</i> pg prints all messages and prompts in the standard output mode (usually inverse video).</p> <p><i>+linenumber</i> Start up at <i>linenumber</i>.</p> <p><i>+/pattern/</i> Start up at the first line containing the regular expression pattern.</p>

USAGE

Commands

The responses that may be typed when **pg** pauses can be divided into three categories: those causing further perusal, those that search, and those that modify the perusal environment.

Commands that cause further perusal normally take a preceding *address*, an optionally signed number indicating the point from which further text should be displayed. This *address* is interpreted in either pages or lines depending on the command. A signed *address* specifies a point relative to the current page or line, and an unsigned *address* specifies an address relative to the beginning of the file. Each command has a default address that is used if none is provided.

The perusal commands and their defaults are as follows:

(+1)<*newline*> or <*blank*>

This causes one page to be displayed. The address is specified in pages.

(+1) **I** With a relative address this causes **pg** to simulate scrolling the screen, forward or backward, the number of lines specified. With an absolute address this command prints a screenful beginning at the specified line.

(+1) **d** or **^D**

Simulates scrolling half a screen forward or backward.

if Skip *i* screens of text.

iz Same as <*newline*> except that *i*, if present, becomes the new default number of lines per screenful.

The following perusal commands take no *address*.

. or **^L** Typing a single period causes the current page of text to be redisplayed.

\$ Displays the last windowful in the file. Use with caution when the input is a pipe.

The following commands are available for searching for text patterns in the text. The regular expressions described in **ed**(1) are available. They must always be terminated by a <*newline*>, even if the *-n* option is specified.

i/pattern/

Search forward for the *ith* (default *i*=1) occurrence of *pattern*. Searching begins immediately after the current page and continues to the end of the current file, without wrap-around.

^i pattern^

i?pattern?

Search backwards for the *ith* (default *i*=1) occurrence of *pattern*. Searching begins immediately before the current page and continues to the beginning of the current file, without wrap-around. The *^* notation is useful for Addis 100 terminals which will not properly handle the *?*.

After searching, **pg** will normally display the line found at the top of the screen. This can be modified by appending **m** or **b** to the search command to leave the line found in the middle or at the bottom of the window from now on. The suffix **t** can be used to restore the original situation.

The user of **pg** can modify the environment of perusal with the following commands:

in Begin perusing the *i*th next file in the command line. The *i* is an unsigned number, default value is 1.

ip Begin perusing the *i*th previous file in the command line. *i* is an unsigned number, default is 1.

iw Display another window of text. If *i* is present, set the window size to *i*.

s filename

Save the input in the named file. Only the current file being perused is saved. The white space between the **s** and *filename* is optional. This command must always be terminated by a *<newline>*, even if the **-n** option is specified.

h Help by displaying an abbreviated summary of available commands.

q or **Q** Quit **pg**.

!command

Command is passed to the shell, whose name is taken from the **SHELL** environment variable. If this is not available, the default shell is used. This command must always be terminated by a *<newline>*, even if the **-n** option is specified.

At any time when output is being sent to the terminal, the user can hit the quit key (normally CTRL-\) or the interrupt (break) key. This causes **pg** to stop sending output, and display the prompt. The user may then enter one of the above commands in the normal manner. Unfortunately, some output is lost when this is done, because any characters waiting in the terminal's output queue are flushed when the quit signal occurs.

If the standard output is not a terminal, then **pg** acts just like **cat(1)**, except that a header is printed before each file (if there is more than one).

EXAMPLES

The following command line uses **pg** to read the system news:

```
example% news | pg -p "(Page %d):"
```

ENVIRONMENT

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

NAME	pkginfo – display software package information
SYNOPSIS	<pre>pkginfo [-q -x -l] [-p -i] [-r] [-a arch] [-v version] [-c category1, [category2 [, ...]]] [pkginst [, pkginst [, ...]]] pkginfo [-d device] [-R root_path] [-q -x -l] [-a arch] [-v version] [-c category1, [category2 [, ...]]] [pkginst [, pkginst [, ...]]]</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>pkginfo displays information about software packages which are installed on the system (with the first synopsis) or which reside on a particular device or directory (with the second synopsis).</p> <p><i>pkginst</i> designates a package by its instance. An instance can be the package abbreviation or a specific instance (for example, inst.1 or inst.beta). All instances of package can be requested by inst.*.</p> <p>Without options, pkginfo lists the primary category, package instance, and the names of all completely installed and partially installed packages. It displays one line for each package selected.</p>
OPTIONS	<p>The -p and -i options are meaningless if used in conjunction with the -d option. The options -q, -x, and -l are mutually exclusive.</p> <p>-q Do not list any information. Used from a program to check whether or not a package has been installed.</p> <p>-x Designate an extracted listing of package information. The listing contains the package abbreviation, package name, package architecture (if available) and package version (if available).</p> <p>-l Specify long format, which includes all available information about the designated package(s).</p> <p>-p Display information for partially installed packages only.</p> <p>-i Display information for fully installed packages only.</p> <p>-r List the installation base for relocatable packages.</p> <p>-a arch Specify the architecture of the package as <i>arch</i>.</p> <p>-v version Specify the version of the package as <i>version</i>. All compatible versions can be requested by preceding the version name with a tilde (~). Multiple white spaces are replaced with a single white space during version comparison.</p> <p>-c category Display packages that match the category <i>category</i>. Categories are defined in the category field of the pkginfo file. If more than one category is supplied, the package needs to match only one category in the list. The match is not case specific.</p>

- d** *device* Defines a device, *device*, on which the software resides. *device* can be an absolute directory pathname or the identifiers for tape, floppy disk, removable disk, and so forth. The special token **spool** may be used to indicate the default installation spool directory (**/var/spool/pkg**).
- R** *root_path* Defines the full path name of a subdirectory to use as the *root_path*. All files, including package system information files, are relocated to a directory tree starting in the specified *root_path*.

SEE ALSO**pkgtrans(1), pkgadd(1M), pkgask(1M), pkgchk(1M), pkgrm(1M)**

NAME	pkgmk – produce an installable package
SYNOPSIS	pkgmk [-o] [-a arch] [-b basedir] [-d device] [-f prototype] [-l limit] [-p pstamp] [-r rootpath] [-v version] [<i>variable=value ...</i>] [<i>pkginst</i>]
DESCRIPTION	pkgmk produces an installable package to be used as input to the pkgadd command. The package contents will be in directory structure format. The command uses the package prototype file as input and creates a pkgmap file. The contents for each entry in the prototype file is copied to the appropriate output location. Information concerning the contents (checksum, file size, modification date) is computed and stored in the pkgmap file, along with attribute information specified in the prototype file.
OPTIONS	<p>-o Overwrite the same instance, package instance will be overwritten if it already exists.</p> <p>-a arch Override the architecture information provided in the pkginfo file with <i>arch</i>.</p> <p>-b basedir Prepend the indicated <i>basedir</i> to locate relocatable objects on the source machine.</p> <p>-d device Create the package on <i>device</i>. <i>device</i> can be an absolute directory path-name or the identifiers for a floppy disk or removable disk (for example, /dev/diskette). The default device is the installation spool directory (/var/spool/pkg).</p> <p>-f prototype Use the file prototype as input to the command. The default prototype filename is [Pp]prototype.</p> <p>-l limit Specify the maximum size in 512 byte blocks of the output device as <i>limit</i>. By default, if the output file is a directory or a mountable device, pkgmk will employ the df command to dynamically calculate the amount of available space on the output device. This option is useful in conjunction with pkgtrans to create package with datastream format.</p> <p>-p pstamp Override the production stamp definition in the pkginfo file with <i>psamp</i>.</p> <p>-r rootpath Ignore destination paths in the prototype file. Instead, use the indicated <i>rootpath</i> with the source pathname appended to locate objects on the source machine.</p> <p>-v version Override the version information provided in the pkginfo file with <i>version</i>.</p> <p><i>variable=value</i> Place the indicated variable in the packaging environment. (See prototype(4) for definitions of packaging variables.)</p> <p><i>pkginst</i> Specifies the package by its instance. An instance can be the package abbreviation or a specific instance (for example, inst.1).</p>

SEE ALSO **pkgparam(1)**, **pkgproto(1)**, **pkgtrans(1)**

NOTES

Architecture information is provided on the command line with the **-a** option or in the **prototype** file. If no architecture information is supplied, **pkgmk** uses the output of **uname -m**.

Version information is provided on the command line with the **-v** option or in the **prototype** file. If no version information is supplied, a default based on the current date will be provided.

Command line definitions for both architecture and version override the **prototype** definitions.

NAME	pkgparam – displays package parameter values
SYNOPSIS	pkgparam [<i>-v</i>] [<i>-R root_path</i>] [<i>-d device</i>] <i>pkginst</i> [<i>param ...</i>] pkgparam <i>-f filename</i> [<i>-v</i>] [<i>param ...</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>pkgparam displays the value associated with the parameter or parameters requested on the command line. The values are located in either the pkginfo file for <i>pkginst</i> or from the specific file named with the <i>-f</i> option.</p> <p>One parameter value is shown per line. Only the value of a parameter is given unless the <i>-v</i> option is used. With this option, the output of the command is in this format:</p> <pre style="margin-left: 40px;">parameter1='value1' parameter2='value2' parameter3='value3'</pre> <p>If no parameters are specified on the command line, values for all parameters associated with the package are shown.</p>
OPTIONS	<p>Options and arguments for this command are:</p> <ul style="list-style-type: none"> <i>-v</i> Verbose mode. Display name of parameter and its value. <i>-R root_path</i> Defines the full path name of a subdirectory to use as the <i>root_path</i>. All files, including package system information files, are relocated to a directory tree starting in the specified <i>root_path</i>. <i>-d device</i> Specify the <i>device</i> on which a <i>pkginst</i> is stored. It can be a directory path-name or the identifiers for tape, floppy disk or removable disk (for example, <i>/var/tmp</i>, <i>/dev/diskette</i>, and <i>/dev/dsk/c1d0s0</i>). The special token spool may be used to represent the default installation spool directory (<i>/var/spool/pkg</i>). <i>-f filename</i> Read <i>filename</i> for parameter values. <i>pkginst</i> Defines a specific package instance for which parameter values should be displayed. <i>param</i> Defines a specific parameter whose value should be displayed.
ERRORS	If parameter information is not available for the indicated package, the command exits with a non-zero status.
SEE ALSO	pkgtrans(1) , pkgmk(1) , pkgparam(1) , pkgproto(1)
NOTES	The <i>-f</i> synopsis allows you to specify the file from which parameter values should be extracted. This file should be in the same format as a pkginfo file. As an example, such a file might be created during package development and used while testing software during this stage.

NAME	pkgproto – generate prototype file entries for input to pkgmk command
SYNOPSIS	pkgproto [-i] [-c class] [<i>path1</i>] pkgproto [-i] [-c class] [<i>path1=path2 ...</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	pkgproto scans the indicated paths and generates prototype file entries that may be used as input to the pkgmk command.
OPTIONS	-i Ignores symbolic links and records the paths as ftype=f (a file) versus ftype=s(symbolic link) -c class Maps the class of all paths to <i>class</i> . <i>path1</i> Pathname where objects are located. <i>=path2</i> Pathname which should be substituted on output for <i>path1</i> . If no paths are specified on the command line, standard input is assumed to be a list of paths. If the pathname listed on the command line is a directory, the contents of the directory is searched. However, if input is read from stdin , a directory specified as a pathname will not be searched.
EXAMPLES	The following two examples show uses of pkgproto and a partial listing of the output produced. Example 1: <pre>example% pkgproto /bin=bin /usr/bin=usrbin /etc=etc f none bin/sed=/bin/sed 0775 bin bin f none bin/sh=/bin/sh 0755 bin daemon f none bin/sort=/bin/sort 0755 bin bin f none usrbin/sdb=/usr/bin/sdb 0775 bin bin f none usrbin/shl=/usr/bin/shl 4755 bin bin d none etc/master.d 0755 root daemon f none etc/master.d/kernel=/etc/master.d/kernel 0644 root daemon f none etc/rc=/etc/rc 0744 root daemon</pre> Example 2: <pre>example% find / -type d -print pkgproto d none / 755 root root d none /bin 755 bin bin d none /usr 755 root root d none /usr/bin 775 bin bin d none /etc 755 root root d none /tmp 777 root root</pre>

SEE ALSO**pkgmk(1)**, **pkgparam(1)**, **pkgtrans(1)****NOTES**

By default, **pkgproto** creates symbolic link entries for any symbolic link encountered (ftype=s). When you use the **-i** option, **pkgproto** creates a file entry for symbolic links (ftype=f). The **prototype** file would have to be edited to assign such file types as "v" (volatile), "e" (editable), or "x" (exclusive directory). **pkgproto** detects linked files. If multiple files are linked together, the first path encountered is considered the source of the link.

By default, **pkgproto** prints prototype entries on the standard output. However, the output should be saved in a file (named **Prototype** or **prototype**, for convenience) to be used as input to the **pkgmk** command.

NAME	pkgtrans – translate package format
SYNOPSIS	pkgtrans [-inos] <i>device1 device2</i> [<i>pkginst1</i> [<i>pkginst2</i>] ...]
AVAILABILITY	SUNWcsu
DESCRIPTION	pkgtrans translates an installable package from one format to another. It translates: <ul style="list-style-type: none"> a file system format to a datastream a datastream to a file system format one file system format to another file system format
OPTIONS	The options and arguments for this command are: <ul style="list-style-type: none"> -i Copy only the pkginfo and pkgmap files. -n Create a new instance of the package on the destination device if any instance of this package already exists, up to the number specified by -o Overwrite the same instance on the destination device; package instance will be overwritten if it already exists. MAXINST. -s Indicates that the package should be written to <i>device2</i> as a datastream rather than as a file system. The default behavior is to write a file system format on devices that support both formats. <p><i>device1</i> Indicates the source device. The package or packages on this device will be translated and placed on <i>device2</i>.</p> <p><i>device2</i> Indicates the destination device. Translated packages will be placed on this device.</p> <p><i>pkginst</i> Specifies which package instance or instances on <i>device1</i> should be translated. The token all may be used to indicate all packages. <i>pkginst.*</i> can be used to indicate all instances of a package. If no packages are defined, a prompt shows all packages on the device and asks which to translate.</p>
EXAMPLES	The following example translates all packages on the floppy drive /dev/diskette and places the translations on /tmp . <pre>example% pkgtrans /dev/diskette /tmp all</pre> <p>The next example translates packages pkg1 and pkg2 on /tmp and places their translations (that is, a datastream) on the 9track1 output device.</p> <pre>example% pkgtrans /tmp 9track1 pkg1 pkg2</pre> <p>The next example translates pkg1 and pkg2 on /tmp and places them on the diskette in a datastream format.</p> <pre>example% pkgtrans -s /tmp /dev/diskette pkg1 pkg2</pre>

ENVIRONMENT	The MAXINST variable is set in the pkginfo file and declares the maximum number of package instances.
SEE ALSO	pkginfo(1) , pkgmk(1) , pkgparam(1) , pkgproto(1) , installf(1M) , pkgadd(1M) , pkgask(1M) , pkgrm(1M) , removef(1M) <i>Application Packaging Developer's Guide</i>
NOTES	Device specifications can be either the special node name (for example, /dev/diskette) or a device alias (for example, diskette1). The device spool indicates the default spool directory. Source and destination devices cannot be the same. By default, pkgtrans will not translate any instance of a package if any instance of that package already exists on the destination device. Using the -n option creates a new instance if an instance of this package already exists. Using the -o option overwrites an instance of this package if it already exists. Neither of these options are useful if the destination device is a datastream.

NAME	plot, aedplot, atoplot, bgplot, crtplot, dumbplot, gigiplot, hpplot, implot, plottoa, t300, t300s, t4013, t450, tek, vplot, hp7221plot – graphics filters for various plotters
SYNOPSIS	<code>/usr/ucb/plot [-T<i>terminal</i>]</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>plot reads plotting instructions (see plot(4B)) from the standard input and produces plotting instructions suitable for a particular <i>terminal</i> on the standard output.</p> <p>If no <i>terminal</i> is specified, the environment variable TERM is used. The default <i>terminal</i> is tek.</p>
ENVIRONMENT	<p>Except for ver, the following terminal-types can be used with 'lpr -g' (see lpr) to produce plotted output:</p> <p>2648 2648a h8 hp2648 hp2648a Hewlett Packard® 2648 graphics terminal.</p> <p>hp7221 hp7 h7 Hewlett Packard® 7221 plotter.</p> <p>300 DASI 300 or GSI terminal (Diablo® mechanism).</p> <p>300s 300S DASI 300s terminal (Diablo® mechanism).</p> <p>450 DASI Hyterm 450 terminal (Diablo® mechanism).</p> <p>4013 Tektronix® 4013 storage scope.</p> <p>4014 tek Tektronix 4014 and 4015 storage scope with Enhanced Graphics Module. (Use 4013 for Tektronix® 4014 or 4015 without the Enhanced Graphics Module).</p> <p>aed AED 512 color graphics terminal.</p> <p>bgplot bitgraph BBN bitgraph graphics terminal.</p> <p>crt Any crt terminal capable of running vi(1).</p> <p>dumb un unknown Dumb terminals without cursor addressing or line printers.</p> <p>gigi vt125 DEC® vt125 terminal.</p> <p>h7 hp7 hp7221 Hewlett Packard® 7221 graphics terminal.</p> <p>implot Imagen plotter.</p> <p>var Benson Varian printer-plotter</p> <p>ver Versatec® D1200A printer-plotter. The output is scan-converted and suitable input to 'lpr -v'.</p>

FILES /usr/ucb/aedplot
/usr/ucb/atoplot
/usr/ucb/bgplot
/usr/ucb/crtplot
/usr/ucb/dumbplot
/usr/ucb/gigiplot
/usr/ucb/hp7221plot
/usr/ucb/hpplot
/usr/ucb/implot
/usr/ucb/plot
/usr/ucb/plottoa
/usr/ucb/t300
/usr/ucb/t300s
/usr/ucb/t4013
/usr/ucb/t450
/usr/ucb/tek
/usr/ucb/vplot
/var/ucb/vplotnnnnnn

SEE ALSO graph(1), vi(1), lpr(1B), plot(4B)

NAME	postdaisy – PostScript translator for Diablo 630 daisy-wheel files
SYNOPSIS	postdaisy [<i>-c num</i>] [<i>-f name</i>] [<i>-h num</i>] [<i>-m num</i>] [<i>-n num</i>] [<i>-o list</i>] [<i>-p mode</i>] [<i>-r num</i>] [<i>-s num</i>] [<i>-v num</i>] [<i>-x num</i>] [<i>-y num</i>] [<i>file ...</i>] <i>/usr/lib/lp/postscript/postdaisy</i>
DESCRIPTION	The postdaisy filter translates Diablo 630 daisy-wheel <i>files</i> into PostScript and writes the results on the standard output. If no <i>files</i> are specified, or if <i>-</i> is one of the input <i>files</i> , the standard input is read.
OPTIONS	<p><i>-c num</i> Print <i>num</i> copies of each page. By default only one copy is printed.</p> <p><i>-f name</i> Print <i>files</i> using font <i>name</i>. Any PostScript font can be used, although the best results will be obtained only with constant-width fonts. The default font is Courier.</p> <p><i>-h num</i> Set the initial horizontal motion index to <i>num</i>. Determines the character advance and the default point size, unless the <i>-s</i> option is used. The default is 12.</p> <p><i>-m num</i> Magnify each logical page by the factor <i>num</i>. Pages are scaled uniformly about the origin, which is located near the upper left corner of each page. The default magnification is 1.0.</p> <p><i>-n num</i> Print <i>num</i> logical pages on each piece of paper, where <i>num</i> can be any positive integer. By default, <i>num</i> is set to 1.</p> <p><i>-o list</i> Print pages whose numbers are given in the comma-separated <i>list</i>. The list contains single numbers <i>N</i> and ranges <i>N1 – N2</i>. A missing <i>N1</i> means the lowest numbered page, a missing <i>N2</i> means the highest. The page range is an expression of logical pages rather than physical sheets of paper. For example, if you are printing two logical pages to a sheet, and you specified a range of 4, then two sheets of paper would print, containing four page layouts. If you specified a page range of 3-4, when requesting two logical pages to a sheet; then <i>only</i> page 3 and page 4 layouts would print, and they would appear on one physical sheet of paper.</p> <p><i>-p mode</i> Print <i>files</i> in either portrait or landscape <i>mode</i>. Only the first character of <i>mode</i> is significant. The default <i>mode</i> is portrait.</p> <p><i>-r num</i> Selects carriage return and line feed behavior. If <i>num</i> is 1, a line feed generates a carriage return. If <i>num</i> is 2, a carriage return generates a line feed. Setting <i>num</i> to 3 enables both modes.</p> <p><i>-s num</i> Use point size <i>num</i> instead of the default value set by the initial horizontal motion index.</p> <p><i>-v num</i> Set the initial vertical motion index to <i>num</i>. The default is 8.</p> <p><i>-x num</i> Translate the origin <i>num</i> inches along the positive x axis. The default coordinate system has the origin fixed near the upper left corner of the page, with positive x to the right and positive y down the page. Positive <i>num</i> moves</p>

everything right. The default offset is **0.25** inches.

-y num Translate the origin *num* inches along the positive y axis. Positive *num* moves text up the page. The default offset is **-0.25** inches.

FILES **/usr/lib/lp/postscript/forms.ps**
/usr/lib/lp/postscript/ps.requests

SEE ALSO **download(1), dpost(1), postdmd(1), postio(1), postmd(1), postprint(1), postreverse(1), posttek(1)**

DIAGNOSTICS An exit status of **0** is returned if *files* were successfully processed.

NAME	postdmd – PostScript translator for DMD bitmap files
SYNOPSIS	<pre>postdmd [-b num] [-c num] [-f] [-m num] [-n num] [-o list] [-p mode] [-x num] [-y num] [file ...] /usr/lib/lp/postscript/postdmd</pre>
DESCRIPTION	<p>postdmd translates DMD bitmap <i>files</i>, as produced by <i>dmdps</i>, or <i>files</i> written in the Ninth Edition bitfile(9.5) format into PostScript and writes the results on the standard output. If no <i>files</i> are specified, or if - is one of the input <i>files</i>, the standard input is read.</p>
OPTIONS	<p>-b num Pack the bitmap in the output file using <i>num</i> byte patterns. A value of 0 turns off all packing of the output file. By default, <i>num</i> is 6.</p> <p>-c num Print <i>num</i> copies of each page. By default only one copy is printed.</p> <p>-f Flip the sense of the bits in <i>files</i> before printing the bitmaps.</p> <p>-m num Magnify each logical page by the factor <i>num</i>. Pages are scaled uniformly about the origin, which by default is located at the center of each page. The default magnification is 1.0.</p> <p>-n num Print <i>num</i> logical pages on each piece of paper, where <i>num</i> can be any positive integer. By default <i>num</i> is set to 1.</p> <p>-o list Print pages whose numbers are given in the comma-separated <i>list</i>. The list contains single numbers <i>N</i> and ranges <i>N1</i> – <i>N2</i>. A missing <i>N1</i> means the lowest numbered page, a missing <i>N2</i> means the highest. The page range is an expression of logical pages rather than physical sheets of paper. For example, if you are printing two logical pages to a sheet, and you specified a range of 4, then two sheets of paper would print, containing four page layouts. If you specified a page range of 3-4, when requesting two logical pages to a sheet; then <i>only</i> page 3 and page 4 layouts would print, and they would appear on one physical sheet of paper.</p> <p>-p mode Print <i>files</i> in either portrait or landscape <i>mode</i>. Only the first character of <i>mode</i> is significant. The default <i>mode</i> is portrait.</p> <p>-x num Translate the origin <i>num</i> inches along the positive x axis. The default coordinate system has the origin fixed at the center of the page, with positive x to the right and positive y up the page. Positive <i>num</i> moves everything right. The default offset is 0 inches.</p> <p>-y num Translate the origin <i>num</i> inches along the positive y axis. Positive <i>num</i> moves everything up the page. The default offset is 0.</p> <p>Only one bitmap is printed on each logical page, and each of the input <i>files</i> must contain complete descriptions of at least one bitmap. Decreasing the pattern size using the -b option may help throughput on printers with fast processors (such as PS-810s), while increasing the pattern size will often be the right move on older models (such as PS-800s).</p>

FILES /usr/lib/lp/postscript/forms.ps
/usr/lib/lp/postscript/ps.requests

SEE ALSO **download(1), dpost(1), postdaisy(1), postio(1), postmd(1), postprint(1), postreverse(1), posttek(1)**

DIAGNOSTICS An exit status of **0** is returned if *files* were successfully processed.

NAME	postio – serial interface for PostScript printers
SYNOPSIS	postio -l <i>line</i> [-D] [-i] [-q] [-t] [-S] [-b <i>speed</i>] [-B <i>num</i>] [-L <i>file</i>] [-P <i>string</i>] [-R <i>num</i>] [<i>file</i> ...] /usr/lib/lp/postscript/postio
DESCRIPTION	postio sends <i>files</i> to the PostScript printer attached to <i>line</i> . If no <i>files</i> are specified the standard input is sent.
OPTIONS	The first group of <i>options</i> should be sufficient for most applications: -D Enable debug mode. Guarantees that everything read on <i>line</i> will be added to the log file (standard error by default). -q Prevents status queries while <i>files</i> are being sent to the printer. When status queries are disabled a dummy message is appended to the log file before each block is transmitted. -b <i>speed</i> Transmit data over <i>line</i> at baud rate <i>speed</i> . Recognized baud rates are 1200, 2400, 4800, 9600, and 19200. The default <i>speed</i> is 9600 baud. -B <i>num</i> Set the internal buffer size for reading and writing <i>files</i> to <i>num</i> bytes. By default <i>num</i> is 2048 bytes. -l <i>line</i> Connect to the printer attached to <i>line</i> . In most cases there is no default and postio must be able to read and write <i>line</i> . If the <i>line</i> does not begin with a / it may be treated as a Datakit destination. -L <i>file</i> Data received on <i>line</i> gets put in <i>file</i> . The default log <i>file</i> is standard error. Printer or status messages that don't show a change in state are not normally written to <i>file</i> but can be forced out using the -D option. -P <i>string</i> Send <i>string</i> to the printer before any of the input files. The default <i>string</i> is simple PostScript code that disables timeouts. -R <i>num</i> Run <i>postio</i> as a single process if <i>num</i> is 1 or as separate read and write processes if <i>num</i> is 2. By default postio runs as a single process. The next two <i>options</i> are provided for users who expect to run postio on their own. Neither is suitable for use in spooler interface programs: -i Run the program in interactive mode. Any <i>files</i> are sent first and followed by the standard input. Forces separate read and write processes and overrides many other options. To exit interactive mode use your interrupt or quit character. To get a friendly interactive connection with the printer type executive on a line by itself. -t Data received on <i>line</i> and not recognized as printer or status information is written to the standard output. Forces separate read and write processes. Convenient if you have a PostScript program that will be returning useful data to the host.

The last option is not generally recommended and should only be used if all else fails to provide a reliable connection:

-S Slow the transmission of data to the printer. Severely limits throughput, runs as a single process, disables the **-q** option, limits the internal buffer size to 1024 bytes, can use an excessive amount of CPU time, and does nothing in interactive mode.

The best performance will usually be obtained by using a large internal buffer (the **-B** option) and by running the program as separate read and write processes (the **-R 2** option). Inability to fork the additional process causes **postio** to continue as a single read/write process. When one process is used, only data sent to the printer is flow controlled.

The *options* are not all mutually exclusive. The **-i** option always wins, selecting its own settings for whatever is needed to run interactive mode, independent of anything else found on the command line. Interactive mode runs as separate read and write processes and few of the other *options* accomplish anything in the presence of the **-i** option. The **-t** option needs a reliable two way connection to the printer and therefore tries to force separate read and write processes. The **-S** option relies on the status query mechanism, so **-q** is disabled and the program runs as a single process.

In most cases **postio** starts by making a connection to *line* and then attempts to force the printer into the IDLE state by sending an appropriate sequence of **^T** (status query), **^C** (interrupt), and **^D** (end of job) characters. When the printer goes IDLE, *files* are transmitted along with an occasional **^T** (unless the **-q** option was used). After all the *files* are sent the program waits until it's reasonably sure the job is complete. Printer generated error messages received at any time except while establishing the initial connection (or when running interactive mode) cause **postio** to exit with a non-zero status. In addition to being added to the log file, printer error messages are also echoed to standard error.

EXAMPLES

Run as a single process at 9600 baud and send *file1* and *file2* to the printer attached to **/dev/tty01**:

```
example% postio -l /dev/tty01 file1 file2
```

Same as above except two processes are used, the internal buffer is set to 4096 bytes, and data returned by the printer gets put in file *log*:

```
example% postio -R 2 -B 4096 -l /dev/tty01 -L log file1 file2
```

Establish an interactive connection with the printer at Datalog destination *my/printer*:

```
example% postio -i -l my/printer
```

Send file program to the printer connected to **/dev/tty22**, recover any data in file *results*, and put log messages in file *log*:

```
example% postio -t -l /dev/tty22 -L log program >results
```

SEE ALSO **download(1)**, **dpost(1)**, **postdaisy(1)**, **postdmd(1)**, **postmd(1)**, **postprint(1)**, **postreverse(1)**, **posttek(1)**

DIAGNOSTICS An exit status of **0** is returned if the files ran successfully. System errors (such as an inability to open the line) set the low order bit in the exit status, while PostScript errors set bit 1. An exit status of **2** usually means the printer detected a PostScript error in the input *files*.

NOTES The input *files* are handled as a single PostScript job. Sending several different jobs, each with their own internal end of job mark (^D) is not guaranteed to work properly. **postio** may quit before all the jobs have completed and could be restarted before the last one finishes.

All the capabilities described above may not be available on every machine or even across the different versions of the UNIX system that are currently supported by the program.

There may be no default *line*, so using the **-l** option is strongly recommended. If omitted, **postio** may attempt to connect to the printer using the standard output. If Datakit is involved, the **-b** option may be ineffective and attempts by **postio** to impose flow control over data in both directions may not work. The **-q** option can help if the printer is connected to RADIANT. The **-S** option is not generally recommended and should be used only if all other attempts to establish a reliable connection fail.

NAME	postmd – matrix display program for PostScript printers
SYNOPSIS	<pre>postmd [-b num] [-c num] [-d dimen] [-g list] [-i list] [-m num] [-n num] [-o list] [-p mode] [-w window] [-x num] [-y num] [file ...]</pre> <p>/usr/lib/lp/postscript/postmd</p>
DESCRIPTION	<p>The postmd filter reads a series of floating point numbers from <i>files</i>, translates them into a PostScript gray scale image, and writes the results on the standard output. In a typical application the numbers might be the elements of a large matrix, written in row major order, while the printed image could help locate patterns in the matrix. If no <i>files</i> are specified, or if – is one of the input <i>files</i>, the standard input is read.</p>
OPTIONS	<p>-b num Pack the bitmap in the output file using <i>num</i> byte patterns. A value of 0 turns off all packing of the output file. By default, <i>num</i> is 6.</p> <p>-c num Print <i>num</i> copies of each page. By default, only one copy is printed.</p> <p>-d dimen Sets the default matrix dimensions for all input <i>files</i> to <i>dimen</i>. The <i>dimen</i> string can be given as rows or rows\timescolumns. If <i>columns</i> is omitted it will be set to rows. By default, postmd assumes each matrix is square and sets the number of rows and columns to the square root of the number of elements in each input file.</p> <p>-g list <i>list</i> is a comma or space separated string of integers, each lying between 0 and 255 inclusive, that assigns PostScript gray scales to the regions of the real line selected by the -i option. 255 corresponds to white, and 0, to black. The postmd filter assigns a default gray scale that omits white (that is, 255) and gets darker as the regions move from left to right along the real line.</p> <p>-i list <i>list</i> is a comma, space or slash(/) separated string of <i>N</i> floating point numbers that partition the real line into $2N+1$ regions. The <i>list</i> must be given in increasing numerical order. The partitions are used to map floating point numbers read from the input <i>files</i> into gray scale integers that are either assigned automatically by postmd or arbitrarily selected using the -g option. The default interval <i>list</i> is -1,0,1, which partitions the real line into seven regions.</p> <p>-m num Magnify each logical page by the factor <i>num</i>. Pages are scaled uniformly about the origin which, by default, is located at the center of each page. The default magnification is 1.0.</p> <p>-n num Print <i>num</i> logical pages on each piece of paper, where <i>num</i> can be any positive integer. By default, <i>num</i> is set to 1.</p> <p>-o list Print pages whose numbers are given in the comma separated <i>list</i>. The list contains single numbers <i>N</i> and ranges <i>N1</i> – <i>N2</i>. A missing <i>N1</i> means the lowest numbered page, a missing <i>N2</i> means the highest. The page range is an expression of logical pages rather than physical sheets of paper. For example, if you are printing two logical pages to a sheet, and you specified a range of 4, then two sheets of paper would print, containing four page layouts. If you</p>

specified a page range of **3-4**, when requesting two logical pages to a sheet; then *only* page 3 and page 4 layouts would print, and they would appear on one physical sheet of paper.

-p mode Print *files* in either portrait or landscape *mode*. Only the first character of *mode* is significant. The default *mode* is portrait.

-w window

Window is a comma or space separated list of four positive integers that select the upper left and lower right corners of a submatrix from each of the input *files*. Row and column indices start at 1 in the upper left corner and the numbers in the input *files* are assumed to be written in row major order. By default, the entire matrix is displayed.

-x num Translate the origin *num* inches along the positive x axis. The default coordinate system has the origin fixed at the center of the page, with positive x to the right and positive y up the page. Positive *num* moves everything right. The default offset is **0** inches.

-y num Translate the origin *num* inches along the positive y axis. Positive *num* moves everything up the page. The default offset is **0**.

Only one matrix is displayed on each logical page, and each of the input *files* must contain complete descriptions of exactly one matrix. Matrix elements are floating point numbers arranged in row major order in each input file. White space, including newlines, is not used to determine matrix dimensions. By default, **postmd** assumes each matrix is square and sets the number of rows and columns to the square root of the number of elements in the input file. Supplying default dimensions on the command line with the **-d** option overrides this default behavior, and in that case the dimensions apply to all input *files*.

An optional header can be supplied with each input file and is used to set the matrix dimensions, the partition of the real line, the gray scale map, and a window into the matrix. The header consists of keyword/value pairs, each on a separate line. It begins on the first line of each input file and ends with the first unrecognized string, which should be the first matrix element. Values set in the header take precedence, but apply only to the current input file. Recognized header keywords are **dimension**, **interval**, **grayscale**, and **window**. The syntax of the value string that follows each keyword parallels what is accepted by the **-d**, **-i**, **-g**, and **-w** options.

EXAMPLES

For example, suppose file initially contains the 1000 numbers in a 20x50 matrix. Then you can produce exactly the same output by completing three steps. First, issue the following command line:

```
example% postmd -d20x50 -i"-100 100" -g0,128,254,128,0 file
```

Second, prepend the following header to *file*:

```
dimension 20x50
interval -100.0 .100e+3
grayscale 0 128 254 128 0
```

Third, issue the following command line:

example% **postmd** *file*

The interval list partitions the real line into five regions and the gray scale list maps numbers less than -100 or greater than 100 into 0 (that is, black), numbers equal to -100 or 100 into 128 (that is, 50 percent black), and numbers between -100 and 100 into 254 (that is, almost white).

FILES

/usr/lib/lp/postscript/forms.ps
/usr/lib/lp/postscript/ps.requests

SEE ALSO

dpost(1), postdaisy(1), postdmd(1), postio(1), postprint(1), postreverse(1), posttek(1)

DIAGNOSTICS

An exit status of **0** is returned if *files* were successfully processed.

NOTES

The largest matrix that can be adequately displayed is a function of the interval and gray scale lists, the printer resolution, and the paper size. A 600x600 matrix is an optimistic upper bound for a two element interval list (that is, five regions) using 8.5x11 inch paper on a 300 dpi printer.

Using white (that is, 255) in a gray scale list is not recommended and won't show up in the legend and bar graph that **postmd** displays below each image.

NAME	postplot – PostScript translator for plot(4) graphics files
SYNOPSIS	postplot [-c <i>num</i>] [-f <i>name</i>] [-m <i>num</i>] [-n <i>num</i>] [-o <i>list</i>] [-p <i>mode</i>] [-w <i>num</i>] [-x <i>num</i>] [-y <i>num</i>] [<i>filename ...</i>] /usr/lib/lp/postscript/postplot
DESCRIPTION	The postplot filter translates plot (1B) graphics <i>filenames</i> into PostScript and writes the results on the standard output. If no <i>filenames</i> are specified, or if - is one of the input <i>filenames</i> , the standard input is read.
OPTIONS	<p>-c <i>num</i> Print <i>num</i> copies of each page. By default, only one copy is printed.</p> <p>-f <i>name</i> Print text using font <i>name</i>. Any PostScript font can be used, although the best results will be obtained only with constant width fonts. The default font is Courier.</p> <p>-m <i>num</i> Magnify each logical page by the factor <i>num</i>. Pages are scaled uniformly about the origin which, by default, is located at the center of each page. The default magnification is 1.0.</p> <p>-n <i>num</i> Print <i>num</i> logical pages on each piece of paper, where <i>num</i> can be any positive integer. By default, <i>num</i> is set to 1.</p> <p>-o <i>list</i> Print pages whose numbers are given in the comma-separated <i>list</i>. The list contains single numbers <i>N</i> and ranges <i>N1</i> – <i>N2</i>. A missing <i>N1</i> means the lowest numbered page, a missing <i>N2</i> means the highest.</p> <p>-p <i>mode</i> Print <i>filenames</i> in either portrait or landscape <i>mode</i>. Only the first character of <i>mode</i> is significant. The default <i>mode</i> is landscape.</p> <p>-w <i>num</i> Set the line width used for graphics to <i>num</i> points, where a point is approximately 1/72 of an inch. By default, <i>num</i> is set to 0 points, which forces lines to be one pixel wide.</p> <p>-x <i>num</i> Translate the origin <i>num</i> inches along the positive x axis. The default coordinate system has the origin fixed at the center of the page, with positive x to the right and positive y up the page. Positive <i>num</i> moves everything right. The default offset is 0.0 inches.</p> <p>-y <i>num</i> Translate the origin <i>num</i> inches along the positive y axis. Positive <i>num</i> moves everything up the page. The default offset is 0.0.</p>
FILES	/usr/lib/lp/postscript/forms.ps /usr/lib/lp/postscript/postplot.ps /usr/lib/lp/postscript/ps.requests

SEE ALSO

download(1), dpost(1), plot(1B), postdaisy(1), postdmd(1), postio(1), postmd(1), postprint(1), postreverse(1)

DIAGNOSTICS

An exit status of 0 is returned if *filenames* were successfully processed.

NOTES

The default line width is too small for write-white print engines, such as the one used by the PS-2400.

NAME	postprint – PostScript translator for text files
SYNOPSIS	<pre>postprint [-c num] [-f name] [-l num] [-m num] [-n num] [-o list] [-p mode] [-r num] [-s num] [-t num] [-x num] [-y num] [file...]</pre> <p>/usr/lib/lp/postscript/postprint</p>
DESCRIPTION	The postprint filter translates text <i>files</i> into PostScript and writes the results on the standard output. If no <i>files</i> are specified, or if – is one of the input <i>files</i> , the standard input is read.
OPTIONS	<p>-c num Print <i>num</i> copies of each page. By default, only one copy is printed.</p> <p>-f name Print <i>files</i> using font <i>name</i>. Any PostScript font can be used, although the best results will be obtained only with constant width fonts. The default font is Courier.</p> <p>-l num Set the length of a page to <i>num</i> lines. By default, <i>num</i> is 66. Setting <i>num</i> to 0 is allowed, and will cause <i>postprint</i> to guess a value, based on the point size that's being used.</p> <p>-m num Magnify each logical page by the factor <i>num</i>. Pages are scaled uniformly about the origin, which is located near the upper left corner of each page. The default magnification is 1.0.</p> <p>-n num Print <i>num</i> logical pages on each piece of paper, where <i>num</i> can be any positive integer. By default, <i>num</i> is set to 1.</p> <p>-o list Print pages whose numbers are given in the comma-separated <i>list</i>. The <i>list</i> contains single numbers <i>N</i> and ranges <i>N1</i> – <i>N2</i>. A missing <i>N1</i> means the lowest numbered page, a missing <i>N2</i> means the highest. The page range is an expression of logical pages rather than physical sheets of paper. For example, if you are printing two logical pages to a sheet, and you specified a range of 4, then two sheets of paper would print, containing four page layouts. If you specified a page range of 3-4, when requesting two logical pages to a sheet; then <i>only</i> page 3 and page 4 layouts would print, and they would appear on one physical sheet of paper.</p> <p>-p mode Print <i>files</i> in either portrait or landscape <i>mode</i>. Only the first character of <i>mode</i> is significant. The default <i>mode</i> is portrait.</p> <p>-r num Selects carriage return behavior. Carriage returns are ignored if <i>num</i> is 0, cause a return to column 1 if <i>num</i> is 1, and generate a newline if <i>num</i> is 2. The default <i>num</i> is 0.</p> <p>-s num Print <i>files</i> using point size <i>num</i>. When printing in landscape mode <i>num</i> is scaled by a factor that depends on the imaging area of the device. The default size for portrait mode is 10. Note that increasing point size increases virtual image size, so you either need to load larger paper, or use the -10 option to scale the number of lines per page.</p>

- t num** Assume tabs are set every *num* columns, starting with the first column. By default, tabs are set every **8** columns.
- x num** Translate the origin *num* inches along the positive x axis. The default coordinate system has the origin fixed near the upper left corner of the page, with positive x to the right and positive y down the page. Positive *num* moves everything to the right. The default offset is **0.25** inches.
- y num** Translate the origin *num* inches along the positive y axis. Positive *num* moves text up the page. The default offset is **-0.25** inches.

A new logical page is started after 66 lines have been printed on the current page, or whenever an ASCII form feed character is read. The number of lines per page can be changed using the **-l** option. Unprintable ASCII characters are ignored, and lines that are too long are silently truncated by the printer.

EXAMPLES

To print *file1* and *file2* in landscape mode, issue the following command:

```
example% postprint -pland file1 file2
```

To print three logical pages on each physical page in portrait mode:

```
example% postprint -n3 file
```

FILES

```
/usr/lib/lp/postscript/forms.ps  
/usr/lib/lp/postscript/ps.requests
```

SEE ALSO

download(1), **dpost(1)**, **postdaisy(1)**, **postdmd(1)**, **postio(1)**, **postmd(1)**, **postreverse(1)**, **posttek(1)**

DIAGNOSTICS

An exit status of **0** is returned if *files* were successfully processed.

NAME	postreverse – reverse the page order in a PostScript file
SYNOPSIS	postreverse [-o <i>list</i>] [-r] [<i>file</i>] /usr/lib/lp/postscript/postreverse
DESCRIPTION	<p>The postreverse filter reverses the page order in files that conform to Adobe's Version 1.0 or Version 2.0 file structuring conventions, and writes the results on the standard output. Only one input <i>file</i> is allowed and if no <i>file</i> is specified, the standard input is read.</p> <p>The postreverse filter can handle a limited class of files that violate page independence, provided all global definitions are bracketed by %%BeginGlobal and %%EndGlobal comments. In addition, files that mark the end of each page with %%EndPage: label ordinal comments will also reverse properly, provided the prologue and trailer sections can be located. If postreverse fails to find an %%EndProlog or %%EndSetup comment, the entire <i>file</i> is copied, unmodified, to the standard output.</p> <p>Because global definitions are extracted from individual pages and put in the prologue, the output file can be minimally conforming, even if the input <i>file</i> was not.</p>
OPTIONS	<p>-o <i>list</i> Select pages whose numbers are given in the comma-separated <i>list</i>. The <i>list</i> contains single numbers <i>N</i> and ranges <i>N1</i> – <i>N2</i>. A missing <i>N1</i> means the lowest numbered page, a missing <i>N2</i> means the highest. The page range is an expression of logical pages rather than physical sheets of paper. For example, if you are printing two logical pages to a sheet, and you specified a range of 4, then two sheets of paper would print, containing four page layouts. If you specified a page range of 3-4, when requesting two logical pages to a sheet; then <i>only</i> page 3 and page 4 layouts would print, and they would appear on one physical sheet of paper.</p> <p>-r Do not reverse the pages in <i>file</i>.</p>
EXAMPLES	<p>To select pages 1 to 100 from <i>file</i> and reverse the pages: example% postreverse -o1-100 file</p> <p>To print four logical pages on each physical page and reverse all the pages: example% postprint -n4 file postreverse</p> <p>To produce a minimally conforming file from output generated by dpost without reversing the pages: example% dpost file postreverse -r</p>
SEE ALSO	download(1) , dpost(1) , postdaisy(1) , postdmd(1) , postio(1) , postmd(1) , postprint(1) , posttek(1)
DIAGNOSTICS	An exit status of 0 is returned if <i>file</i> was successfully processed.

NOTES

No attempt has been made to deal with redefinitions of global variables or procedures. If standard input is used, the input *file* will be read three times before being reversed.

NAME	posttek – PostScript translator for Tektronix 4014 files
SYNOPSIS	<pre>posttek [-c num] [-f name] [-m num] [-n num] [-o list] [-p mode] [-w num] [-x num] [-y num] [file ...]</pre> <p>/usr/lib/lp/postscript/posttek</p>
DESCRIPTION	The posttek filter translates Tektronix 4014 graphics <i>files</i> into PostScript and writes the results on the standard output. If no <i>files</i> are specified, or if - is one of the input <i>files</i> , the standard input is read.
OPTIONS	<p>-c num Print <i>num</i> copies of each page. By default, only one copy is printed.</p> <p>-f name Print text using font <i>name</i>. Any PostScript font can be used, although the best results will be obtained only with constant width fonts. The default font is Courier.</p> <p>-m num Magnify each logical page by the factor <i>num</i>. Pages are scaled uniformly about the origin which, by default, is located at the center of each page. The default magnification is 1.0.</p> <p>-n num Print <i>num</i> logical pages on each piece of paper, where <i>num</i> can be any positive integer. By default, <i>num</i> is set to 1.</p> <p>-o list Print pages whose numbers are given in the comma-separated <i>list</i>. The <i>list</i> contains single numbers <i>N</i> and ranges <i>N1</i> – <i>N2</i>. A missing <i>N1</i> means the lowest numbered page, a missing <i>N2</i> means the highest. The page range is an expression of logical pages rather than physical sheets of paper. For example, if you are printing two logical pages to a sheet, and you specified a range of 4, then two sheets of paper would print, containing four page layouts. If you specified a page range of 3-4, when requesting two logical pages to a sheet; then <i>only</i> page 3 and page 4 layouts would print, and they would appear on one physical sheet of paper.</p> <p>-p mode Print <i>files</i> in either portrait or landscape <i>mode</i>. Only the first character of <i>mode</i> is significant. The default <i>mode</i> is landscape.</p> <p>-w num Set the line width used for graphics to <i>num</i> points, where a point is approximately 1/72 of an inch. By default, <i>num</i> is set to 0 points, which forces lines to be one pixel wide.</p> <p>-x num Translate the origin <i>num</i> inches along the positive x axis. The default coordinate system has the origin fixed at the center of the page, with positive x to the right and positive y up the page. Positive <i>num</i> moves everything right. The default offset is 0.0 inches.</p> <p>-y num Translate the origin <i>num</i> inches along the positive y axis. Positive <i>num</i> moves everything up the page. The default offset is 0.0.</p>

FILES	/usr/lib/lp/postscript/forms.ps /usr/lib/lp/postscript/ps.requests
SEE ALSO	download(1), dpost(1), postdaisy(1), postdmd(1), postio(1), postmd(1), postprint(1), postreverse(1)
DIAGNOSTICS	An exit status of 0 is returned if <i>files</i> were successfully processed.
NOTES	The default line width is too small for write-white print engines, such as the one used by the PS-2400.

NAME	pr – print files														
SYNOPSIS	<pre>pr [[-columns] [-wwidth] [-a]] [-eck] [-ick] [-drftp] [+page] [-nck] [-offset] [-llength] [-sseparator] [-hheader] [-F] [filename ...] pr [[-m] [-wwidth]] [-eck] [-ick] [-drftp] [+page] [-nck] [-offset] [-llength] [-sseparator] [-hheader] [-F] [filename1 filename2 ...]</pre>														
AVAILABILITY	SUNWcsu														
DESCRIPTION	<p>The pr command formats and prints the contents of a file. If <i>filename</i> is -, or if no files are specified, pr assumes standard input. pr prints the named files on standard output.</p> <p>By default, the listing is separated into pages, each headed by the page number, the date and time that the file was last modified, and the name of the file. Page length is 66 lines which includes 10 lines of header and trailer output. The header is composed of 2 blank lines, 1 line of text (can be altered with -h), and 2 blank lines; the trailer is 5 blank lines. For single column output, line width may not be set and is unlimited. For multicolumn output, line width may be set and the default is 72 columns. Diagnostic reports (failed options) are reported at the end of standard output associated with a terminal, rather than interspersed in the output. Pages are separated by series of line feeds rather than form feed characters.</p> <p>By default, columns are of equal width, separated by at least one space; lines which do not fit are truncated. If the -s option is used, lines are not truncated and columns are separated by the <i>separator</i> character.</p> <p>Either -columns or -m should be used to produce multi-column output. -a should only be used with -columns and not -m.</p>														
OPTIONS	<table border="0"> <tr> <td style="vertical-align: top;">+page</td> <td>Begin printing with page numbered <i>page</i> (default is 1).</td> </tr> <tr> <td style="vertical-align: top;">-columns</td> <td>Print <i>columns</i> columns of output (default is 1). For example pr -3 produces the following output:</td> </tr> <tr> <td></td> <td style="text-align: center;"> <pre>Print of in the one three lines file columns.</pre> </td> </tr> <tr> <td></td> <td>Columns are not balanced; if, for example, there are as many lines in the file as there are lines on the page, only one column will be printed.</td> </tr> <tr> <td></td> <td>Output appears as if -e and -i are on for multi-column output. May not use with -m.</td> </tr> <tr> <td style="vertical-align: top;">-a</td> <td>Print multi-column output across the page one line per column. <i>columns</i> must be greater than one. If a line is too long to fit in a column, it is truncated.</td> </tr> <tr> <td style="vertical-align: top;">-m</td> <td>Merge and print all files simultaneously, one per column. The maximum number of files that may be specified is eight. If a line is too long to fit in a column, it is truncated. May not use with -column.</td> </tr> </table>	+page	Begin printing with page numbered <i>page</i> (default is 1).	-columns	Print <i>columns</i> columns of output (default is 1). For example pr -3 produces the following output:		<pre>Print of in the one three lines file columns.</pre>		Columns are not balanced; if, for example, there are as many lines in the file as there are lines on the page, only one column will be printed.		Output appears as if -e and -i are on for multi-column output. May not use with -m .	-a	Print multi-column output across the page one line per column. <i>columns</i> must be greater than one. If a line is too long to fit in a column, it is truncated.	-m	Merge and print all files simultaneously, one per column. The maximum number of files that may be specified is eight. If a line is too long to fit in a column, it is truncated. May not use with -column .
+page	Begin printing with page numbered <i>page</i> (default is 1).														
-columns	Print <i>columns</i> columns of output (default is 1). For example pr -3 produces the following output:														
	<pre>Print of in the one three lines file columns.</pre>														
	Columns are not balanced; if, for example, there are as many lines in the file as there are lines on the page, only one column will be printed.														
	Output appears as if -e and -i are on for multi-column output. May not use with -m .														
-a	Print multi-column output across the page one line per column. <i>columns</i> must be greater than one. If a line is too long to fit in a column, it is truncated.														
-m	Merge and print all files simultaneously, one per column. The maximum number of files that may be specified is eight. If a line is too long to fit in a column, it is truncated. May not use with -column .														

- d** Double-space the output. Blank lines that result from double-spacing are dropped when they occur at the top of a page.
- eck** Expand input tabs to character positions $k+1$, $2*k+1$, $3*k+1$, etc. If k is 0 or is omitted, default tab settings at every eighth position are assumed. Tab characters in the input are expanded into the appropriate number of spaces. If c (any non-digit character) is given, it is treated as the input tab character (default for c is the tab character).
- ick** In output, replace white space wherever possible by inserting tabs to character positions $k+1$, $2*k+1$, $3*k+1$, etc. If k is 0 or is omitted, default tab settings at every eighth position are assumed. If c (any non-digit character) is given, it is treated as the output tab character (default for c is the tab character).
- nck** Provide k -digit line numbering (default for k is 5). The number occupies the first $k+1$ character positions of each column of single column output or each line of **-m** output. If c (any non-digit character) is given, it is appended to the line number to separate it from whatever follows (default for c is a tab).
- width** Set the width of a line to $width$ character positions (default is 72). This is effective only for multi-column output (**-column** and **-m**). There is no line limit for single column output.
- offset** Offset each line by $offset$ character positions (default is 0). The number of character positions per line is the sum of the width and offset.
- length** Set the length of a page to $length$ lines (default is 66). A $length$ of 0 specifies the default length. By default, output contains 5 lines of header and 5 lines of trailer leaving 56 lines for user-supplied text. When **-length** is used and $length$ exceeds 10, then $length-10$ lines are left per page for user supplied text. When $length$ is 10 or less, header and trailer output is omitted to make room for user supplied text; see the **-t** option.
- h header** Use *header* as the text line of the header to be printed instead of the file name. **-h** is ignored when **-t** is specified or **-length** is specified and the value of $length$ is 10 or less. (**-h** is the only **pr** option requiring space between the option and argument.)
- p** Pause before beginning each page if the output device is a terminal. **pr** rings the terminal bell and waits for a carriage return.
- f** Use a single form-feed character for new pages (default is to use a sequence of line feeds). Pause before beginning the first page if the standard output is associated with a terminal.
- r** Print no diagnostic reports on files that cannot be opened.

- t** Print neither the five-line identifying header nor the five-line trailer normally supplied for each page. Quit printing after the last line of each file without spacing to the end of the page. Use of **-t** overrides the **-h** option.
- separator** Separate columns by the single character *separator* instead of by the appropriate number of spaces (default for *separator* is a tab). Prevents truncation of lines on multicolumn output unless **-w** is specified.
- F** Fold the lines of the input file. When used in multi-column mode (with the **-a** or **-m** options) lines will be folded to fit the current column's width, otherwise they will be folded to fit the current line width (80 columns).

EXAMPLES

Print **filename1** and **filename2** as a double-spaced, three-column listing headed by "file list":

```
example% pr -3dh "file list" filename1 filename2
```

Copy **filename1** to **filename2**, expanding tabs to columns 10, 19, 28, 37, ... :

```
example% pr -e9 -t < filename1 > filename2
```

Print **filename1** and **filename2** simultaneously in a two-column listing with no header or trailer where both columns have line numbers:

```
example% pr -t -n filename1 | pr -t -m -n filename2 -
```

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

LC_TIME

Determines how **pr** handles date and time formats. In the "C" locale, date and time handling follows the U.S. rules.

FILES | **/dev/tty*** | If standard output is directed to one of the special files **/dev/tty***, then other output directed to this terminal is delayed until standard output is completed. This prevents error messages from being interspersed throughout the output.

SEE ALSO | **cat(1), fold(1), more(1), pg(1), environ(5)**

NAME	print – shell built-in function to output characters to the screen or window
SYNOPSIS ksh	print [-Rnrpsu [<i>n</i>]] [<i>arg</i> . . .]
DESCRIPTION ksh	<p>The shell output mechanism. With no flags or with flag - or --, the arguments are printed on standard output as described by echo(1). The exit status is 0, unless the output file is not open for writing.</p> <p>-n suppresses new-line from being added to the output.</p> <p>-R</p> <p>-r (raw mode) ignore the escape conventions of echo. The -R option will print all subsequent arguments and options other than -n.</p> <p>-p causes the arguments to be written onto the pipe of the process spawned with & instead of standard output.</p> <p>-s causes the arguments to be written onto the history file instead of standard output.</p> <p>-u [<i>n</i>] flag can be used to specify a one digit file descriptor unit number <i>n</i> on which the output will be placed. The default is 1.</p>
SEE ALSO	echo (1), ksh (1)

NAME	printenv – display environment variables currently set
SYNOPSIS	/usr/ucb/printenv [<i>variable</i>]
AVAILABILITY	SUNWscpu
DESCRIPTION	printenv prints out the values of the variables in the environment. If a <i>variable</i> is specified, only its value is printed.
SEE ALSO	cs (1), echo (1), sh (1), stty (1), tset (1B), environ (5)
DIAGNOSTICS	If a <i>variable</i> is specified and it is not defined in the environment, printenv returns an exit status of 1 .

NAME	printf – print formatted output
SYNOPSIS	printf <i>format</i> [<i>argument</i> . . .]
AVAILABILITY	SUNWloc
DESCRIPTION	The printf command converts, formats, and prints its <i>arguments</i> under control of the <i>format</i> . It fully supports conversion specifications for strings (%s descriptor); however, the results are undefined for the other conversion specifications supported by printf(3S) .
OPTIONS	<p><i>format</i> A character string that contains three types of objects: 1) plain characters, which are simply copied to the output stream; 2) conversion specifications, each of which results in fetching zero or more <i>arguments</i>; and 3) C-language escape sequences, which are translated into the corresponding characters.</p> <p><i>argument</i> String(s) to be printed under the control of <i>format</i>. The results are undefined if there are insufficient <i>arguments</i> for the format. If the format is exhausted while <i>arguments</i> remain, the excess <i>arguments</i> are simply ignored.</p> <p>Each conversion specification is introduced by the character %. After the %, the following appear in sequence:</p> <ul style="list-style-type: none">An optional field, consisting of a decimal digit string followed by a \$, specifying the next <i>argument</i> to be converted. If this field is not provided, the <i>argument</i> following the last <i>argument</i> converted is used.An optional decimal digit string specifying a minimum <i>field width</i>. If the converted value has fewer characters than the field width, it is padded on the left (or right, if the left-adjustment flag ‘-’ has been given) to the field width. The padding is with blanks unless the field width digit string starts with a zero, in which case the padding is with zeros.An optional <i>precision</i> that gives the maximum number of characters to be printed from a string in %s conversion. The precision takes the form of a period (.) followed by a decimal digit string; a null digit string is treated as zero (nothing is printed). Padding specified by the precision overrides the padding specified by the field width. That is, if <i>precision</i> is specified, its value is used to control the number of characters printed. <p>A field width or precision or both may be indicated by an asterisk (*) instead of a digit string. In this case, an integer <i>argument</i> supplies the field width or precision. The <i>argument</i> that is actually converted is not fetched until the conversion letter is seen, so the <i>arguments</i> specifying field width or precision must appear <i>before</i> the <i>argument</i> (if any) to be converted. A negative field width argument is taken as a ‘-’ (left-adjustment) flag followed by a positive field width. If the precision argument is negative, it is changed to zero (nothing is printed). In no case does a non-existent or small field width cause truncation of a field; if the result of a conversion is wider than the field width, the field is simply expanded to contain the conversion result.</p>

The conversion characters and their meanings are:

- %s** The *argument* is taken to be a string and characters from the string are printed until a null character (`\0`) is encountered or the number of characters indicated by the precision specification is reached. If the precision is missing, it is taken to be infinite, so all characters up to the first null character are printed. A null value for *argument* yields undefined results.
- %%** Print a %; no argument is converted.

EXAMPLES

The command

```
example% printf '%s %s %s\n' Good Morning World
```

results in the output:

```
Good Morning World
```

The following command produces the same output.

```
example% printf '%2$s %s %1$s\n' World Good Morning
```

Here is an example that prints the first 6 characters of `$PATH` left-adjusted in a 10-character field:

```
example% printf 'First 6 chars of %s are %-10.6s.\n' $PATH $PATH
```

If `$PATH` has the value `/usr/bin:/usr/local/bin`, then the above command would print the following output:

```
First 6 chars of /usr/bin:/usr/local/bin are /usr/b .
```

SEE ALSO

printf(3S)

NAME	priocntl – display or set scheduling parameters of specified process(es)
SYNOPSIS	<p>priocntl -l priocntl -d [-i <i>idtype</i>] [<i>idlist</i>] priocntl -s [-c <i>class</i>] [<i>class-specific options</i>] [-i <i>idtype</i>] [<i>idlist</i>] priocntl -e [-c <i>class</i>] [<i>class-specific options</i>] <i>command</i> [<i>argument(s)</i>]</p>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The priocntl command displays or sets scheduling parameters of the specified process(es). It can also be used to display the current configuration information for the system's process scheduler or execute a command with specified scheduling parameters. Processes fall into distinct classes with a separate scheduling policy applied to each class. The process classes currently supported are the real-time class, time-sharing class, and the interactive class. The characteristics of these classes and the class-specific options they accept are described below in the USAGE section under the headings Real-Time Class, Time-Sharing Class, and Inter-Active Class. With appropriate permissions, the priocntl command can change the class and other scheduling parameters associated with a running process.</p> <p>In the default configuration, a runnable real-time process runs before any other process. Therefore, inappropriate use of real-time processes can have a dramatic negative impact on system performance.</p> <p>If an <i>idlist</i> is present it must appear last on the command line and the elements of the list must be separated by white space. If no <i>idlist</i> is present an <i>idtype</i> argument of pid, ppid, pgid, sid, class, uid, or gid specifies the process ID, parent process ID, process group ID, session ID, class, user ID, or group ID respectively of the priocntl command itself.</p> <p>The command</p> <p style="padding-left: 40px;">priocntl -d [-i <i>idtype</i>] [<i>idlist</i>]</p> <p>displays the class and class-specific scheduling parameters of the process(es) specified by <i>idtype</i> and <i>idlist</i>.</p> <p>The command</p> <p style="padding-left: 40px;">priocntl -s [-c <i>class</i>] [<i>class-specific options</i>] [-i <i>idtype</i>] [<i>idlist</i>]</p> <p>sets the class and class-specific parameters of the specified processes to the values given on the command line. The -c class option specifies the class to be set. (The valid <i>class</i> arguments are RT for real-time TS for time-sharing or IA for inter-active.)</p> <p>The class-specific parameters to be set are specified by the class-specific options as explained under the appropriate heading below. If the -c class option is omitted, <i>idtype</i> and <i>idlist</i> must specify a set of processes which are all in the same class, otherwise an error results. If no class-specific options are specified the process's class-specific parameters are set to the default values for the class specified by -c class (or to the default parameter values for the process's current class if the -c class option is also omitted).</p>

In order to change the scheduling parameters of a process using **prioctl** the real or effective user ID of the user invoking **prioctl** must match the real or effective user ID of the receiving process or the effective user ID of the user must be super-user. These are the minimum permission requirements enforced for all classes. An individual class may impose additional permissions requirements when setting processes to that class or when setting class-specific scheduling parameters.

When *idtype* and *idlist* specify a set of processes, **prioctl** acts on the processes in the set in an implementation-specific order. If **prioctl** encounters an error for one or more of the target processes, it may or may not continue through the set of processes, depending on the nature of the error.

If the error is related to permissions, **prioctl** prints an error message and then continue through the process set, resetting the parameters for all target processes for which the user has appropriate permissions. If **prioctl** encounters an error other than permissions, it does not continue through the process set but prints an error message and exits immediately.

A special **sys** scheduling class exists for the purpose of scheduling the execution of certain special system processes (such as the swapper process). It is not possible to change the class of any process to **sys**. In addition, any processes in the **sys** class that are included in the set of processes specified by *idtype* and *idlist* are disregarded by **prioctl**. For example, if *idtype* were **uid**, an *idlist* consisting of a zero would specify all processes with a UID of zero except processes in the **sys** class and (if changing the parameters using the **-s** option) the **init** process.

The **init** process (process ID 1) is a special case. In order for the **prioctl** command to change the class or other scheduling parameters of the **init** process, *idtype* must be **pid** and *idlist* must be consist of only a 1. The **init** process may be assigned to any class configured on the system, but the time-sharing class is almost always the appropriate choice. (Other choices may be highly undesirable; see the *File System Administration* for more information.)

The command

```
prioctl -e [-c class] [class-specific options] command [argument . . .]
```

executes the specified command with the class and scheduling parameters specified on the command line (*arguments* are the arguments to the command). If the **-c class** option is omitted the command is run in the user's current class.

OPTIONS

- l** Display a list of the classes currently configured in the system along with class-specific information about each class. The format of the class-specific information displayed is described under **USAGE**.
- d** Display the scheduling parameters associated with a set of processes.
- s** Set the scheduling parameters associated with a set of processes.
- e** Execute a specified command with the class and scheduling parameters associated with a set of processes.

- i idtype** This option together with the *idlist* arguments (if any), specify one or more processes to which the **prioctl** command is to apply. The interpretation of *idlist* depends on the value of *idtype*. The valid *idtype* arguments and corresponding interpretations of *idlist* are as follows:
- i pid** *idlist* is a list of process IDs. The **prioctl** command applies to the specified processes.
 - i ppid** *idlist* is a list of parent process IDs. The **prioctl** command applies to all processes whose parent process ID is in the list.
 - i pgid** *idlist* is a list of process group IDs. The **prioctl** command applies to all processes in the specified process groups.
 - i sid** *idlist* is a list of session IDs. The **prioctl** command applies to all processes in the specified sessions.
 - i class** *idlist* consists of a single class name (**RT** for real-time or **TS** for time-sharing or **IA** for inter-active). The **prioctl** command applies to all processes in the specified class.
 - i uid** *idlist* is a list of user IDs. The **prioctl** command applies to all processes with an effective user ID equal to an ID from the list.
 - i gid** *idlist* is a list of group IDs. The **prioctl** command applies to all processes with an effective group ID equal to an ID from the list.
 - i all** The **prioctl** command applies to all existing processes. No *idlist* should be specified (if one is it is ignored). The permission restrictions described below still apply.

If the **-i idtype** option is omitted when using the **-d** or **-s** options the default *idtype* of **pid** is assumed.

- c class** Specifies the *class* to be set. (The valid *class* arguments are **RT** for real-time or **TS** for time-sharing or **IA** for inter-active.) If the specified class is not already configured, it will automatically be configured.

The valid class-specific options for setting real-time parameters are:

- p rtpri** Set the real-time priority of the specified process(es) to *rtpri*.

- t tqntm** [-r *res*]

Set the time quantum of the specified process(es) to *tqntm*. You may optionally specify a resolution as explained below.

The valid class-specific options for setting time-sharing parameters are:

- m tsuprilim**

Set the user priority limit of the specified process(es) to *tsuprilim*.

- p tsupri** Set the user priority of the specified process(es) to *tsupri*.

The valid class-specific options for setting inter-active parameters are:

- m iamode** Mark the specified process(es) as currently interactive, or not. *iamode*.

USAGE
Real-Time Class

The real-time class provides a fixed priority preemptive scheduling policy for those processes requiring fast and deterministic response and absolute user/application control of scheduling priorities. If the real-time class is configured in the system it should have exclusive control of the highest range of scheduling priorities on the system. This ensures that a runnable real-time process is given CPU service before any process belonging to any other class.

The real-time class has a range of real-time priority (*rtpri*) values that may be assigned to processes within the class. Real-time priorities range from 0 to *x*, where the value of *x* is configurable and can be displayed for a specific installation that has already configured a real-time scheduler, by using the command

prioctl -l

The real-time scheduling policy is a fixed priority policy. The scheduling priority of a real-time process never changes except as the result of an explicit request by the user/application to change the *rtpri* value of the process.

For processes in the real-time class, the *rtpri* value is, for all practical purposes, equivalent to the scheduling priority of the process. The *rtpri* value completely determines the scheduling priority of a real-time process relative to other processes within its class. Numerically higher *rtpri* values represent higher priorities. Since the real-time class controls the highest range of scheduling priorities in the system it is guaranteed that the runnable real-time process with the highest *rtpri* value is always selected to run before any other process in the system.

In addition to providing control over priority, **prioctl** provides for control over the length of the time quantum allotted to processes in the real-time class. The time quantum value specifies the maximum amount of time a process may run assuming that it does not complete or enter a resource or event wait state (**sleep**). Note: If another process becomes runnable at a higher priority the currently running process may be preempted before receiving its full time quantum.

The command

prioctl -d [-i *idtype*] [*idlist*]

displays the real-time priority and time quantum (in millisecond resolution) for each real-time process in the set specified by *idtype* and *idlist*.

Any combination of the **-p** and **-t** options may be used with **prioctl -s** or **prioctl -e** for the real-time class. If an option is omitted and the process is currently real-time the associated parameter is unaffected. If an option is omitted when changing the class of a process to real-time from some other class, the associated parameter is set to a default value. The default value for *rtpri* is 0 and the default for time quantum is dependent on the value of *rtpri* and on the system configuration; see **rt_dptbl(4)**.

When using the `-t tqntm` option you may optionally specify a resolution using the `-r res` option. (If no resolution is specified, millisecond resolution is assumed.) If *res* is specified it must be a positive integer between 1 and 1,000,000,000 inclusive and the resolution used is the reciprocal of *res* in seconds. For example, specifying `-t 10 -r 100` would set the resolution to hundredths of a second and the resulting time quantum length would be 10/100 seconds (one tenth of a second). Although very fine (nanosecond) resolution may be specified, the time quantum length is rounded up by the system to the next integral multiple of the system clock's resolution. Requests for time quanta of zero or quanta greater than the (typically very large) implementation-specific maximum quantum result in an error.

In order to change the class of a process to real-time (from any other class) the user invoking **priocntl** must have super-user privilege. In order to change the *rtpri* value or time quantum of a real-time process the user invoking **priocntl** must either be super-user, or must currently be in the real-time class (shell running as a real-time process) with a real or effective user ID matching the real or effective user ID of the target process.

The real-time priority and time quantum are inherited across the `fork(2)` and `exec(2)` system calls.

Time-Sharing Class

The time-sharing scheduling policy provides for a fair and effective allocation of the CPU resource among processes with varying CPU consumption characteristics. The objectives of the time-sharing policy are to provide good response time to interactive processes and good throughput to CPU-bound jobs while providing a degree of user/application control over scheduling.

The time-sharing class has a range of time-sharing user priority (*tsupri*) values that may be assigned to processes within the class. User priorities range from `-x` to `+x`, where the value of *x* is configurable. The range for a specific installation can be displayed by using the command

```
priocntl -l
```

The purpose of the user priority is to provide some degree of user/application control over the scheduling of processes in the time-sharing class. Raising or lowering the *tsupri* value of a process in the time-sharing class raises or lowers the scheduling priority of the process. It is not guaranteed, however, that a time-sharing process with a higher *tsupri* value will run before one with a lower *tsupri* value. This is because the *tsupri* value is just one factor used to determine the scheduling priority of a time-sharing process. The system may dynamically adjust the internal scheduling priority of a time-sharing process based on other factors such as recent CPU usage.

In addition to the system-wide limits on user priority (displayed with **priocntl -l**), there is a per process user priority limit (*tsuprilim*), which specifies the maximum *tsupri* value that may be set for a given process.

The command

```
priocntl -d [-i idtype] [idlist]
```

displays the user priority and user priority limit for each time-sharing process in the set specified by *idtype* and *idlist*.

Any time-sharing process may lower its own *tsuprilim* (or that of another process with the same user ID). Only a time-sharing process with super-user privilege may raise a *tsuprilim*. When changing the class of a process to time-sharing from some other class, super-user privilege is required in order to set the initial *tsuprilim* to a value greater than zero.

Any time-sharing process may set its own *tsupri* (or that of another process with the same user ID) to any value less than or equal to the process's *tsuprilim*. Attempts to set the *tsupri* above the *tsuprilim* (and/or set the *tsuprilim* below the *tsupri*) result in the *tsupri* being set equal to the *tsuprilim*.

Any combination of the **-m** and **-p** options may be used with **prioctl -s** or **prioctl -e** for the time-sharing class. If an option is omitted and the process is currently time-sharing the associated parameter is normally unaffected. The exception is when the **-p** option is omitted and **-m** is used to set a *tsuprilim* below the current *tsupri*. In this case the *tsupri* is set equal to the *tsuprilim* which is being set. If an option is omitted when changing the class of a process to time-sharing from some other class, the associated parameter is set to a default value. The default value for *tsuprilim* is 0 and the default for *tsupri* is to set it equal to the *tsuprilim* value which is being set.

The time-sharing user priority and user priority limit are inherited across the **fork(2)** and **exec(2)** system calls.

Inter-Active Class

The inter-active scheduling policy provides for a fair and effective allocation of the CPU resource among processes with varying CPU consumption characteristics while providing good responsiveness for user interaction. The objectives of the inter-active policy are to provide good response time to interactive processes and good throughput to CPU-bound jobs. Only the super user has access to the inter-active class, the user has no control over scheduling policies.

EXAMPLES

Real-Time Class examples follow:

```
example% prioctl -s -c RT -t 1 -r 10 -i idtype idlist
```

The above example sets the class of any non-real-time processes selected by *idtype* and *idlist* to real-time and sets their real-time priority to the default value of 0. The real-time priorities of any processes currently in the real-time class are unaffected. The time quantum of all of the specified processes are set to 1/10 seconds.

```
example% prioctl -e -c RT -p 15 -t 20 command
```

This example executes *command* in the real-time class with a real-time priority of 15 and a time quantum of 20 milliseconds.

Time-Sharing Class examples follow:

```
example% prioctl -s -c TS -i idtype idlist
```

The above example sets the class of any non-time-sharing processes selected by *idtype* and *idlist* to time-sharing and sets both their user priority limit and user priority to 0. Processes already in the time-sharing class are unaffected.

This example executes *command* with the arguments *arguments* in the time-sharing class with a user priority limit of 0 and a user priority of -15.

```
example% priocntl -e -c TS -m 0 -p -15 command [arguments]
```

SEE ALSO

nice(1), ps(1), exec(2), fork(2), priocntl(2), rt_dptbl(4)

DIAGNOSTICS

priocntl prints the following error messages:

Process(es) not found

None of the specified processes exists.

Specified processes from different classes

The **-s** option is being used to set parameters, the **-c class** option is not present, and processes from more than one class are specified.

Invalid option or argument

An unrecognized or invalid option or option argument is used.

NAME	prof – display profile data
SYNOPSIS	prof [-a c n t] [-o x] [-g I] [-C] [-h] [-m mdata] [-s] [-V prog] [-z]
DESCRIPTION	The prof command interprets a profile file produced by the monitor function. The symbol table in the object file <i>prog</i> (a.out by default) is read and correlated with a profile file (mon.out by default). For each external text symbol the percentage of time spent executing between the address of that symbol and the address of the next is printed, together with the number of times that function was called and the average number of milliseconds per call.
OPTIONS	<p>The mutually exclusive options -a, -c, -n, and -t determine the type of sorting of the output lines:</p> <ul style="list-style-type: none"> -a Sort by increasing symbol address. -c Sort by decreasing number of calls. -n Sort lexically by symbol name. -t Sort by decreasing percentage of total time (default). <p>The mutually exclusive options -o and -x specify the printing of the address of each symbol monitored:</p> <ul style="list-style-type: none"> -o Print each symbol address (in octal) along with the symbol name. -x Print each symbol address (in hexadecimal) along with the symbol name. <p>The mutually exclusive options -g and -I control the type of symbols to be reported. The -I option must be used with care; it applies the time spent in a static function to the preceding (in memory) global function, instead of giving the static function a separate entry in the report. If all static functions are properly located (see example below), this feature can be very useful. If not, the resulting report may be misleading.</p> <p>Assume that A and B are global functions and only A calls static function S. If S is located immediately after A in the source code (that is, if S is properly located), then, with the -I option, the amount of time spent in A can easily be determined, including the time spent in S. If, however, both A and B call S, then, if the -I option is used, the report will be misleading; the time spent during B's call to S will be attributed to A, making it appear as if more time had been spent in A than really had. In this case, function S cannot be properly located.</p> <ul style="list-style-type: none"> -g Include static (non-global) functions. -I Do not include static (non-global) functions (default). <p>The following options may be used in any combination:</p> <ul style="list-style-type: none"> -C Demangle C++ symbol names before printing them out. -h Suppress the heading normally printed on the report. This is useful if the report is to be processed further.

-m mdata

Use file *mdata* instead of **mon.out** as the input profile file.

-s Print a summary of several of the monitoring parameters and statistics on the standard error output.

-V Print **prof** version information on the standard error output.

-z Include all symbols in the profile range, even if associated with zero number of calls and zero time.

A program creates a profile file if it has been link edited with the **-p** option of **cc(1B)**. This option to the **cc(1B)** command arranges for calls to **monitor** at the beginning and end of execution. It is the call to **monitor** at the end of execution that causes the system to write a profile file. The number of calls to a function is tallied if the **-p** option was used when the file containing the function was compiled.

A single function may be split into subfunctions for profiling by means of the **MARK** macro (see **prof(5)**).

ENVIRONMENT

PROFDIR The name of the file created by a profiled program is controlled by the environment variable **PROFDIR**. If **PROFDIR** is not set, **mon.out** is produced in the directory current when the program terminates. If **PROFDIR=string**, *string/pid.progname* is produced, where *progname* consists of **argv[0]** with any path prefix removed, and *pid* is the process ID of the program. If **PROFDIR** is set, but null, no profiling output is produced.

FILES

mon.out	default profile file
a.out	default namelist (object) file

SEE ALSO

cc(1B), **exit(2)**, **profil(2)**, **malloc(3C)**, **monitor(3C)**, **malloc(3X)**, **prof(5)**

The **lprof** section in *Programming Utilities Guide*

NOTES

The times reported in successive identical runs may show variances because of varying cache-hit ratios that result from sharing the cache with other processes. Even if a program seems to be the only one using the machine, hidden background or asynchronous processes may blur the data. In rare cases, the clock ticks initiating recording of the program counter may “beat” with loops in a program, grossly distorting measurements. Call counts are always recorded precisely, however.

Only programs that call **exit** or return from **main** are guaranteed to produce a profile file, unless a final call to **monitor** is explicitly coded.

The times for static functions are attributed to the preceding external text symbol if the **-g** option is not used. However, the call counts for the preceding function are still correct; that is, the static function call counts are not added to the call counts of the external function.

If more than one of the options **-t**, **-c**, **-a**, and **-n** is specified, the last option specified is used and the user is warned.

Profiling may be used with dynamically linked executables, but care must be applied. Currently, shared objects cannot be profiled with **prof**. Thus, when a profiled, dynamically linked program is executed, only the “main” portion of the image is sampled. This means that all time spent outside of the “main” object, that is, time spent in a shared object, will not be included in the profile summary; the total time reported for the program may be less than the total time used by the program.

Because the time spent in a shared object cannot be accounted for, the use of shared objects should be minimized whenever a program is profiled with **prof**. If desired, the program should be linked to the profiled version of a library (or to the standard archive version if no profiling version is available), instead of the shared object to get profile information on the functions of a library. Versions of profiled libraries may be supplied with the system in the **/usr/lib/libp** directory. Refer to compiler driver documentation on profiling.

Consider an extreme case. A profiled program dynamically linked with the shared C library spends 100 units of time in some **libc** routine, say, **malloc()**. Suppose **malloc()** is called only from routine **B** and **B** consumes only 1 unit of time. Suppose further that routine **A** consumes 10 units of time, more than any other routine in the “main” (profiled) portion of the image. In this case, **prof** will conclude that most of the time is being spent in **A** and almost no time is being spent in **B**. From this it will be almost impossible to tell that the greatest improvement can be made by looking at routine **B** and not routine **A**. The value of the profiler in this case is severely degraded; the solution is to use archives as much as possible for profiling.

NAME	ps – report process status
SYNOPSIS	ps [-acdefjl] [-g <i>grplist</i>] [-p <i>proclist</i>] [-s <i>sidlist</i>] [-t <i>term</i>] [-u <i>uidlist</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>ps prints information about active processes. Without <i>options</i>, ps prints information about processes associated with the controlling terminal. The output contains only the process ID, terminal identifier, cumulative execution time, and the command name. Otherwise, the information that is displayed is controlled by the <i>options</i>.</p> <p>Some options accept lists as arguments. Items in a list can be either separated by commas or else enclosed in double quotes and separated by commas or spaces. Values for <i>proclist</i> and <i>grplist</i> must be numeric.</p>
OPTIONS	<p>-a Print information about all processes most frequently requested: all those except process group leaders and processes not associated with a terminal.</p> <p>-c Print information in a format that reflects scheduler properties as described in prioctl(1). The -c option affects the output of the -f and -l options, as described below.</p> <p>-d Print information about all processes except session leaders.</p> <p>-e Print information about every process now running.</p> <p>-f Generate a full listing. (See below for significance of columns in a full listing.)</p> <p>-j Print session ID and process group ID.</p> <p>-l Generate a long listing. (See below.)</p> <p>-g <i>grplist</i> List only process data whose group leader's ID number(s) appears in <i>grplist</i>. (A group leader is a process whose process ID number is identical to its process group ID number.)</p> <p>-p <i>proclist</i> List only process data whose process ID numbers are given in <i>proclist</i>.</p> <p>-s <i>sidlist</i> List information on all session leaders whose IDs appear in <i>sidlist</i>.</p> <p>-t <i>term</i> List only process data associated with <i>term</i>. Terminal identifiers are specified as a device file name, and an identifier. For example, term/a, or pts/0.</p> <p>-u <i>uidlist</i> List only process data whose user ID number or login name is given in <i>uidlist</i>. In the listing, the numerical user ID will be printed unless you give the -f option, which prints the login name.</p>

**DISPLAY
FORMATS**

Under the **-f** option, **ps** tries to determine the command name and arguments given when the process was created by examining the user block. Failing this, the command name is printed, as it would have appeared without the **-f** option, in square brackets.

The column headings and the meaning of the columns in a **ps** listing are given below; the letters **f** and **l** indicate the option (full or long, respectively) that causes the corresponding heading to appear; **all** means that the heading always appears. Note: These two options determine only what information is provided for a process; they do not determine which processes will be listed.

- F** (l) Flags (hexadecimal and additive) associated with the process. These flags are available for historical purposes; no meaning should be currently ascribed to them.
- S** (l) The state of the process:
- O** Process is running on a processor.
 - S** Sleeping: process is waiting for an event to complete.
 - R** Runnable: process is on run queue.
 - I** Idle: process is being created.
 - Z** Zombie state: process terminated and parent not waiting.
 - T** Traced: process stopped by a signal because parent is tracing it.
 - X** SXBRK state: process is waiting for more primary memory.
- UID** (f,l) The user ID number of the process owner (the login name is printed under the **-f** option).
- PID** (all) The process ID of the process (this datum is necessary in order to kill a process).
- PPID** (f,l) The process ID of the parent process.
- C** (f,l) Processor utilization for scheduling. Not printed when the **-c** option is used.
- CLS** (f,l) Scheduling class. Printed only when the **-c** option is used.
- PRI** (l) The priority of the process. Without the **-c** option, higher numbers mean lower priority. With the **-c** option, higher numbers mean higher priority.
- NI** (l) Nice value, used in priority computation. Not printed when the **-c** option is used. Only processes in the time-sharing class have a nice value.
- ADDR** (l) The memory address of the process.
- SZ** (l) The size (in pages or clicks) of the swappable process's image in main memory.

- WCHAN** (l) The address of an event for which the process is sleeping, or in SXBRK state, (if blank, the process is running).
- STIME** (f) The starting time of the process, given in hours, minutes, and seconds. (A process begun more than twenty-four hours before the **ps** inquiry is executed is given in months and days.)
- TTY** (all) The controlling terminal for the process (the message, **?**, is printed when there is no controlling terminal).
- TIME** (all) The cumulative execution time for the process.
- COMMAND**(all) The command name (the full command name and its arguments are printed under the **-f** option).

A process that has exited and has a parent, but has not yet been waited for by the parent, is marked **<defunct>**.

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

LC_TIME

Determines how **ps** handles date and time formats. In the "C" locale, date and time handling follows the U.S. rules.

FILES

/dev
/dev/sxt/*
/dev/tty*
/dev/xt/* terminal ("tty") names searcher files
/dev/kmem kernel virtual memory
/dev/mem memory
/dev/swap the default swap device
/etc/passwd UID information supplier

/tmp/ps_data internal data structure

SEE ALSO

kill(1), nice(1), priocntl(1), getty(1M), environ(5)

NOTES

Things can change while **ps** is running; the snap-shot it gives is true only for a split-second, and it may not be accurate by the time you see it. Some data printed for defunct processes is irrelevant.

If no *termlist*, *proclist*, *uidlist*, or *grplist* is specified, **ps** checks **stdin**, **stdout**, and **stderr** in that order, looking for the controlling terminal and will attempt to report on processes associated with the controlling terminal. In this situation, if **stdin**, **stdout**, and **stderr** are all redirected, **ps** will not find a controlling terminal, so there will be no report.

On a heavily loaded system, **ps** may report an **lseek** error and exit. **ps** may seek to an invalid user area address: having obtained the address of a process' user area, **ps** may not be able to seek to that address before the process exits and the address becomes invalid.

ps -ef may not report the actual start of a tty login session, but rather an earlier time, when a getty was last respawned on the tty line.

NAME	ps – display the status of current processes
SYNOPSIS	/usr/ucb/ps [-acglnrSuUvwx] [-t <i>term</i>] [<i>num</i>]
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>The ps command displays information about processes. Normally, only those processes that are running with your effective user ID and are attached to a controlling terminal (see termio(7)) are shown. Additional categories of processes can be added to the display using various options. In particular, the -a option allows you to include processes that are not owned by you (that do not have your user ID), and the -x option allows you to include processes without control terminals. When you specify both -a and -x, you get processes owned by anyone, with or without a control terminal. The -r option restricts the list of processes printed to running and runnable processes.</p> <p>ps displays the process ID, under PID; the control terminal (if any), under TT; the cpu time used by the process so far, including both user and system time, under TIME; the state of the process, under S; and finally, an indication of the COMMAND that is running.</p> <p>The state is given by a single letter from the following:</p> <ul style="list-style-type: none"> O Process is running on a processor. S Sleeping. Process is waiting for an event to complete. R Runnable. Process is on run queue. I Idle. Process is being created. Z Zombie state. Process terminated and parent not waiting. T Traced. Process stopped by a signal because parent is tracing it. X SXBRK state. Process is waiting for more primary memory.
OPTIONS	<p>The following options must all be combined to form the first argument:</p> <ul style="list-style-type: none"> -a Include information about processes owned by others. -c Display the command name, as stored internally in the system for purposes of accounting, rather than the command arguments, which are kept in the process' address space. This is more reliable, if less informative, since the process is free to destroy the latter information. -g Display all processes. Without this option, ps only prints interesting processes. Processes are deemed to be uninteresting if they are process group leaders. This normally eliminates top-level command interpreters and processes waiting for users to login on free terminals. -l Display a long listing, with fields F, PPID, CP, PRI, NI, SZ, RSS and WCHAN as described below. -n Produce numerical output for some fields. In a user listing, the USER field is replaced by a UID field. -r Restrict output to running and runnable processes.

- S Display accumulated CPU time used by this process and all of its reaped children.
- u Display user-oriented output. This includes fields **USER** , **SZ** , **RSS** and **START** as described below.
- U Update a private database where **ps** keeps system information.
- v Display a version of the output containing virtual memory. This includes fields **SIZE** and **RSS** , described below.
- w Use a wide output format (132 columns rather than 80); if repeated, that is, **-ww** , use arbitrarily wide output. This information is used to decide how much of long commands to print.
- x Include processes with no controlling terminal.
- t *term* List only process data associated with the terminal, *term*. Terminal identifiers may be specified in one of two forms: the device's file name (for example, **tty04** or **term/14**) or, if the device's file name starts with **tty**, just the digit identifier (for example, **04**).
- num* A process number may be given, in which case the output is restricted to that process. This option must be supplied last.

DISPLAY FORMATS

Fields that are not common to all output formats:

- USER** Name of the owner of the process.
- NI** Process scheduling increment (see **getpriority(3B)** and **nice(3B)**).
- SIZE**
- SZ** The combined size of the data and stack segments (in kilobyte units)
- RSS** Real memory (resident set) size of the process (in kilobyte units).
- UID** Numerical user-ID of process owner.
- PPID** Numerical ID of parent of process.
- CP** Short-term CPU utilization factor (used in scheduling).
- PRI** The priority of the process (higher numbers mean lower priority).
- START** The starting time of the process, given in hours, minutes, and seconds. A process begun more than 24 hours before the **ps** inquiry is executed is given in months and days.
- WCHAN** The address of an event for which the process is sleeping, or in **SXBRK** state (if blank, the process is running).
- F** Flags (hexadecimal and additive) associated with the process:
 - 00** Process has terminated. Process table now available.
 - 01** A system process, always in primary memory.
 - 02** Parent is tracing process.
 - 04** Tracing parent's signal has stopped process. Parent is waiting, see **ptrace(2)**.
 - 08** Process is currently in primary memory.

10 Process currently in primary memory, locked until an event is completed.

A process that has exited and has a parent, but has not yet been waited for by the parent is marked < **defunct** >; otherwise, **ps** tries to determine the command name and arguments given when the process was created by examining the user block.

FILES

/dev	
/dev/kmem	kernel virtual memory
/dev/mem	memory
/dev/swap	default swap device
/dev/sxt/*	
/dev/tty*	
/dev/xt/*	terminal (tty) names searcher files
/etc/passwd	UID information supplier
/etc/ps_data	internal data structure

SEE ALSO

kill(1), **whodo(1M)**, **lseek(2)**, **getpriority(3B)**, **nice(3B)**

NOTES

Things can change while **ps** is running; the picture it gives is only a close approximation to the current state. Some data printed for defunct processes is irrelevant.

If no *term* or *num* is specified, **ps** checks the standard input, the standard output, and the standard error in that order, looking for the controlling terminal and will attempt to report on processes associated with the controlling terminal. In this situation, if the standard input, the standard output, and the standard error are all redirected, **ps** will not find a controlling terminal, so there will be no report.

On a heavily loaded system, **ps** may report an **lseek(2)** error and exit. **ps** may seek to an invalid user area address, having obtained the address of process' user area, **ps** may not be able to seek to that address before the process exits and the address becomes invalid.

NAME	pwd – working directory name
SYNOPSIS	pwd
AVAILABILITY	SUNWcsu
DESCRIPTION	pwd prints the path name of the working (current) directory.
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of pwd for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how pwd behaves.</p> <p>LC_CTYPE Determines how pwd handles characters. When LC_CTYPE is set to a valid value, pwd can display and handle text and filenames containing valid characters for that locale. pwd can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. pwd can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p> <p>LC_MESSAGES Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).</p>
SEE ALSO	cd(1) , shell_builtins(1) , environ(5)
DIAGNOSTICS	“ Cannot open .. ” and “ Read error in .. ” indicate possible file system trouble and should be referred to a UNIX system administrator.
NOTES	If you move the current directory or one above it, pwd may not give the correct response. Use the cd(1) command with a full path name to correct this situation.

NAME	rcp – remote file copy
SYNOPSIS	rcp [-p] <i>filename1 filename2</i> rcp [-pr] <i>filename...directory</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The rcp command copies files between machines. Each <i>filename</i> or <i>directory</i> argument is either a remote file name of the form:</p> <p style="padding-left: 40px;"><i>hostname:path</i></p> <p>or a local file name (containing no : characters, or a / before any : characters).</p> <p>If a <i>filename</i> is not a full path name, it is interpreted relative to your home directory on <i>hostname</i>. A <i>path</i> on a remote host may be quoted (using \, " , or ') so that the metacharacters are interpreted remotely.</p> <p>rcp does not prompt for passwords; your current local user name must exist on <i>hostname</i> and allow remote command execution by rsh(1).</p> <p>rcp handles third party copies, where neither source nor target files are on the current machine. Hostnames may also take the form</p> <p style="padding-left: 40px;"><i>username@hostname:filename</i></p> <p>to use <i>username</i> rather than your current local user name as the user name on the remote host. rcp also supports Internet domain addressing of the remote host, so that:</p> <p style="padding-left: 40px;"><i>username@host.domain:filename</i></p> <p>specifies the username to be used, the hostname, and the domain in which that host resides. Filenames that are not full path names will be interpreted relative to the home directory of the user named <i>username</i>, on the remote host.</p>
OPTIONS	<p>-p Attempt to give each copy the same modification times, access times, and modes as the original file.</p> <p>-r Copy each subtree rooted at <i>filename</i>; in this case the destination must be a directory.</p>
FILES	\$HOME/.profile
SEE ALSO	cpio (1), ftp (1), rlogin (1), rsh (1), tar (1), hosts.equiv (4)
NOTES	<p>rcp is meant to copy between different hosts; attempting to rcp a file onto itself, as with:</p> <p style="padding-left: 40px;">rcp tmp/file myhost:/tmp/file</p> <p>results in a severely corrupted file.</p> <p>rcp may not correctly fail when the target of a copy is a file instead of a directory.</p>

rcp can become confused by output generated by commands in a **\$HOME/.profile** on the remote host.

rcp requires that the source host have permission to execute commands on the remote host when doing third-party copies.

rcp does not properly handle symbolic links. Use **tar** (see **tar(1)**) or **cpio** (see **cpio(1)**) piped to **rsh** to obtain remote copies of directories containing symbolic links or named pipes.

If you forget to quote metacharacters intended for the remote host you get an incomprehensible error message.

NAME	rdist – remote file distribution program
SYNOPSIS	<pre>rdist [-b] [-D] [-h] [-i] [-n] [-q] [-R] [-v] [-w] [-y] [-d macro = value] [-f distfile] [-m host] ... [package ...] rdist [-b] [-D] [-h] [-i] [-n] [-q] [-R] [-v] [-w] [-y] -c pathname ... [login@] hostname [:destpath]</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>rdist maintains copies of files on multiple hosts. It preserves the owner, group, mode, and modification time of the master copies, and can update programs that are executing. Normally, a copy on a remote host is updated if its size or modification time differs from the original on the local host. rdist reads the indicated <i>distfile</i> for instructions on updating files and/or directories. If <i>distfile</i> is '-', the standard input is used. If no -f option is present, rdist first looks in its working directory for distfile, and then for Distfile, for instructions.</p> <p>rdist updates each <i>package</i> specified on the command line; if none are given, all packages are updated according to their entries in the <i>distfile</i>.</p> <p>In order to be able to use rdist across machines, each host machine must have a <i>/etc/host.equiv</i> file, or the user must have an entry in the <i>.rhosts</i> file in the home directory. See hosts.equiv(4) for more information.</p>
OPTIONS	<p>-b Binary comparison. Perform a binary comparison and update files if they differ, rather than merely comparing dates and sizes.</p> <p>-D Enable debugging.</p> <p>-h Follow symbolic links. Copy the file that the link points to rather than the link itself.</p> <p>-i Ignore unresolved links. rdist will normally try to maintain the link structure of files being transferred and warn the user if all the links cannot be found.</p> <p>-n Print the commands without executing them. This option is useful for debugging a <i>distfile</i>.</p> <p>-q Quiet mode. Do not display the files being updated on the standard output.</p> <p>-R Remove extraneous files. If a directory is being updated, remove files on the remote host that do not correspond to those in the master (local) directory. This is useful for maintaining truly identical copies of directories.</p> <p>-v Verify that the files are up to date on all the hosts. Any files that are out of date are displayed, but no files are updated, nor is any mail sent.</p> <p>-w Whole mode. The whole file name is appended to the destination directory name. Normally, only the last component of a name is used when</p>

- renaming files. This preserves the directory structure of the files being copied, instead of flattening the directory structure. For instance, renaming a list of files such as **dir1/dir2** to **dir3** would create files **dir3/dir1** and **dir3/dir2** instead of **dir3** and **dir3**. When the **-w** option is used with a filename that begins with **~**, everything except the home directory is appended to the destination name.
- y** Younger mode. Do not update remote copies that are younger than the master copy, but issue a warning message instead.
- d macro=value** Define *macro* to have *value*. This option is used to define or override macro definitions in the distfile. *value* can be the empty string, one name, or a list of names surrounded by parentheses and separated by white space.
- c pathname ...** [*login@*]hostname[:*destpath*]
Update each *pathname* on the named host. (Relative filenames are taken as relative to your home directory.) If the '*login@*' prefix is given, the update is performed with the user ID of *login*. If the '*:destpath*' is given, the remote file is installed as that pathname.
- f distfile** Use the description file *distfile*. A '-' as the *distfile* argument denotes the standard input.
- m host** Limit which machines are to be updated. Multiple **-m** arguments can be given to limit updates to a subset of the hosts listed in the distfile.

USAGE

Packages

A typical package begins with a label composed of the package name followed by a colon:

package:

This label allows you to group any number of file-to-host and file-to-timestamp mappings into a single distribution package. If no package label appears in the distfile, the default package includes all mappings in the file.

A file-to-host mapping specifies a list of files or directories to distribute, their destination host(s), and any **rdist** primitives to use in performing the update. A mapping of this sort takes the form:

(*pathname ...*) → (*hostname ...*) *primitive* ; [*primitive* ;] ...

In this case, each *pathname* is the full pathname of a local file or directory to distribute; each *hostname* is the name of a remote host on which those files are to be copied, and *primitive* is one of the **rdist** primitive listed under *Primitives*, below. A *hostname* can also take the form:

login@hostname

in which case the update is performed as the user named *login*.

A file-to-timestamp mapping is used to monitor which local files are updated with respect to a local "timestamp" file. This mapping takes the form:

	<p>(<i>filename . . .</i>) :: <i>timestamp-file primitive</i> ; [<i>primitive</i> ;] . . .</p> <p>In this case, <i>timestamp-file</i> is the name of a file, the modification time of which is compared with each named file on the local host. If a file is newer than <i>time-stamp-file</i>, rdist displays a message to that effect.</p>
White Space Characters	<p>NEWLINE, TAB, and SPACE characters are all treated as white space; a mapping continues across input lines until the start of the next mapping: either a single <i>filename</i> followed by a '→' or the opening parenthesis of a filename list.</p>
Comments	<p>Comments begin with # and end with a NEWLINE.</p>
Macros	<p>rdist has a limited macro facility. Macros are only expanded in filename or hostname lists, and in the argument lists of certain primitives. Macros cannot be used to stand for primitives or their options, or the '→' or '::' symbols.</p> <p>A macro definition is a line of the form:</p> <p style="padding-left: 40px;"><i>macro = value</i></p> <p>A macro reference is a string of the form:</p> <p style="padding-left: 40px;">\${<i>macro</i>}</p> <p>although (as with make(1S)) the braces can be omitted if the macro name consists of just one character.</p>
Metacharacters	<p>The shell meta-characters: [,], {, }, *, and ? are recognized and expanded (on the local host only) just as they are with csh(1). Metacharacters can be escaped by prepending a backslash.</p> <p>The ~ character is also expanded in the same way as with csh, however, it is expanded separately on the local and destination hosts.</p>
Filenames	<p>File names that do not begin with '/' or '~' are taken to be relative to user's home directory on each destination host; they are <i>not</i> relative to the current working directory. Multiple file names must be enclosed within parentheses.</p>
Primitives	<p>The following primitives can be used to specify actions rdist is to take when updating remote copies of each file.</p> <p>install [-b] [-h] [-i] [-R] [-v] [-w] [-y] [<i>newname</i>]</p> <p>Copy out-of-date files and directories (recursively). If no install primitive appears in the package entry, or if no <i>newname</i> option is given, the name of the local file is given to the remote host's copy. If absent from the remote host, parent directories in a filename's path are created. To help prevent disasters, a non-empty directory on a target host is not replaced with a regular file or a symbolic link by rdist. However, when using the -R option, a non-empty directory is removed if the corresponding filename is completely absent on the master host.</p>

The options for **install** have the same semantics as their command line counterparts, but are limited in scope to a particular map. The login name used on the destination host is the same as the local host unless the destination name is of the format *login@host*. In that case, the update is performed under the username *login*.

notify *address* ...

Send mail to the indicated TCP/IP *address* of the form:

user@host

that lists the files updated and any errors that may have occurred. If an address does not contain a '@host' suffix, **rdist** uses the name of the destination host to complete the address.

except *filename* ...

Omit from updates the files named as arguments.

except_pat *pattern* ...

Omit from updates the filenames that match each regular-expression *pattern* (see **ed**(1) for more information on regular expressions). Note that '\ ' and '\$' characters must be escaped in the distfile. Shell variables can also be used within a pattern, however shell filename expansion is not supported.

special [*filename*] ... "*command-line*"

Specify a Bourne shell, **sh**(1) command line to execute on the remote host after each named file is updated. If no *filename* argument is present, the *command-line* is performed for every updated file, with the shell variable **FILE** set to the file's name on the local host. The quotation marks allow *command-line* to span input lines in the distfile; multiple shell commands must be separated by semicolons (;).

The default working directory for the shell executing each *command-line* is the user's home directory on the remote host.

EXAMPLES

The following sample distfile instructs **rdist** to maintain identical copies of a shared library, a shared-library initialized data file, several include files, and a directory, on hosts named **hermes** and **magus**. On **magus**, commands are executed as root. **rdist** notifies **merlin@druid** whenever it discovers that a local file has changed relative to a timestamp file.

```
HOSTS = ( hermes root@magus )
```

```
FILES = ( /usr/local/lib/libcant.so.1.1
          /usrlocal/lib/libcant.sa.1.1 /usr/local/include/{*.h}
          /usr/local/bin )
```

```
(${FILES}) → (${HOSTS})
install -R ;
```

```
${FILES} :: /usr/local/lib/timestamp
notify merlin@druid ;
```

FILES	~/rhosts user's trusted hosts and users /etc/host.equiv system trusted hosts and users /tmp/rdist* temporary file for update lists
SEE ALSO	cs(1), ed(1), make(1S), sh(1), stat(2)
DIAGNOSTICS	A complaint about mismatch of rdist version numbers may really stem from some problem with starting your shell, for example, you are in too many groups.
WARNINGS	root does not have its accustomed access privileges on NFS mounted file systems. Using rdist to copy to such a file system may fail, or the copies may be owned by user "nobody."
BUGS	Source files must reside or be mounted on the local host. There is no easy way to have a special command executed only once after all files in a directory have been updated. Variable expansion only works for name lists; there should be a general macro facility. rdist aborts on files that have a negative modification time (before Jan 1, 1970). There should be a "force" option to allow replacement of non-empty directories by regular files or symlinks. A means of updating file modes and owners of otherwise identical files is also needed.

NAME	read – shell built-in function to receive from standard input (keyboard)
SYNOPSIS	
sh	read <i>name</i> ...
cs h	set <i>variable</i> = \$<
ksh	read [-prsu [<i>n</i>]] [<i>name?prompt</i>] [<i>name</i> ...]
DESCRIPTION	
sh	One line is read from the standard input and, using the internal field separator, IFS (normally space or tab), to delimit word boundaries, the first word is assigned to the first <i>name</i> , the second word to the second <i>name</i> , etc., with leftover words assigned to the last <i>name</i> . Lines can be continued using \newline . Characters other than newline can be quoted by preceding them with a backslash. These backslashes are removed before words are assigned to <i>names</i> , and no interpretation is done on the character that follows the backslash. The return code is 0 , unless an EOF is encountered.
cs h	The notation set <i>variable</i> = \$< loads one line of standard input as the value for <i>variable</i> . (See cs h (1)).
ksh	The shell input mechanism. One line is read and is broken up into fields using the characters in IFS as separators. The escape character, (\), is used to remove any special meaning for the next character and for line continuation. In raw mode, -r , the \ character is not treated specially. The first field is assigned to the first <i>name</i> , the second field to the second <i>name</i> , etc., with leftover fields assigned to the last <i>name</i> . The -p option causes the input line to be taken from the input pipe of a process spawned by the shell using & . If the -s flag is present, the input will be saved as a command in the history file. The flag -u can be used to specify a one digit file descriptor unit <i>n</i> to read from. The file descriptor can be opened with the exec special command. The default value of <i>n</i> is 0. If <i>name</i> is omitted then REPLY is used as the default <i>name</i> . The exit status is 0 unless the input file is not open for reading or an end-of-file is encountered. An end-of-file with the -p option causes cleanup for this process so that another can be spawned. If the first argument contains a ? , the remainder of this word is used as a <i>prompt</i> on standard error when the shell is interactive. The exit status is 0 unless an end-of-file is encountered.
SEE ALSO	cs h (1), ksh (1), sh (1), set (1)

NAME	readfile, longline – reads file, gets longest line
SYNOPSIS	readfile <i>filename</i> longline [<i>filename</i>]
DESCRIPTION	The readfile function reads <i>filename</i> and copies it to <i>stdout</i> . No translation of NEWLINE is done. It keeps track of the longest line it reads and if there is a subsequent call to longline , the length of that line, including the NEWLINE character, is returned. The longline function returns the length, including the NEWLINE character, of the longest line in <i>filename</i> . If <i>filename</i> is not specified, it uses the file named in the last call to readfile .
EXAMPLES	Here is a typical use of readfile and longline in a text frame definition file: <pre> . . . text="`readfile myfile`" columns=`longline` . . . </pre>
SEE ALSO	cat (1)
DIAGNOSTICS	If <i>filename</i> does not exist, readfile will return FALSE (that is, the expression will have an error return). longline returns 0 if a readfile has not previously been issued.
NOTES	More than one descriptor can call readfile in the same frame definition file. In text frames, if one of those calls is made from the text descriptor, then a subsequent use of longline will always get the longest line of the file read by the readfile associated with the text descriptor, even if it was not the most recent use of readfile .

NAME	readonly – shell built-in function to protect the value of the given variable from reassignment
SYNOPSIS	
sh	readonly [<i>name</i> ...]
ksh	†† readonly [<i>name</i> [= <i>value</i>]] ...
DESCRIPTION	
sh	The given <i>names</i> are marked <i>readonly</i> and the values of these <i>names</i> may not be changed by subsequent assignment. If no arguments are given, a list of all <i>readonly</i> names is printed.
ksh	The given <i>names</i> are marked readonly and these names cannot be changed by subsequent assignment. On this man page, ksh (1) commands that are preceded by one or two † (daggers) are treated specially in the following ways: <ol style="list-style-type: none">1. Variable assignment lists preceding the command remain in effect when the command completes.2. I/O redirections are processed after variable assignments.3. Errors cause a script that contains them to abort.4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.
SEE ALSO	ksh (1), sh (1), typeset (1)

NAME	refer – expand and insert references from a bibliographic database
SYNOPSIS	refer [-ben] [-ar] [-cstring] [-kx] [-lm,n] [-p filename] [-skeys] <i>filename</i> . .
AVAILABILITY	SUNWdoc
DESCRIPTION	<p>refer is a preprocessor for nroff(1), or troff(1), that finds and formats references. The input files (standard input by default) are copied to the standard output, except for lines between .[and .] command lines. Such lines are assumed to contain keywords as for lookbib(1), and are replaced by information from a bibliographic data base. The user can avoid the search, override fields from it, or add new fields. The reference data, from whatever source, is assigned to a set of troff strings. Macro packages such as ms(5) print the finished reference text from these strings. A flag is placed in the text at the point of reference. By default, the references are indicated by numbers.</p> <p>When refer is used with eqn(1), neqn, or tbl(1), refer should be used first in the sequence, to minimize the volume of data passed through pipes.</p>
OPTIONS	<p>-b Bare mode — do not put any flags in text (neither numbers or labels).</p> <p>-e Accumulate references instead of leaving the references where encountered, until a sequence of the form:</p> <pre style="margin-left: 40px;">.[\$LISTS .]</pre> <p>is encountered, and then write out all references collected so far. Collapse references to the same source.</p> <p>-n Do not search the default file.</p> <p>-ar Reverse the first <i>r</i> author names (Jones, J. A. instead of J. A. Jones). If <i>r</i> is omitted, all author names are reversed.</p> <p>-cstring Capitalize (with SMALL CAPS) the fields whose key-letters are in <i>string</i>.</p> <p>-kx Instead of numbering references, use labels as specified in a reference data line beginning with the characters <i>%x</i>; By default, <i>x</i> is L.</p> <p>-lm,n Instead of numbering references, use labels from the senior author's last name and the year of publication. Only the first <i>m</i> letters of the last name and the last <i>n</i> digits of the date are used. If either of <i>m</i> or <i>n</i> is omitted, the entire name or date, respectively, is used.</p> <p>-p filename Take the next argument as a file of references to be searched. The default file is searched last.</p>

-skeys Sort references by fields whose key-letters are in the *keys* string, and permute reference numbers in the text accordingly. Using this option implies the **-e** option. The key-letters in *keys* may be followed by a number indicating how many such fields are used, with a + sign taken as a very large number. The default is **AD**, which sorts on the senior author and date. To sort on all authors and then the date, for instance, use the options '**-sA+T**'.

FILES **/usr/lib/refer** directory of programs
/usr/lib/refer/papers directory of default publication lists and indexes

SEE ALSO **addbib(1), eqn(1), indxbib(1), lookbib(1), nroff(1), roffb(1), sortbib(1), tbl(1), troff(1)**

NAME	regcmp – regular expression compile
SYNOPSIS	regcmp [-] <i>filename</i> . . .
DESCRIPTION	The regcmp command performs a function similar to regcmp and, in most cases, precludes the need for calling regcmp from C programs. Bypassing regcmp saves on both execution time and program size. The command regcmp compiles the regular expressions in <i>filename</i> and places the output in <i>filename.i</i> .
OPTIONS	– If the – option is used, the output is placed in <i>filename.c</i> . The format of entries in <i>filename</i> is a name (C variable) followed by one or more blanks followed by one or more regular expressions enclosed in double quotes. The output of regcmp is C source code. Compiled regular expressions are represented as extern char vectors. <i>filename.i</i> files may thus be #included in C programs, or <i>filename.c</i> files may be compiled and later loaded. In the C program that uses the regcmp output, regex(abc,line) applies the regular expression named abc to line . Diagnostics are self-explanatory.
EXAMPLES	<pre>name "[A-Za-z][A-Za-z0-9_]*\$0" telno "\({0,1}([2-9][01][1-9])\$0\) {0,1} *" "([2-9][0-9]{2})\$1[-]{0,1}" "([0-9]{4})\$2"</pre> <p>The three arguments to telno shown above must all be entered on one line.</p> <p>In the C program that uses the regcmp output,</p> <pre> regex(telno, line, area, exch, rest)</pre> <p>applies the regular expression named telno to line.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of regcmp for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how regcmp behaves.</p> <p>LC_CTYPE</p> <p>Determines how regcmp handles characters. When LC_CTYPE is set to a valid value, regcmp can display and handle text and filenames containing valid characters for that locale. regcmp can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. regcmp can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p>

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

regcmp(3G), environ(5)

NAME	regex – match patterns against a string
SYNOPSIS	regex [-e] [-v "string"] [pattern template] ... pattern [template]
DESCRIPTION	<p>The regex command takes a string from <i>the standard input</i>, and a list of <i>pattern</i> / <i>template</i> pairs, and runs regex () to compare the string against each <i>pattern</i> until there is a match. When a match occurs, regex writes the corresponding <i>template</i> to <i>the standard output</i> and returns TRUE. The last (or only) <i>pattern</i> does not need a template. If that is the pattern that matches the string, the function simply returns TRUE. If no match is found, regex returns FALSE.</p> <p>The argument <i>pattern</i> is a regular expression of the form described in regex (). In most cases <i>pattern</i> should be enclosed in single quotes to turn off special meanings of characters. Note that only the final <i>pattern</i> in the list may lack a <i>template</i> .</p> <p>The argument <i>template</i> may contain the strings \$m0 through \$m9 , which will be expanded to the part of <i>pattern</i> enclosed in (...)\$0 through (...)\$9 constructs (see examples below). Note that if you use this feature, you must be sure to enclose <i>template</i> in single quotes so that FMLI does not expand \$m0 through \$m9 at parse time. This feature gives regex much of the power of cut(1), paste(1), and grep(1), and some of the capabilities of sed(1). If there is no <i>template</i>, the default is \$m0\$m1\$m2\$m3\$m4\$m5\$m6\$m7\$m8\$m9 .</p>
OPTIONS	<p>-e Evaluate the corresponding template and write the result to <i>the standard output</i>.</p> <p>-v "string" Use <i>string</i> instead of <i>the standard input</i> to match against patterns.</p>
EXAMPLES	<p>To cut the 4th through 8th letters out of a string (this example will output strin and return TRUE):</p> <pre style="margin-left: 40px;">`regex -v "my string is nice" `.{3}(.{5})\$0' '\$m0`</pre> <p>In a form, to validate input to field 5 as an integer:</p> <pre style="margin-left: 40px;">valid=`regex -v "\$F5" `[0-9]+\$`</pre> <p>In a form, to translate an environment variable which contains one of the numbers 1, 2, 3, 4, 5 to the letters a, b, c, d, e:</p> <pre style="margin-left: 40px;">value=`regex -v "\$VAR1" 1 a 2 b 3 c 4 d 5 e '.*' 'Error`</pre> <p>Note the use of the pattern '.*' to mean "anything else".</p> <p>In the example below, all three lines constitute a single backquoted expression. This expression, by itself, could be put in a menu definition file. Since backquoted expressions are expanded as they are parsed, and output from a backquoted expression (the cat command, in this example) becomes part of the definition file being parsed, this expression would read /etc/passwd and make a dynamic menu of all the login ids on the system.</p> <pre style="margin-left: 40px;">`cat /etc/passwd regex `^([:]*)\$0.*\$` name=\$m0 action=`message "\$m0 is a user"``</pre>

DIAGNOSTICS

If none of the patterns match, **regex** returns FALSE, otherwise TRUE.

NOTES

Patterns and templates must often be enclosed in single quotes to turn off the special meanings of characters. Especially if you use the **\$m0** through **\$m9** variables in the template, since FMLI will expand the variables (usually to "") before **regex** even sees them.

Single characters in character classes (inside []) must be listed before character ranges, otherwise they will not be recognized. For example, [**a-zA-Z_**] will not find underscores (**_**) or slashes (**/**), but [**_/a-zA-Z**] will.

The regular expressions accepted by **regcmp** differ slightly from other utilities (that is, **sed**, **grep**, **awk**, **ed**, etc.).

regex with the **-e** option forces subsequent commands to be ignored. In other words if a backquoted statement appears as follows:

```
`regex -e ...; command1; command2`
```

command1 and *command2* would never be executed. However, dividing the expression into two:

```
`regex -e ...`command1; command2`
```

would yield the desired result.

SEE ALSO

awk(1), **cut(1)**, **grep(1)**, **paste(1)**, **sed(1)**, **regcmp(3G)**

NAME reinit – runs an initialization file

SYNOPSIS **reinit** *filename*

DESCRIPTION The **reinit** command is used to change the values of descriptors defined in the initialization file that was named when **fml**i was invoked and/or define additional descriptors. FMLI will parse and evaluate the descriptors in *filename*, and then continue running the current application. The argument *filename* must be the name of a valid FMLI initialization file.

The **reinit** command does not re-display the introductory frame or change the layout of screen labels for function keys.

NAME	renice – alter priority of running processes
SYNOPSIS	<i>/usr/ucb/renice priority pid ...</i> <i>/usr/ucb/renice priority [-p pid ...] [-g pgrp ...] [-u username ...]</i>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>The renice command alters the scheduling priority of one or more running processes. By default, the processes to be affected are specified by their process IDs. <i>priority</i> is the new priority value.</p> <p>Users other than the privileged user may only alter the priority of processes they own, and can only monotonically increase their “nice value” within the range 0 to 20. This prevents overriding administrative flats. The privileged user may alter the priority of any process and set the priority to any value in the range -20 to 20. Useful priorities are: 19 (the affected processes will run only when nothing else in the system wants to), 0 (the “base” scheduling priority) and any negative value (to make things go very fast).</p> <p>If only the priority is specified, the current process (alternatively, process group or user) is used.</p>
OPTIONS	<p>-p pid ... Specify a list of process IDs.</p> <p>-g pgrp ... Specify a list of process group IDs. The processes in the specified process groups have their scheduling priority altered.</p> <p>-u user ... Specify a list of user IDs or usernames. All processes owned by each <i>user</i> have their scheduling altered.</p>
FILES	<i>/etc/passwd</i> map user names to user ID's
SEE ALSO	prioctl(1)
NOTES	<p>If you make the priority very negative, then the process cannot be interrupted. To regain control you must make the priority greater than zero.</p> <p>Users other than the privileged user cannot increase scheduling priorities of their own processes, even if they were the ones that decreased the priorities in the first place. The prioctl command subsumes the function of renice .</p>

NAME	repeat – shell built-in function to execute a command more than once
SYNOPSIS csh	repeat <i>count</i> <i>command</i>
DESCRIPTION csh	Repeat <i>command</i> <i>count</i> times. <i>command</i> is restricted to a one-line statement.
SEE ALSO	csh(1)

NAME	reset – reset the current form field to its default values
SYNOPSIS	reset
DESCRIPTION	The reset function changes the entry in a field of a form to its default value; that is, the value displayed when the form was opened.

NAME	rlogin – remote login
SYNOPSIS	rlogin [-L] [-8] [-ec] [-l <i>username</i>] <i>hostname</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>rlogin establishes a remote login session from your terminal to the remote machine named <i>hostname</i>.</p> <p>Hostnames are listed in the <i>hosts</i> database, which may be contained in the /etc/hosts file, the Network Information Service (NIS) hosts map, the Internet domain name server, or a combination of these. Each host has one official name (the first name in the database entry), and optionally one or more nicknames. Either official hostnames or nicknames may be specified in <i>hostname</i>.</p> <p>Each remote machine may have a file named /etc/hosts.equiv containing a list of trusted hostnames with which it shares usernames. Users with the same username on both the local and remote machine may rlogin from the machines listed in the remote machine's /etc/hosts.equiv file without supplying a password. Individual users may set up a similar private equivalence list with the file .rhosts in their home directories. Each line in this file contains two names: a <i>hostname</i> and a <i>username</i> separated by a space. An entry in a remote user's .rhosts file permits the user named <i>username</i> who is logged into <i>hostname</i> to log in to the remote machine as the remote user without supplying a password. If the name of the local host is not found in the /etc/hosts.equiv file on the remote machine, and the local username and hostname are not found in the remote user's .rhosts file, then the remote machine will prompt for a password. Hostnames listed in /etc/hosts.equiv and .rhosts files must be the official hostnames listed in the hosts database; nicknames may not be used in either of these files.</p> <p>For security reasons, the .rhosts file must be owned by either the remote user or by root.</p> <p>The remote terminal type is the same as your local terminal type (as given in your environment TERM variable). The terminal or window size is also copied to the remote system if the server supports the option, and changes in size are reflected as well. All echoing takes place at the remote site, so that (except for delays) the remote login is transparent. Flow control using Ctrl-S and Ctrl-Q and flushing of input and output on interrupts are handled properly.</p>
OPTIONS	<p>-L Allow the rlogin session to be run in “litout” mode.</p> <p>-8 Pass eight-bit data across the net instead of seven-bit data.</p> <p>-ec Specify a different escape character, <i>c</i>, for the line used to disconnect from the remote host.</p> <p>-l <i>username</i> Specify a different <i>username</i> for the remote login. If you do not use this option, the remote username used is the same as your local username.</p>

Escape Sequences

Lines that you type which start with the tilde character are “escape sequences” (the escape character can be changed using the **-e** options):

- ~. Disconnect from the remote host — this is not the same as a logout, because the local host breaks the connection with no warning to the remote end.
- ~**susp** Suspend the login session (only if you are using a shell with Job Control). **susp** is your “suspend” character, usually Control-Z; see **tty(1)**.
- ~**dsusp** Suspend the input half of the login, but output will still be seen (only if you are using a shell with Job Control). **dsusp** is your “deferred suspend” character, usually Control-Y; see **tty(1)**.

FILES

/etc/passwd	
/usr/hosts/*	for <i>hostname</i> version of the command
/etc/hosts.equiv	list of trusted hostnames with shared usernames
\$HOME/.rhosts	private list of trusted hostname/username combinations

SEE ALSO

rsh(1), **stty(1)**, **tty(1)**, **in.named(1M)**, **hosts(4)**, **hosts.equiv(4)**

NOTES

When a system is listed in **hosts.equiv**, its security must be as good as local security. One insecure system listed in **hosts.equiv** can compromise the security of the entire system. The Network Information Service (NIS) was formerly known as Sun Yellow Pages (YP). The functionality of the two remains the same; only the name has changed. This implementation can only use the TCP network service.

NAME	rm , rmdir – remove files or directories
SYNOPSIS	rm [-f] [-i] <i>filename</i> ... rm -r [-f] [-i] <i>dirname</i> ... [<i>filename</i> ...] rmdir [-p] [-s] <i>dirname</i> ...
AVAILABILITY	SUNWcsu
DESCRIPTION	rm removes the entries for one or more files from a directory. If a file has no write permission and the standard input is a terminal, the full set of permissions (in octal) for the file are printed followed by a question mark. This is a prompt for confirmation. If the answer begins with y (for yes), the file is deleted, otherwise the file remains. If <i>filename</i> is a symbolic link, the link will be removed, but the file or directory to which it refers will not be deleted. A user does not need write permission on a symbolic link to remove it, provided they have write permissions in the directory. Note: If the standard input is not a terminal, the command will operate as if the -f option is in effect.
OPTIONS	The following options apply to rm : -f Remove all files (whether write-protected or not) in a directory without prompting the user. In a write-protected directory, however, files are never removed (whatever their permissions are), but no messages are displayed. If the removal of a write-protected directory is attempted, this option will not suppress an error message. -i Interactive. With this option, rm prompts for confirmation before removing any write-protected files. It overrides the -f option and remains in effect even if the standard input is not a terminal. -r Recursively remove directories and subdirectories in the argument list. The directory will be emptied of files and removed. Note: The user is normally prompted for removal of any write-protected files which the directory contains. The write-protected files are removed without prompting, however, if the -f option is used, or if the standard input is not a terminal and the -i option is not used. Symbolic links that are encountered with this option will not be traversed. If the removal of a non-empty, write-protected directory is attempted, the command will always fail (even if the -f option is used), resulting in an error message. The following options apply to rmdir : -p Allow users to remove the directory <i>dirname</i> and its parent directories which become empty. A message is printed on the standard output about whether the whole path is removed or part of the path remains for some reason.

-s Suppress the message printed on the standard error when **-p** is in effect.

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

rmdir(2), **unlink(2)**, **environ(5)**

DIAGNOSTICS

All messages are generally self-explanatory.

It is forbidden to remove the files "." and ".." in order to avoid the consequences of inadvertently doing something like the following:

```
rm -r .*
```

Both **rm** and **rmdir** return exit codes of 0 if all the specified directories are removed successfully. Otherwise, they return a non-zero exit code.

NOTES

A **--** permits the user to mark explicitly the end of any command line options, allowing **rm** to recognize filename arguments that begin with a **-**. As an aid to BSD migration, **rm** will accept **-** as a synonym for **--**. This migration aid may disappear in a future release. If a **--** and a **-** both appear on the same command line, the second will be interpreted as a filename.

NAME	roffbib – format and print a bibliographic database
SYNOPSIS	roffbib [-e] [-h] [-m filename] [-np] [-olist] [-Q] [-raN] [-sN] [-Tterm] [-V] [-x] [<i>filename</i>] . . .
AVAILABILITY	SUNWdoc
DESCRIPTION	roffbib prints out all records in a bibliographic database, in bibliography format rather than as footnotes or endnotes. Generally it is used in conjunction with sortbib (1): example% sortbib database roffbib
OPTIONS	roffbib accepts all options understood by nroff (1) except -i and -q . -e Produce equally-spaced words in adjusted lines using full terminal resolution. -h Use output tabs during horizontal spacing to speed output and reduce output character count. TAB settings are assumed to be every 8 nominal character widths. -m filename Prepend the macro file /usr/share/lib/tmac/tmac.name to the input files. There should be a space between the -m and the macro filename. This set of macros will replace the ones defined in /usr/share/lib/tmac/tmac.bib . -np Number first generated page <i>p</i> . -olist Print only page numbers that appear in the comma-separated <i>list</i> of numbers and ranges. A range <i>N–M</i> means pages <i>N</i> through <i>M</i> ; an initial -N means from the beginning to page <i>N</i> ; a final N– means from page <i>N</i> to end. -Q Queue output for the phototypesetter. Page offset is set to 1 inch. -raN Set register <i>a</i> (one-character) to <i>N</i> . The command-line argument -rN1 will number the references starting at 1. Four command-line registers control formatting style of the bibliography, much like the number registers of ms (5). The flag -rV2 will double space the bibliography, while -rV1 will double space references but single space annotation paragraphs. The line length can be changed from the default 6.5 inches to 6 inches with the -rL6i argument, and the page offset can be set from the default of 0 to one inch by specifying -rO1i (capital O, not zero). -sN Halt prior to every <i>N</i> pages for paper loading or changing (default <i>N</i> =1). To resume, enter NEWLINE or RETURN. -Tterm Specify <i>term</i> as the terminal type. -V Send output to the Versatec. Page offset is set to 1 inch. -x If abstracts or comments are entered following the %X field key, roffbib will format them into paragraphs for an annotated bibliography. Several %X fields may be given if several annotation paragraphs are desired.

FILES /usr/share/lib/tmac/tmac.bib file of macros used by **nroff/troff**

SEE ALSO **addbib(1)**, **indxbib(1)**, **lookbib(1)**, **nroff(1)** **refer(1)**, **sortbib(1)**, **troff(1)**

BUGS Users have to rewrite macros to create customized formats.

NAME	rpcgen – an RPC protocol compiler
SYNOPSIS	<p>rpcgen <i>infile</i></p> <p>rpcgen [-a] [-A] [-b] [-C] [-D <i>name</i> [= <i>value</i>]] [-i <i>size</i>] [-I [-K <i>seconds</i>]] [-L] [-M] [-N] [-T] [-Y <i>pathname</i>] <i>infile</i></p> <p>rpcgen [-c -h -l -m -t -Sc -Ss -Sm] [-o <i>outfile</i>] [<i>infile</i>]</p> <p>rpcgen [-s <i>nettype</i>] [-o <i>outfile</i>] [<i>infile</i>]</p> <p>rpcgen [-n <i>netid</i>] [-o <i>outfile</i>] [<i>infile</i>]</p>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>rpcgen is a tool that generates C code to implement an RPC protocol. The input to rpcgen is a language similar to C known as RPC Language (Remote Procedure Call Language).</p> <p>rpcgen is normally used as in the first synopsis where it takes an input file and generates three output files. If the <i>infile</i> is named proto.x, then rpcgen generates a header in proto.h, XDR routines in proto_xdr.c, server-side stubs in proto_svc.c, and client-side stubs in proto_clnt.c. With the -T option, it also generates the RPC dispatch table in proto_tbl.i.</p> <p>rpcgen can also generate sample client and server files that can be customized to suit a particular application. The -Sc, -Ss and -Sm options generate sample client, server and makefile, respectively. The -a option generates all files, including sample files. If the <i>infile</i> is proto.x, then the client side sample file is written to proto_client.c, the server side sample file to proto_server.c and the sample makefile to makefile.proto.</p> <p>The server created can be started both by the port monitors (for example, inetd or listen) or by itself. When it is started by a port monitor, it creates servers only for the transport for which the file descriptor 0 was passed. The name of the transport must be specified by setting up the environment variable PM_TRANSPORT. When the server generated by rpcgen is executed, it creates server handles for all the transports specified in NET_PATH environment variable, or if it is unset, it creates server handles for all the visible transports from /etc/netconfig file. Note: the transports are chosen at run time and not at compile time. When the server is self-started, it backgrounds itself by default. A special define symbol RPC_SVC_FG can be used to run the server process in foreground.</p> <p>The second synopsis provides special features which allow for the creation of more sophisticated RPC servers. These features include support for user provided #defines and RPC dispatch tables. The entries in the RPC dispatch table contain:</p> <ul style="list-style-type: none"> • pointers to the service routine corresponding to that procedure, • a pointer to the input and output arguments • the size of these routines

A server can use the dispatch table to check authorization and then to execute the service routine; a client library may use it to deal with the details of storage management and XDR data conversion.

The other three synopses shown above are used when one does not want to generate all the output files, but only a particular one. See the EXAMPLES section below for examples of **rpcgen** usage. When **rpcgen** is executed with the **-s** option, it creates servers for that particular class of transports. When executed with the **-n** option, it creates a server for the transport specified by *netid*. If *infile* is not specified, **rpcgen** accepts the standard input.

The C preprocessor, **cc -E** is run on the input file before it is actually interpreted by **rpcgen**. For each type of output file, **rpcgen** defines a special preprocessor symbol for use by the **rpcgen** programmer:

RPC_HDR defined when compiling into headers
RPC_XDR defined when compiling into XDR routines
RPC_SVC defined when compiling into server-side stubs
RPC_CLNT defined when compiling into client-side stubs
RPC_TBL defined when compiling into RPC dispatch tables

Any line beginning with “%” is passed directly into the output file, uninterpreted by **rpcgen**. To specify the path name of the C preprocessor use **-Y** flag.

For every data type referred to in *infile*, **rpcgen** assumes that there exists a routine with the string **xdr_** prepended to the name of the data type. If this routine does not exist in the RPC/XDR library, it must be provided. Providing an undefined data type allows customization of XDR routines.

OPTIONS

- a** Generate all files, including sample files.
- A** Enable the Automatic MT mode in the server main program. In this mode, the RPC library automatically creates threads to service client requests. This option generates multithread-safe stubs by implicitly turning on the **-M** option. Server multithreading modes and parameters can be set using the **rpc_control()** call. **rpcgen** generated code does not change the default values for the Automatic MT mode.
- b** Backward compatibility mode. Generate transport specific RPC code for older versions of the operating system.
- c** Compile into XDR routines.
- C** Generate header and stub files which can be used with ANSI C compilers. Headers generated with this flag can also be used with C++ programs.
- Dname[=value]** Define a symbol *name*. Equivalent to the **#define** directive in the source. If no *value* is given, *value* is defined as **1**. This option may be specified more than once.
- h** Compile into C data-definitions (a header). **-T** option can be used in conjunction to produce a header which supports RPC dispatch tables.

- i** *size* Size at which to start generating inline code. This option is useful for optimization. The default size is 5.
- I** Compile support for **inetd**(1M) in the server side stubs. Such servers can be self-started or can be started by **inetd**. When the server is self-started, it backgrounds itself by default. A special define symbol **RPC_SVC_FG** can be used to run the server process in foreground, or the user may simply compile without the **-I** option.
- If there are no pending client requests, the **inetd** servers exit after 120 seconds (default). The default can be changed with the **-K** option. All of the error messages for **inetd** servers are always logged with **syslog**(3).
- Note: This option is supported for backward compatibility only. It should always be used in conjunction with the **-b** option which generates backward compatibility code. By default (i.e., when **-b** is not specified), **rpcgen** generates servers that can be invoked through portmonitors.
- K** *seconds* By default, services created using **rpcgen** and invoked through portmonitors wait **120** seconds after servicing a request before exiting. That interval can be changed using the **-K** flag. To create a server that exits immediately upon servicing a request, use **-K 0**. To create a server that never exits, the appropriate argument is **-K -1**.
- When monitoring for a server, some portmonitors, like **listen**(1M), *always* spawn a new process in response to a service request. If it is known that a server will be used with such a monitor, the server should exit immediately on completion. For such servers, **rpcgen** should be used with **-K 0**.
- l** Compile into client-side stubs.
- L** When the servers are started in foreground, use **syslog**(3) to log the server errors instead of printing them on the standard error.
- m** Compile into server-side stubs, but do not generate a “main” routine. This option is useful for doing callback-routines and for users who need to write their own “main” routine to do initialization.
- M** Generate multithread-safe stubs for passing arguments and results between **rpcgen** generated code and user written code. This option is useful for users who want to use threads in their code.
- N** This option allows procedures to have multiple arguments. It also uses the style of parameter passing that closely resembles C. So, when passing an argument to a remote procedure, you do not have to pass a pointer to the argument, but can pass the argument itself. This behavior is different from the old style of **rpcgen** generated code. To maintain backward compatibility, this option is not the default.

-n <i>netid</i>	Compile into server-side stubs for the transport specified by <i>netid</i> . There should be an entry for <i>netid</i> in the netconfig database. This option may be specified more than once, so as to compile a server that serves multiple transports.
-o <i>outfile</i>	Specify the name of the output file. If none is specified, standard output is used (-c , -h , -l , -m , -n , -s , -Sc , -Sm , -Ss , and -t modes only).
-s <i>nettype</i>	Compile into server-side stubs for all the transports belonging to the class <i>nettype</i> . The supported classes are netpath , visible , circuit_n , circuit_v , datagram_n , datagram_v , tcp , and udp (see rpc(3N) for the meanings associated with these classes). This option may be specified more than once. Note: the transports are chosen at run time and not at compile time.
-Sc	Generate sample client code that uses remote procedure calls.
-Sm	Generate a sample Makefile which can be used for compiling the application.
-Ss	Generate sample server code that uses remote procedure calls.
-t	Compile into RPC dispatch table.
-T	Generate the code to support RPC dispatch tables. The options -c , -h , -l , -m , -s , -Sc , -Sm , -Ss , and -t are used exclusively to generate a particular type of file, while the options -D and -T are global and can be used with the other options.
-Y <i>pathname</i>	Give the name of the directory where rpcgen will start looking for the C-preprocessor.

EXAMPLES

The following example:

```
example% rpcgen -T prot.x
```

generates all the five files: **prot.h**, **prot_clnt.c**, **prot_svc.c**, **prot_xdr.c** and **prot_tbl.i**.

The following example sends the C data-definitions (header) to the standard output.

```
example% rpcgen -h prot.x
```

To send the test version of the **-DTEST**, server side stubs for all the transport belonging to the class **datagram_n** to standard output, use:

```
example% rpcgen -s datagram_n -DTEST prot.x
```

To create the server side stubs for the transport indicated by *netid* **tcp**, use:

```
example% rpcgen -n tcp -o prot_svc.c
```

SEE ALSO

cc(1B), **inetd(1M)**, **listen(1M)**, **syslog(3)**, **rpc(3N)**, **rpc_svc_calls(3N)**

The **rpcgen** chapter in the *Network Interfaces Programmer's Guide* manual.

NAME	rsh, remsh, remote_shell – remote shell
SYNOPSIS	rsh [-n] [-l username] hostname command rsh hostname [-n] [-l username] command remsh [-n] [-l username] hostname command remsh hostname [-n] [-l username] command hostname [-n] [-l username] command
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>rsh connects to the specified <i>hostname</i> and executes the specified <i>command</i>. rsh copies its standard input to the remote command, the standard output of the remote command to its standard output, and the standard error of the remote command to its standard error. Interrupt, quit and terminate signals are propagated to the remote command; rsh normally terminates when the remote command does.</p> <p>If you omit <i>command</i>, instead of executing a single command, rsh logs you in on the remote host using rlogin(1). Shell metacharacters which are not quoted are interpreted on the local machine, while quoted metacharacters are interpreted on the remote machine. See EXAMPLES.</p> <p>Hostnames are given in the <i>hosts</i> database, which may be contained in the <i>/etc/hosts</i> file, the Internet domain name database, or both. Each host has one official name (the first name in the database entry) and optionally one or more nicknames. Official hostnames or nicknames may be given as <i>hostname</i>.</p> <p>If the name of the file from which rsh is executed is anything other than rsh, rsh takes this name as its <i>hostname</i> argument. This allows you to create a symbolic link to rsh in the name of a host which, when executed, will invoke a remote shell on that host. By creating a directory and populating it with symbolic links in the names of commonly used hosts, then including the directory in your shell's search path, you can run rsh by typing <i>hostname</i> to your shell.</p> <p>If rsh is invoked with the basename remsh, rsh will check for the existence of the file <i>/usr/bin/remsh</i>. If this file exists, rsh will behave as if remsh is an alias for rsh. If <i>/usr/bin/remsh</i> does not exist, rsh will behave as if remsh is a host name.</p> <p>Each remote machine may have a file named <i>/etc/hosts.equiv</i> containing a list of trusted hostnames with which it shares usernames. Users with the same username on both the local and remote machine may rsh from the machines listed in the remote machine's <i>/etc/hosts</i> file. Individual users may set up a similar private equivalence list with the file <i>.rhosts</i> in their home directories. Each line in this file contains two names: a <i>hostname</i> and a <i>username</i> separated by a space. The entry permits the user named <i>username</i> who is logged into <i>hostname</i> to use rsh to access the remote machine as the remote user. If the name of the local host is not found in the <i>/etc/hosts.equiv</i> file on the remote machine, and the local username and hostname are not found in the remote user's <i>.rhosts</i> file, then the access is denied. The hostnames listed in the <i>/etc/hosts.equiv</i> and <i>.rhosts</i> files must be</p>

the official hostnames listed in the **hosts** database; nicknames may not be used in either of these files.

rsh will not prompt for a password if access is denied on the remote machine unless the *command* argument is omitted.

- OPTIONS**
- l *username*** Use *username* as the remote username instead of your local username. In the absence of this option, the remote username is the same as your local username.
 - n** Redirect the input of **rsh** to **/dev/null**. You sometimes need this option to avoid unfortunate interactions between **rsh** and the shell which invokes it. For example, if you are running **rsh** and invoke a **rsh** in the background without redirecting its input away from the terminal, it will block even if no reads are posted by the remote command. The **-n** option will prevent this.

The type of remote shell (**sh**, **rsh**, or other) is determined by the user's entry in the file **/etc/passwd** on the remote system.

EXIT CODES Returns **0** upon successful completion, **1** otherwise.

EXAMPLES The following command:

```
example% rsh lizard cat lizard.file >> example.file
```

appends the remote file **lizard.file** from the machine called "lizard" to the file called **example.file** on the machine called "example," while the command:

```
example% rsh lizard cat lizard.file ">>" lizard.file2
```

appends the file **lizard.file** on the machine called "lizard" to the file **another.lizard.file** which also resides on the machine called "lizard."

FILES **/etc/hosts**
/etc/passwd

SEE ALSO **rlogin(1)**, **vi(1)**, **in.named(1M)**, **hosts(4)**, **hosts.equiv(4)**

NOTES When a system is listed in **hosts.equiv**, its security must be as good as local security. One insecure system listed in **hosts.equiv** can compromise the security of the entire system. You cannot run an interactive command (such as **vi(1)**); use **rlogin** if you wish to do so. Stop signals stop the local **rsh** process only; this is arguably wrong, but currently hard to fix for reasons too complicated to explain here. The current local environment is not passed to the remote shell. Sometimes the **-n** option is needed for reasons that are less than obvious. For example, the command:

example% rsh somehost dd if=/dev/nrmt0 bs=20b | tar xvpBf -

will put your shell into a strange state. Evidently, what happens is that the **tar** terminates before the **rsh**. The **rsh** then tries to write into the “broken pipe” and, instead of terminating neatly, proceeds to compete with your shell for its standard input. Invoking **rsh** with the **-n** option avoids such incidents.

This bug occurs only when **rsh** is at the beginning of a pipeline and is not reading standard input. Do not use the **-n** if **rsh** actually needs to read standard input. For example,

example% tar cf - . | rsh sundial dd of=/dev/rmt0 obs=20b

does not produce the bug. If you were to use the **-n** in a case like this, **rsh** would incorrectly read from **/dev/null** instead of from the pipe.

NAME	run – run an executable
SYNOPSIS	run [g-s] [-e] [-n] [-t string] <i>program</i>
DESCRIPTION	The grun function runs <i>program</i> , using the PATH variable to find it. By default, when <i>program</i> has completed, the user is prompted (Press ENTER to continue:), before being returned to FMLI. The argument <i>program</i> is a system executable followed by its options (if any).
OPTIONS	<p>g-e If g-e is specified the user will be prompted before returning to FMLI only if there is an error condition</p> <p>g-n If g-n is specified the user will never be prompted before returning to FMLI (useful for programs like gvi, in which the user must do some specific action to exit in the first place).</p> <p>g-s The g-s option means "silent", implying that the screen will not have to be repainted when <i>program</i> has completed. NOTE: The g-s option should only be used when <i>program</i> does not write to the terminal. In addition, when g-s is used, <i>program</i> cannot be interrupted, even if it recognizes interrupts.</p> <p>g-tstring If g-t is specified, <i>string</i> is the name this process will have in the pop-up menu generated by the gfrm-list command. This feature requires the executable gfacesuspend, (See face(1)), to suspend the process and return to the FMLI application.</p>
EXAMPLE	<p>Here is a menu that uses grun:</p> <pre> gmenu="Edit special System files" name="Password file" action=`run -e vi /etc/passwd` name="Group file" action=`run -e vi /etc/group` name="My .profile" action=`run -n vi \$HOME/.profile` </pre>

NAME	rup – show host status of remote machines (RPC version)
SYNOPSIS	rup [-hlt] rup [<i>host...</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	<p>rup gives a status similar to uptime for remote machines. It broadcasts on the local network, and displays the responses it receives.</p> <p>Normally, the listing is in the order that responses are received, but this order can be changed by specifying one of the options listed below.</p> <p>When <i>host</i> arguments are given, rather than broadcasting rup will only query the list of specified hosts.</p> <p>A remote host will only respond if it is running the rstatd daemon, which is normally started up from inetd(1M).</p>
OPTIONS	-h Sort the display alphabetically by host name. -l Sort the display by load average. -t Sort the display by up time.
FILES	<i>/etc/servers</i>
SEE ALSO	runtime (1), inetd (1M) <i>SPARC: Installing Solaris Software</i> <i>x86: Installing Solaris Software</i>
BUGS	Broadcasting does not work through gateways.

NAME	rup – show host status of remote machines (RPC version)
SYNOPSIS	rup [-hlt] rup [<i>host...</i>]
DESCRIPTION	<p>rup gives a status similar to uptime for remote machines. It broadcasts on the local network, and displays the responses it receives.</p> <p>Normally, the listing is in the order that responses are received, but this order can be changed by specifying one of the options listed below.</p> <p>When <i>host</i> arguments are given, rather than broadcasting rup only queries the list of specified hosts.</p> <p>A remote host will only respond if it is running the rstatd daemon, which is normally started up from inetd(1M).</p>
OPTIONS	-h Sort the display alphabetically by host name. -l Sort the display by load average. -t Sort the display by up time.
SEE ALSO	runtime (1), inetd (1M)
BUGS	Broadcasting does not work through gateways.

NAME	ruptime – show host status of local machines
SYNOPSIS	ruptime [-alrtu]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>ruptime gives a status line like uptime for each machine on the local network; these are formed from packets broadcast by each host on the network once a minute.</p> <p>Machines for which no status report has been received for 5 minutes are shown as being down.</p> <p>Normally, the listing is sorted by host name, but this order can be changed by specifying one of the options listed below.</p>
OPTIONS	<p>-a Count even those users who have been idle for an hour or more.</p> <p>-l Sort the display by load average.</p> <p>-r Reverse the sorting order.</p> <p>-t Sort the display by up time.</p> <p>-u Sort the display by number of users.</p>
FILES	/var/spool/rwho/whod.* data files
SEE ALSO	rwwho(1) , in.rwhod(1M)

NAME	rusage – print resource usage for a command																																										
SYNOPSIS	<i>/usr/ucb/rusage command</i>																																										
AVAILABILITY	SUNWscpu																																										
DESCRIPTION	The rusage command is similar to time (1). It runs the given <i>command</i> , which must be specified; that is, <i>command</i> is not optional as it is in the C shell's timing facility. When the command is complete, rusage displays the real (wall clock), the system CPU, and the user CPU times which elapsed during execution of the command, plus other fields in the rusage structure, all on one long line. Times are reported in seconds and hundredths of a second.																																										
EXAMPLES	<p>The example below shows the format of rusage output.</p> <pre>example% rusage wc /usr/share/man/man1/csh (1) 3045 13423 78071 /usr/share/man/man1/csh (1) 2.26 real 0.80 user 0.36 sys 11 pf 38 pr 0 sw 11 rb 0 wb 16 vcx 37 icx 24 mx 0 ix 1230 id 9 is example%</pre> <p>Each of the fields identified corresponds to an element of the rusage structure, as described in getrusage(3B), as follows:</p> <table border="0"> <tr> <td>real</td> <td></td> <td>elapsed real time</td> </tr> <tr> <td>user</td> <td>ru_utime</td> <td>user time used</td> </tr> <tr> <td>sys</td> <td>ru_stime</td> <td>system time used</td> </tr> <tr> <td>pf</td> <td>ru_majflt</td> <td>page faults requiring physical I/O</td> </tr> <tr> <td>pr</td> <td>ru_minflt</td> <td>page faults not requiring physical I/O</td> </tr> <tr> <td>sw</td> <td>ru_nswap</td> <td>swaps</td> </tr> <tr> <td>rb</td> <td>ru_inblock</td> <td>block input operations</td> </tr> <tr> <td>wb</td> <td>ru_oublock</td> <td>block output operations</td> </tr> <tr> <td>vcx</td> <td>ru_nvcsw</td> <td>voluntary context switches</td> </tr> <tr> <td>icx</td> <td>ru_nivcsw</td> <td>involuntary context switches</td> </tr> <tr> <td>mx</td> <td>ru_maxrss</td> <td>maximum resident set size</td> </tr> <tr> <td>ix</td> <td>ru_ixrss</td> <td>currently 0</td> </tr> <tr> <td>id</td> <td>ru_idrss</td> <td>integral resident set size</td> </tr> <tr> <td>is</td> <td>ru_isrss</td> <td>currently 0</td> </tr> </table>	real		elapsed real time	user	ru_utime	user time used	sys	ru_stime	system time used	pf	ru_majflt	page faults requiring physical I/O	pr	ru_minflt	page faults not requiring physical I/O	sw	ru_nswap	swaps	rb	ru_inblock	block input operations	wb	ru_oublock	block output operations	vcx	ru_nvcsw	voluntary context switches	icx	ru_nivcsw	involuntary context switches	mx	ru_maxrss	maximum resident set size	ix	ru_ixrss	currently 0	id	ru_idrss	integral resident set size	is	ru_isrss	currently 0
real		elapsed real time																																									
user	ru_utime	user time used																																									
sys	ru_stime	system time used																																									
pf	ru_majflt	page faults requiring physical I/O																																									
pr	ru_minflt	page faults not requiring physical I/O																																									
sw	ru_nswap	swaps																																									
rb	ru_inblock	block input operations																																									
wb	ru_oublock	block output operations																																									
vcx	ru_nvcsw	voluntary context switches																																									
icx	ru_nivcsw	involuntary context switches																																									
mx	ru_maxrss	maximum resident set size																																									
ix	ru_ixrss	currently 0																																									
id	ru_idrss	integral resident set size																																									
is	ru_isrss	currently 0																																									
SEE ALSO	csh (1), time (1), getrusage (3B)																																										
BUGS	When the command being timed is interrupted, the timing values displayed may be inaccurate.																																										

NAME	rusers – who's logged in on remote machines
SYNOPSIS	rusers [-ahilu] <i>host</i> ...
AVAILABILITY	SUNWesu
DESCRIPTION	<p>The rusers command produces output similar to who(1), but for remote machines. The listing is in the order that responses are received, but this order can be changed by specifying one of the options listed below.</p> <p>The default is to print out the names of the users logged in. When the -l flag is given, additional information is printed for each user:</p> <p style="padding-left: 40px;"><i>userid hostname:terminal login date login time idle time login host</i></p> <p>If <i>hostname</i> and <i>login host</i> are the same value, the <i>login host</i> field is not displayed. Likewise, if <i>hostname</i> is not idle, the <i>idle time</i> is not displayed.</p> <p>A remote host will only respond if it is running the rusersd daemon, which may be started up from inetd(1M) or listen(1M).</p>
OPTIONS	<p>-a Give a report for a machine even if no users are logged on.</p> <p>-h Sort alphabetically by host name.</p> <p>-i Sort by idle time.</p> <p>-l Give a longer listing in the style of who(1).</p> <p>-u Sort by number of users.</p>
SEE ALSO	who (1), inetd (1M), listen (1M), pmadm (1M), sacadm (1M)

NAME	rwho – who's logged in on local machines
SYNOPSIS	rwho [-a]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The rwho command produces output similar to who(1), but for all machines on your network. If no report has been received from a machine for 5 minutes, rwho assumes the machine is down, and does not report users last known to be logged into that machine.</p> <p>If a user has not typed to the system for a minute or more, rwho reports this idle time. If a user has not typed to the system for an hour or more, the user is omitted from the output of rwho unless the -a flag is given.</p>
OPTIONS	-a Report all users whether or not they have typed to the system in the past hour.
FILES	/var/spool/rwho/whod.* information about other machines
SEE ALSO	finger (1), ruptime (1), who (1), in.rwhod (1M)
NOTES	<p>rwho does not work through gateways.</p> <p>The directory /var/spool/rwho must exist on the host from which rwho is run.</p> <p>This service takes up progressively more network bandwidth as the number of hosts on the local net increases. For large networks, the cost becomes prohibitive.</p> <p>The rwho service daemon, in.rwhod(1M), must be enabled for this command to return useful results.</p>

NAME	sag – system activity graph
SYNOPSIS	sag [<i>-e time</i>] [<i>-f file</i>] [<i>-i sec</i>] [<i>-s time</i>] [<i>-T term</i>] [<i>-x spec</i>] [<i>-y spec</i>]
DESCRIPTION	sag graphically displays the system activity data stored in a binary data file by a previous sar (1) run. Any of the sar data items may be plotted singly, or in combination; as cross plots, or versus time. Simple arithmetic combinations of data may be specified. sag invokes sar and finds the desired data by string-matching the data column header (run sar to see what is available).
OPTIONS	<p>These <i>options</i> are passed through to sar:</p> <p><i>-e time</i> Select data up to <i>time</i>. Default is 18:00.</p> <p><i>-f file</i> Use <i>file</i> as the data source for sar. Default is the current daily data file <i>/usr/adm/sa/sadd</i>.</p> <p><i>-i sec</i> Select data at intervals as close as possible to <i>sec</i> seconds.</p> <p><i>-s time</i> Select data later than <i>time</i> in the form <i>hh [:mm]</i>. Default is 08:00.</p> <p>Other <i>options</i>:</p> <p><i>-T term</i> Produce output suitable for terminal <i>term</i>. Default for <i>term</i> is \$TERM.</p> <p><i>-x spec</i> x axis specification with <i>spec</i> in the form: <i>name [op name] ... [lo hi]</i></p> <p><i>name</i> is either a string that will match a column header in the sar report, with an optional device name in square brackets, for example, r+w/s[dsk-1], or an integer value. <i>op</i> is + - * or / surrounded by blank spaces. Up to five names may be specified. Parentheses are not recognized. Contrary to custom, + and - have precedence over * and /. Evaluation is left to right. Thus, A / A + B * 100 is evaluated as (A/(A+B))*100, and A + B / C + D is (A+B)/(C+D). <i>lo</i> and <i>hi</i> are optional numeric scale limits. If unspecified, they are deduced from the data.</p> <p>Enclose <i>spec</i> in double-quotes ("<i> </i>") if it includes white space.</p> <p>A single <i>spec</i> is permitted for the x axis. If unspecified, <i>time</i> is used.</p> <p><i>-y spec</i> y axis specification with <i>spec</i> in the same form as for <i>-x</i>. Up to 5 <i>spec</i>'s separated by a semi-colon (;) may be given for <i>-y</i>. The <i>-y</i> default is:</p> <p style="padding-left: 40px;">-y "%usr 0 100; %usr + %sys 0 100; %usr + %sys + %wio 0 100"</p>

EXAMPLES

To see today's CPU utilization:

```
example% sag
```

To see activity over 15 minutes of all disk drives:

```
example% set TS=`date +%H:%M`
```

```
example% sar -o tempfile 60 15
```

```
example% set TE=`date +%H:%M`
```

```
example% sag -f tempfile -s $TS -e $TE -y "r+w/s[dsk]"
```

FILES

/usr/adm/sa/sadd daily data file for day *dd*

SEE ALSO

sar(1)

NAME	sar – system activity reporter
SYNOPSIS	sar [-aAbcdgkmpqruvw y] [-o filename] <i>t</i> [<i>n</i>] sar [-aAbcdgkmpqruvw y] [-e time] [-f filename] [-i sec] [-s time]
DESCRIPTION	<p>In the first instance sar samples cumulative activity counters in the operating system at <i>n</i> intervals of <i>t</i> seconds, where <i>t</i> should be 5 or greater. If <i>t</i> is specified with more than one option, all headers are printed together and the output may be difficult to read. (If the sampling interval is less than 5, the activity of sar itself may effect the sample.) If the -o option is specified, it saves the samples in <i>filename</i> in binary format. The default value of <i>n</i> is 1.</p> <p>In the second instance no sampling interval is specified. sar extracts data from a previously recorded <i>filename</i>, either the one specified by the -f option or, by default, the standard system activity daily data file <code>/var/adm/sa/sadd</code> for the current day <i>dd</i>. The starting and ending times of the report can be bounded using the -e and -s arguments with <i>time</i> specified in the form <i>hh[:mm[:ss]]</i>. The -i option selects records at <i>sec</i> second intervals. Otherwise, all intervals found in the data file are reported.</p>
OPTIONS	<p>The following options modify the subsets of information reported by sar.</p> <p>-a Report use of file access system routines: iget/s, namei/s, dirblk/s.</p> <p>-A Report all data. Equivalent to -abcdgkmpqruvwy.</p> <p>-b Report buffer activity: bread/s, bwrit/s – transfers per second of data between system buffers and disk or other block devices; lread/s, lwrit/s – accesses of system buffers; %rcache, %wcache – cache hit ratios, that is, (1–bread/lread) as a percentage; pread/s, pwrit/s – transfers using raw (physical) device mechanism.</p> <p>-c Report system calls: scall/s – system calls of all types; sread/s, swrit/s, fork/s, exec/s – specific system calls; rchar/s, wchar/s – characters transferred by read and write system calls. No incoming or outgoing exec(2) and fork(2) calls are reported.</p> <p>-d Report activity for each block device (for example, disk or tape drive) with the exception of XDC disks and tape drives. When data is displayed, the device specification <i>dsk</i>- is generally used to represent a disk drive. The device specification used to represent a tape drive is machine dependent. The activity data reported is: %busy, avque – portion of time device was busy servicing a transfer request, average number of requests outstanding during that time; read/s, write/s, blks/s – number of read/write transfers from or to device, number of bytes transferred in 512-byte units; avseek – number of milliseconds per average seek.</p>

For more general system statistics, use **iostat(1M)**, **sar(1M)**, or **vmstat(1M)**.

See *Peripherals Administration* for naming conventions for disks.

- g** Report paging activities:
 pgout/s – page-out requests per second;
 ppgout/s – pages paged-out per second;
 pgfree/s – pages per second placed on the free list by the page stealing daemon;
 pgscan/s – pages per second scanned by the page stealing daemon.
 %ufs_ipf – the percentage of UFS inodes taken off the freelist by iget which had reusable pages associated with them. These pages are flushed and cannot be reclaimed by processes. Thus this is the percentage of igets with page flushes.
- k** Report kernel memory allocation (KMA) activities:
 sml_mem, alloc, fail – information about the memory pool reserving and allocating space for small requests: the amount of memory in bytes KMA has for the small pool, the number of bytes allocated to satisfy requests for small amounts of memory, and the number of requests for small amounts of memory that were not satisfied (failed);
 lg_mem, alloc, fail – information for the large memory pool (analogous to the information for the small memory pool);
 ovsz_alloc, fail – the amount of memory allocated for oversize requests and the number of oversize requests which could not be satisfied (because oversized memory is allocated dynamically, there is not a pool).
- m** Report message and semaphore activities:
 msg/s, sema/s – primitives per second.
- p** Report paging activities:
 atch/s – page faults per second that are satisfied by reclaiming a page currently in memory (attaches per second);
 pgin/s – page-in requests per second;
 ppgin/s – pages paged-in per second;
 pflt/s – page faults from protection errors per second (illegal access to page or "copy-on-writes");
 vflt/s – address translation page faults per second (valid page not in memory);
 slock/s – faults per second caused by software lock requests requiring physical I/O.
- q** Report average queue length while occupied, and % of time occupied:
 runq-sz, %runocc – run queue of processes in memory and runnable;
 swpq-sz, %swpocc – these are no longer reported by sar.
- r** Report unused memory pages and disk blocks:
 freemem – average pages available to user processes;
 freeswap – disk blocks available for page swapping.

- u** Report CPU utilization (the default):
%usr, %sys, %wio, %idle – portion of time running in user mode, running in system mode, idle with some process waiting for block I/O, and otherwise idle.
- v** Report status of process, i-node, file tables:
proc-sz, inod-sz, file-sz, lock-sz – entries/size for each table, evaluated once at sampling point;
ov – overflows that occur between sampling points for each table.
- w** Report system swapping and switching activity:
swpin/s, swpot/s, bswin/s, bswot/s – number of transfers and number of 512-byte units transferred for swapins and swapouts (including initial loading of some programs);
pswch/s – process switches.
- y** Report TTY device activity:
rawch/s, canch/s, outch/s – input character rate, input character rate processed by canon, output character rate;
rcvin/s, xmtin/s, mdmin/s – receive, transmit and modem interrupt rates.
- e time** Select data up to *time*. Default is 18:00.
- f filename** Use *filename* as the data source for **sar**. Default is the current daily data file **/usr/adm/sa/sadd**.
- i sec** Select data at intervals as close as possible to *sec* seconds.
- o filename**
Save samples in file, *filename*, in binary format.
- s time** Select data later than *time* in the form *hh[:mm]*. Default is 08:00.

EXAMPLES

To see today's CPU activity so far:

```
example% sar
```

To watch CPU activity evolve for 10 minutes and save data:

```
example% sar -o temp 60 10
```

To later review disk and tape activity from that period:

```
example% sar -d -f temp
```

FILES

/var/adm/sa/sadd daily data file, where *dd* are digits representing the day of the month

SEE ALSO

sag(1), **iostat(1M)**, **sar(1M)**, **vmstat(1M)**, **exec(2)**, **fork(2)**

Security, Performance, and Accounting Administration

Peripherals Administration

NAME	sccs-admin, admin – create and administer SCCS history files
SYNOPSIS	<pre> /usr/ccs/bin/admin [-bhnz] [-ausername groupid] ... [-dflag] ... [-eusername groupid] ... [-fflag [value]] ... [-i [filename]] [-l a release[,release..]] [-m mr-list] [-rrelease] [-t [description-file]] [-y[comment]] s.filename ... </pre>
DESCRIPTION	<p>admin creates or modifies the flags and other parameters of SCCS history files. Filenames of SCCS history files begin with the 's.' prefix, and are referred to as s.files, or "history" files.</p> <p>The named s.file is created if it does not exist already. Its parameters are initialized or modified according to the options you specify. Parameters not specified are given default values when the file is initialized, otherwise they remain unchanged.</p> <p>If a directory name is used in place of the <i>s.filename</i> argument, the admin command applies to all s.files in that directory. Unreadable s.files produce an error. The use of '-' as the <i>s.filename</i> argument indicates that the names of files are to be read from the standard input, one s.file per line.</p>
OPTIONS	<p>-b Force encoding of binary data. Files that contain ASCII NUL or other control characters, or that do not end with a NEWLINE, are recognized as binary data files. The contents of such files are stored in the history file in encoded form. See uuencode(1C) for details about the encoding. This option is normally used in conjunction with -i to force admin to encode initial versions not recognized as containing binary data.</p> <p>-h Check the structure of an existing s.file (see sccsfile(4)), and compare a newly computed check-sum with one stored in the first line of that file. -h inhibits writing on the file; and so nullifies the effect of any other options.</p> <p>-n Create a new SCCS history file.</p> <p>-z Recompute the file check-sum and store it in the first line of the s.file. Caution: it is important to verify the contents of the history file (see sccs-val(1), and the print subcommand in sccs(1)), since using -z on a truly corrupted file may prevent detection of the error.</p> <p>-ausername groupid Add a user name, or a numerical group ID, to the list of users who may check deltas in or out. If the list is empty, any user is allowed to do so.</p> <p>-dflag Delete the indicated <i>flag</i> from the SCCS file. The -d option may be specified only for existing s.files. See -f for the list of recognized flags.</p> <p>-eusername groupid Erase a user name or group ID from the list of users allowed to make deltas.</p>

-fflag [*value*]

Set the indicated *flag* to the (optional) *value* specified. The following flags are recognized:

- b** Enable branch deltas. When **b** is set, branches can be created using the **-b** option of the SCCS **get** command (see **sccs-get(1)**).
- ceil** Set a ceiling on the releases that can be checked out. *ceil* is a number less than or equal to 9999. If **c** is not set, the ceiling is 9999.
- ffloor** Set a floor on the releases that can be checked out. The floor is a number greater than 0 but less than 9999. If **f** is not set, the floor is 1.
- dsid** The default delta number, or SID, to be used by an SCCS **get** command.
- i** Treat the 'No id keywords (ge6)' message issued by an SCCS **get** or **delta** command as an error rather than a warning.
- j** Allow concurrent updates.
- la**
lrelease[,*release* . . .]
Lock the indicated list of releases against deltas. If **a** is used, lock out deltas to all releases. An SCCS '**get -e**' command fails when applied against a locked release.
- n** Create empty releases when releases are skipped. These null (empty) deltas serve as anchor points for branch deltas.
- qvalue** Supply a *value* to which the **%Q%** keyword is to expand when a read-only version is retrieved with the SCCS **get** command.
- mmodule**
Supply a value for the module name to which the **%M%** keyword is to expand. If the **m** flag is not specified, the value assigned is the name of the SCCS file with the leading **s**. removed.
- ttype** Supply a value for the module type to which the **%Y%** keyword is to expand.
- v** [*program*]
Specify a validation *program* for the MR numbers associated with a new delta. The optional *program* specifies the name of an MR number validity checking *program*. If this flag is set when creating an SCCS file, the **-m** option must also be used, in which case the list of MRs may be empty.

-i [*filename*]

Initialize the history file with text from the indicated file. This text constitutes the initial delta, or set of checked-in changes. If *filename* is omitted, the initial text is obtained from the standard input. Omitting the **-i** option altogether creates an empty *s.file*. You can only initialize one *s.file* with text using **-i**. This option implies the **-n** option.

-la**-l** [*release[,release...]*]

Unlock the specified releases so that deltas can be checked in. If **a** is specified, allow deltas to be checked in for all releases.

-m [*mr-list*]

Insert the indicated Modification Request (MR) numbers into the commentary for the initial version. When specifying more than one MR number on the command line, *mr-list* takes the form of a quoted, space-separated list. A warning results if the **v** flag is not set or the MR validation fails.

-rrelease

Specify the release for the initial delta. **-r** may be used only in conjunction with **-i**. The initial delta is inserted into release 1 if this option is omitted. The level of the initial delta is always 1; initial deltas are named 1.1 by default.

-t [*description-file*]

Insert descriptive text from the file *description-file*. When **-t** is used in conjunction with **-n**, or **-i** to initialize a new *s.file*, the *description-file* must be supplied. When modifying the description for an existing file: a **-t** option without a *description-file* removes the descriptive text, if any; a **-t** option with a *description-file* replaces the existing text.

-y [*comment*]

Insert the indicated *comment* in the "Comments:" field for the initial delta. Valid only in conjunction with **-i** or **-n**. If **-y** option is omitted, a default comment line is inserted that notes the date and time the history file was created.

FILES

s.*	history file
SCCS/s.*	history file in SCCS subdirectory
z.*	temporary lock file

SEE ALSO

sccs(1), **sccs-cdc(1)**, **sccs-delta(1)**, **sccs-get(1)**, **sccs-help(1)**, **sccs-rmdel(1)**, **sccs-val(1)**, **sccsfile(4)**

Programming Utilities Guide

DIAGNOSTICS

Use the SCCS **help** command for explanations (see **sccs-help(1)**).

WARNINGS

The last component of all SCCS filenames must have the 's.' prefix. New SCCS files are given mode 444 (see **chmod(1)**). All writing done by **admin** is to a temporary file with an **x.** prefix, created with mode 444 for a new SCCS file, or with the same mode as an existing SCCS file. After successful execution of **admin**, the existing **s.** file is removed and replaced with the **x.file**. This ensures that changes are made to the SCCS file only when no errors have occurred.

It is recommended that directories containing SCCS files have permission mode 755, and that the **s.files** themselves have mode 444. The mode for directories allows only the owner to modify the SCCS files contained in the directories, while the mode of the **s.files** prevents all modifications except those performed using SCCS commands.

If it should be necessary to patch an SCCS file for any reason, the mode may be changed to 644 by the owner to allow use of a text editor. However, extreme care must be taken when doing this. The edited file should *always* be processed by an '**admin -h**' to check for corruption, followed by an '**admin -z**' to generate a proper check-sum. Another '**admin -h**' is recommended to ensure that the resulting **s.file** is valid.

admin also uses a temporary lock **s.file**, starting with the 'z.' prefix, to prevent simultaneous updates to the **s.file**. See **sccs-get(1)** for further information about the 'z.file'.

NAME	sccs-cdc, cdc – change the delta commentary of an SCCS delta
SYNOPSIS	<code>/usr/ccs/bin/cdc -rsid [-mmr-list] [-y [comment]] s.filename . . .</code>
DESCRIPTION	<p>cdc annotates the delta commentary for the SCCS delta ID (SID) specified by the -r option in each named s.file.</p> <p>If the v flag is set in the s.file, you can also use cdc to update the Modification Request (MR) list.</p> <p>If you checked in the delta, or, if you own the file and directory and have write permission, you can use cdc to annotate the commentary.</p> <p>Rather than replacing the existing commentary, cdc inserts the new comment you supply, followed by a line of the form:</p> <pre style="margin-left: 4em;">*** CHANGED *** yy/mm/dd hh/mm/ss username</pre> <p>above the existing commentary.</p> <p>If a directory is named as the s.filename argument, the cdc command applies to all s.files in that directory. Unreadable s.files produce an error; processing continues with the next file (if any). If '-' is given as the s.filename argument, each line of the standard input is taken as the name of an SCCS history file to be processed, and the -m and -y options must be used.</p>
OPTIONS	<p>-rsid Specify the SID of the delta to change.</p> <p>-mmr-list Specify one or more MR numbers to add or delete. When specifying more than one MR on the command line, <i>mr-list</i> takes the form of a quoted, space-separated list. To delete an MR number, precede it with a ! character (an empty MR list has no effect). A list of deleted MRs is placed in the comment section of the delta commentary. If -m is not used and the standard input is a terminal, cdc prompts with MRs? for the list (before issuing the comments? prompt). -m is only useful when the v flag is set in the s.file. If that flag has a value, it is taken to be the name of a program to validate the MR numbers. If that validation program returns a non-zero exit status, cdc terminates and the delta commentary remains unchanged.</p> <p>-y[comment] Use <i>comment</i> as the annotation in the delta commentary. The previous comments are retained; the <i>comment</i> is added along with a notation that the commentary was changed. A null <i>comment</i> leaves the commentary unaffected. If -y is not specified and the standard input is a terminal, cdc prompts with comments? for the text of the notation to be added. An unescaped NEWLINE character terminates the annotation text.</p>

EXAMPLES	<p>The following command: example% cdc -r1.6 -y"corrected commentary" s.program.c produces the following annotated commentary for delta 1.6 in s.program.c:</p> <pre>D 1.6 88/07/05 23:21:07 username 9 0 00001/00000/00000 MRs: COMMENTS: corrected commentary *** CHANGED *** 88/07/07 14:09:41 username performance enhancements in main()</pre>
FILES	<p>z.file temporary lock file</p>
SEE ALSO	<p>sccs(1), sccs-admin(1), sccs-comb(1), sccs-delta(1), sccs-help(1), sccs-prs(1), sccs-prt(1), sccs-rmdel(1), what(1), sccsfile(4) <i>Programming Utilities Guide</i></p>
DIAGNOSTICS	<p>Use the SCCS help command for explanations (see sccs-help(1)).</p>

NAME	sccs-comb, comb – combine SCCS deltas				
SYNOPSIS	<code>/usr/ccs/bin/comb [-os] [-csid-list] [-psid] s.filename ...</code>				
DESCRIPTION	<p>comb generates a shell script (see sh(1)) that you can use to reconstruct the indicated s.files. This script is written to the standard output.</p> <p>If a directory name is used in place of the <i>s.filename</i> argument, the comb command applies to all s.files in that directory. Unreadable s.files produce an error; processing continues with the next file (if any). The use of '-' as the <i>s.filename</i> argument indicates that the names of files are to be read from the standard input, one s.file per line.</p> <p>If no options are specified, comb preserves only the most recent (leaf) delta in a branch, and the minimal number of ancestors needed to preserve the history.</p>				
OPTIONS	<p>-o For each 'get -e' generated, access the reconstructed file at the release of the delta to be created. Otherwise, the reconstructed file is accessed at the most recent ancestor. The use of -o may decrease the size of the reconstructed s.file. It may also alter the shape of the delta tree of the original file.</p> <p>-s Generate scripts to gather statistics, rather than combining deltas. When run, the shell scripts report: the file name, size (in blocks) after combining, original size (also in blocks), and the percentage size change, computed by the formula: $100 * (original - combined) / original$ This option can be used to calculate the space that will be saved, before actually doing the combining.</p> <p>-csid-list Include the indicated list of deltas. All other deltas are omitted. <i>sid-list</i> is a comma-separated list of SCCS delta IDs (SIDs). To specify a range of deltas, use a '-' separator instead of a comma, between two SIDs in the list.</p> <p>-pSID The SID of the oldest delta to be preserved.</p>				
FILES	<table border="0"> <tr> <td>s.COMB</td> <td>reconstructed SCCS file</td> </tr> <tr> <td>comb?????</td> <td>temporary file</td> </tr> </table>	s.COMB	reconstructed SCCS file	comb?????	temporary file
s.COMB	reconstructed SCCS file				
comb?????	temporary file				
SEE ALSO	<p>sccs(1), sccs-admin(1), sccs-cdc(1), sccs-delta(1), sccs-help(1), sccs-prs(1), sccs-prt(1), sccs-rmdel(1), sccs-sccsdiff(1), what(1), sccsfile(4)</p> <p><i>Programming Utilities Guide</i></p>				
DIAGNOSTICS	Use the SCCS help command for explanations (see sccs-help (1)).				
BUGS	comb may rearrange the shape of the tree of deltas. It may not save any space; in fact, it is possible for the reconstructed file to actually be larger than the original.				

NAME	sccs-delta, delta – make a delta to an SCCS file
SYNOPSIS	<code>/usr/ccs/bin/delta [-nps] [-gsid-list] [-mmr-list] [-rsid] [-y[comment]] s.filename ...</code>
DESCRIPTION	<p>delta checks in a record of the line-by-line differences made to a checked-out version of a file under SCCS control. These changes are taken from the writable working copy that was retrieved using the SCCS get command (see sccs-get(1)). This working copy does not have the 's.' prefix, and is also referred to as a g-file.</p> <p>If a directory name is used in place of the <i>s.filename</i>, argument, the delta command applies to all s.files in that directory. Unreadable s.files produce an error; processing continues with the next file (if any). The use of '-' as the <i>s.filename</i> argument indicates that the names of files are to be read from the standard input, one s.file per line (requires -y, and in some cases, -m).</p> <p>delta may issue prompts on the standard output depending upon the options specified and the flags that are set in the s.file (see sccs-admin(1), and the -m and -y options below, for details).</p>
OPTIONS	<p>-n Retain the edited g-file, which is normally removed at the completion of processing.</p> <p>-p Display line-by-line differences (in diff(1) format) on the standard output.</p> <p>-s Silent. Do not display warning or confirmation messages. Do not suppress error messages (which are written to standard error).</p> <p>-gsid-list Specify a list of deltas to omit when the file is accessed at the SCCS version ID (SID) created by this delta. <i>sid-list</i> is a comma-separated list of SIDs. To specify a range of deltas, use a '-' separator instead of a comma, between two SIDs in the list.</p> <p>-m [mr-list] If the SCCS file has the v flag set (see sccs-admin(1)), you must supply one or more Modification Request (MR) numbers for the new delta. When specifying more than one MR number on the command line, <i>mr-list</i> takes the form of a quoted, space-separated list. If -m is not used and the standard input is a terminal, delta prompts with MRs? for the list (before issuing the comments? prompt). If the v flag in the s.file has a value, it is taken to be the name of a program to validate the MR numbers. If that validation program returns a non-zero exit status, delta terminates without checking in the changes.</p> <p>-rsid When two or more versions are checked out, specify the version to check in. This SID value can be either the SID specified on the get command line, or the SID of the new version to be checked in as reported by get. A diagnostic results if the specified SID is ambiguous, or if one is required but not supplied.</p>

-y*[comment]*

Supply a comment for the delta table (version log). A null comment is accepted, and produces an empty commentary in the log. If **-y** is not specified and the standard input is a terminal, **delta** prompts with '**comments?**'. An unescaped NEWLINE terminates the comment.

FILES

d.file temporary file of differences
p.file lock file for a checked-out version
q.file temporary file
s.file SCCS history file
x.file temporary copy of the **s.file**
z.file temporary file

SEE ALSO

sccs(1), **sccs-admin(1)**, **sccs-cdc(1)**, **sccs-get(1)**, **sccs-help(1)**, **sccs-prs(1)**, **sccs-prt(1)**, **sccs-rmdel(1)**, **sccs-sccsdiff(1)**, **sccs-unget(1)**, **what(1)**, **sccsfile(4)**

Programming Utilities Guide

DIAGNOSTICS

Use the SCCS **help** command for explanations (see **sccs-help(1)**).

WARNINGS

Lines beginning with an ASCII SOH character (binary 001) cannot be placed in the SCCS file unless the SOH is escaped. This character has special meaning to SCCS (see **sccsfile(4)**) and produces an error.

NAME	sccs-get, get – retrieve a version of an SCCS file
SYNOPSIS	<code>/usr/ccs/bin/get [-begkmpst] [-l [p]] [-asequence] [-cdate-time] [-Gg-file] [-isid-list] [-rsid] [-xsid-list] s.filename . . .</code>
DESCRIPTION	<p>get retrieves a working copy from the SCCS history file, according to the specified options.</p> <p>For each <i>s.filename</i> argument, get displays the SCCS delta ID (SID) and number of lines retrieved.</p> <p>If a directory name is used in place of the <i>s.filename</i> argument, the get command applies to all <i>s.files</i> in that directory. Unreadable <i>s.files</i> produce an error; processing continues with the next file (if any). The use of ‘-’ as the <i>s.filename</i> argument indicates that the names of files are to be read from the standard input, one <i>s.file</i> per line.</p> <p>The retrieved file normally has the same filename base as the <i>s.file</i>, less the prefix, and is referred to as the <i>g-file</i>.</p> <p>For each file processed, get responds (on the standard output) with the SID being accessed, and with the number of lines retrieved from the <i>s.file</i>.</p>
OPTIONS	<p>-b Create a new branch. Used with the -e option to indicate that the new delta should have an SID in a new branch. Instead of incrementing the level for version to be checked in, get indicates in the <i>p.file</i> that the delta to be checked in should either initialize a new branch and sequence (if there is no existing branch at the current level), or increment the branch component of the SID. If the b flag is not set in the <i>s.file</i>, this option is ignored.</p> <p>-e Retrieve a version for editing. With this option, get places a lock on the <i>s.file</i>, so that no one else can check in changes to the version you have checked out. If the j flag is set in the <i>s.file</i>, the lock is advisory: get issues a warning message. Concurrent use of ‘get -e’ for different SIDs is allowed, however, get will not check out a version of the file if a writable version is present in the directory. All SCCS file protections stored in the <i>s.file</i>, including the release ceiling, floor, and authorized user list, are honored by ‘get -e’.</p> <p>-g Get the SCCS version ID, without retrieving the version itself. Used to verify the existence of a particular SID.</p> <p>-k Suppress expansion of ID keywords. -k is implied by the -e.</p> <p>-m Precede each retrieved line with the SID of the delta in which it was added to the file. The SID is separated from the line with a TAB.</p> <p>-n Precede each line with the %M% ID keyword and a TAB. When both the -m and -n options are used, the ID keyword precedes the SID, and the line of text.</p> <p>-p Write the text of the retrieved version to the standard output. All messages that normally go to the standard output are written to the standard error instead.</p>

- s** Suppress all output normally written on the standard output. However, fatal error messages (which always go to the standard error) remain unaffected.
- t** Retrieve the most recently created (top) delta in a given release (for example: **-r1**).
- l [p]** Retrieve a summary of the delta table (version log) and write it to a listing file, with the 'l.' prefix (called 'l.file'). When **-lp** is used, write the summary onto the standard output.
- a sequence**
Retrieve the version corresponding to the indicated delta sequence number. This option is used primarily by the SCCS **comb** command (see **sccs-comb(1)**); for users, **-r** is an easier way to specify a version. **-a** supersedes **-r** when both are used.
- cdate-time**
Retrieve the latest version checked in prior to the date and time indicated by the *date-time* argument. *date-time* takes the form: *yy[mm[dd[hh[mm[ss]]]]]*. Units omitted from the indicated date and time default to their maximum possible values; that is **-c7502** is equivalent to **-c750228235959**. Any number of non-numeric characters may separate the various 2 digit components. If white-space characters occur, the *date-time* specification must be quoted.
- Gnewname**
Use *newname* as the name of the retrieved version.
- isid-list**
Specify a list of deltas to include in the retrieved version. The included deltas are noted in the standard output message. *sid-list* is a comma-separated list of SIDs. To specify a range of deltas, use a '-' separator instead of a comma, between two SIDs in the list.
- rsid** Retrieve the version corresponding to the indicated SID (delta).
The SID for a given delta is a number, in Dewey decimal format, composed of two or four fields: the *release* and *level* fields, and for branch deltas, the *branch* and *sequence* fields. For instance, if **1.2** is the SID, **1** is the release, and **2** is the level number. If **1.2.3.4** is the SID, **3** is the branch and **4** is the sequence number.
You need not specify the entire SID to retrieve a version with **get**. When you omit **-r** altogether, or when you omit both release and level, **get** normally retrieves the highest release and level. If the **d** flag is set to an SID in the **s.file** and you omit the SID, **get** retrieves the default version indicated by that flag.
When you specify a release but omit the level, **get** retrieves the highest level in that release. If that release does not exist, **get** retrieves highest level from the next-highest existing release.
Similarly with branches, if you specify a release, level and branch, **get** retrieves the highest sequence in that branch.

-xsid-list

Exclude the indicated deltas from the retrieved version. The excluded deltas are noted in the standard output message. *sid-list* is a comma-separated list of SIDs. To specify a range of deltas, use a '-' separator instead of a comma, between two SIDs in the list.

USAGE**ID Keywords**

In the absence of **-e** or **-k**, **get** expands the following ID keywords by replacing them with the indicated values in the text of the retrieved source.

<i>Keyword</i>	<i>Value</i>
%A%	Shorthand notation for an ID line with data for what(1) : %Z% %Y% %M% %I% %Z%
%B%	SID branch component
%C%	Current line number. Intended for identifying messages output by the program such as “ <i>this shouldn't have happened</i> ” type errors. It is <i>not</i> intended to be used on every line to provide sequence numbers.
%D%	Current date: <i>yy/mm/dd</i>
%E%	Date newest applied delta was created: <i>yy/mm/dd</i>
%F%	SCCS <i>s.file</i> name
%G%	Date newest applied delta was created: <i>mm/dd/yy</i>
%H%	Current date: <i>mm/dd/yy</i>
%I%	SID of the retrieved version: %R%.%L%.%B%.%S%
%L%	SID level component
%M%	Module name: either the value of the m flag in the <i>s.file</i> (see sccs-admin(1)), or the name of the <i>s.file</i> less the prefix
%P%	Fully qualified <i>s.file</i> name
%Q%	Value of the q flag in the <i>s.file</i>
%R%	SID Release component
%S%	SID Sequence component
%T%	Current time: <i>hh:mm:ss</i>
%U%	Time the newest applied delta was created: <i>hh:mm:ss</i>
%W%	Shorthand notation for an ID line with data for what : %Z% %M% %I%
%Y%	Module type: value of the t flag in the <i>s.file</i>
%Z%	4-character string: '@(#)', recognized by what .

ID String

The table below explains how the SCCS identification string is determined for retrieving and creating deltas.

Determination of SCCS Identification String				
SID* Specified	-b Option Used†	Other Conditions	SID Retrieved	SID of Delta to be Created
none‡	no	R defaults to mR	mR.mL	mR.(mL+1)
none‡	yes	R defaults to mR	mR.mL	mR.mL.(mB+1).1
R	no	R > mR	mR.mL	R.1***
R	no	R = mR	mR.mL	mR.(mL+1)
R	yes	R > mR	mR.mL	mR.mL.(mB+1).1
R	yes	R = mR	mR.mL	mR.mL.(mB+1).1
R	-	R < mR and R does <i>not</i> exist	hR.mL**	hR.mL.(mB+1).1
R	-	Trunk succ.# in release > R and R exists	R.mL	R.mL.(mB+1).1
R.L	no	No trunk succ.	R.L	R.(L+1)
R.L	yes	No trunk succ.	R.L	R.L.(mB+1).1
R.L	-	Trunk succ. in release ≥ R	R.L	R.L.(mB+1).1
R.L.B	no	No branch succ.	R.L.B.mS	R.L.B.(mS+1)
R.L.B	yes	No branch succ.	R.L.B.mS	R.L.(mB+1).1
R.L.B.S	no	No branch succ.	R.L.B.S	R.L.B.(S+1)
R.L.B.S	yes	No branch succ.	R.L.B.S	R.L.(mB+1).1
R.L.B.S	-	Branch succ.	R.L.B.S	R.L.(mB+1).1

- * 'R', 'L', 'B', and 'S' are the 'release', 'level', 'branch', and 'sequence' components of the SID, respectively; 'm' means 'maximum'. Thus, for example, 'R.mL' means 'the maximum level number within release R'; 'R.L.(mB+1).1' means 'the first sequence number on the *new* branch (that is, maximum branch number plus one) of level L within release R'. Note: if the SID specified is of the form 'R.L', 'R.L.B', or 'R.L.B.S', each of the specified components *must* exist.
- ** 'hR' is the highest *existing* release that is lower than the specified, *nonexistent*, release R.
- *** Forces creation of the *first* delta in a *new* release.
- # Successor.
- † The -b option is effective only if the b flag is present in the file. An entry of '-' means 'irrelevant'.
- ‡ This case applies if the d (default SID) flag is *not* present in the file. If the d flag is present in the file, the SID obtained from the d flag is interpreted as if it had been specified on the command line. Thus, one of the other cases in this table applies.

FILES	“g-file” l.file p.file z.file	version retrieved by get file containing extracted delta table info permissions (lock) file temporary copy of s.file
SEE ALSO	sccs(1) , sccs-admin(1) , sccs-delta(1) , sccs-help(1) , sccs-prs(1) , sccs-prt(1) , sccs-sact(1) , sccs-unget(1) , what(1) , sccsfile(4) <i>Programming Utilities Guide</i>	
DIAGNOSTICS	Use the SCCS help command for explanations (see sccs-help(1)).	
BUGS	If the effective user has write permission (either explicitly or implicitly) in the directory containing the SCCS files, but the real user does not, only one file may be named when using -e .	

NAME	sccs-help, help – ask for help regarding SCCS error or warning messages
SYNOPSIS	/usr/ccs/bin/help [<i>argument</i>] ...
DESCRIPTION	<p>help retrieves information to further explain errors messages and warnings from SCCS commands. It also provides some information about SCCS command usage. If no arguments are given, help prompts for one.</p> <p>An <i>argument</i> may be a message number (which normally appears in parentheses following each SCCS error or warning message), or an SCCS command name. help responds with an explanation of the message or a usage line for the command.</p> <p>When all else fails, try '/usr/ccs/bin/help stuck'.</p>
FILES	/usr/ccs/lib/help directory containing files of message text
SEE ALSO	sccs(1), sccs-admin(1), sccs-cdc(1), sccs-comb(1), sccs-delta(1), sccs-get(1), sccs-prs(1), sccs-prt(1), sccs-rmdel(1), sccs-sact(1), sccs-sccsdiff(1), sccs-unget(1), sccs-val(1), what(1), sccsfile(4)

NAME	sccs-prs, prs – display selected portions of an SCCS history
SYNOPSIS	<code>/usr/ccs/bin/prs [-ael] [-cdate-time] [-ddataspec] [-rsid] s.filename ...</code>
DESCRIPTION	<p>prs displays part or all of the SCCS file (see sccsfile(4)) in a user supplied format. If a directory name is used in place of the <i>s.filename</i> argument, the prs command applies to all <i>s.files</i> in that directory. Unreadable <i>s.files</i> produce an error; processing continues with the next file (if any). The use of ‘-’ as the <i>s.filename</i> argument indicates that the names of files are to be read from the standard input, one <i>s.file</i> per line.</p>
OPTIONS	<p>In the absence of options, prs displays the delta table (version log). In the absence of -d, or -l, prs displays the entry for each delta indicated by the other options.</p> <p>-a Include all deltas, including those marked as removed (see sccs-rmdel(1)).</p> <p>-e Request information for all deltas created <i>earlier</i> than, and including, the delta indicated with -r or -c.</p> <p>-l Request information for all deltas created <i>later</i> than, and including, the delta indicated with -r or -c.</p> <p>-cdate-time Display information on the latest delta checked in prior to the date and time indicated by the <i>date-time</i> argument. <i>date-time</i> takes the form: <code>yy[mm[dd[hh[mm[ss]]]]]</code>. Units omitted from the indicated date and time default to their maximum possible values; that is -c7502 is equivalent to -c750228235959. Any number of non-numeric characters may separate the various 2 digit components. If white-space characters occur, the <i>date-time</i> specification must be quoted.</p> <p>-ddataspec Produce a report according to the indicated data specification. <i>dataspec</i> consists of a (quoted) text string that includes embedded data keywords of the form: ‘:key:’ (see <i>Data Keywords</i>, below). prs expands these keywords in the output it produces. To specify a TAB character in the output, use <code>\t</code>; to specify a NEWLINE in the output, use <code>\n</code>.</p> <p>-rsid Specify the SCCS delta ID (SID) of the delta for which information is desired. If no SID is specified, the most recently created delta is used.</p>

USAGE
Data Keywords

Data keywords specify which parts of an SCCS file are to be retrieved. All parts of an SCCS file (see **sccsfile(4)**) have an associated data keyword. A data keyword may appear any number of times in a data specification argument to **-d**. These data keywords are listed in the table below:

<i>Keyword</i>	<i>Data Item</i>	<i>File Section*</i>	<i>Value</i>	<i>Format**</i>
:A:	a format for the what string:	N/A	:Z::Y: :M: :I::Z:	S
:B:	branch number	D	<i>nnnn</i>	S
:BD:	body	B	<i>text</i>	M
:BF:	branch flag	F	yes or no	S
:CB:	ceiling boundary	F	:R:	S
:C:	comments for delta	D	<i>text</i>	M
:D:	date delta created	D	:Dy:/:Dm:/:Dd:	S
:Dd:	day delta created	D	<i>nn</i>	S
:Dg:	deltas ignored (seq #)	D	:DS: :DS: . . .	S
:DI:	seq-no. of deltas included, excluded, ignored	D	:Dn:/:Dx:/:Dg:	S
:DL:	delta line statistics	D	:Li:/:Ld:/:Lu:	S
:Dm:	month delta created	D	<i>nn</i>	S
:Dn:	deltas included (seq #)	D	:DS: :DS: . . .	S
:DP:	predecessor delta seq-no.	D	<i>nnnn</i>	S
:Ds:	default SID	F	:I:	S
:DS:	delta sequence number	D	<i>nnnn</i>	S
:Dt:	delta information	D	:DT: :I: :D: :T: :P: :DS: :DP:	S
:DT:	delta type	D	D or R	S
:Dx:	deltas excluded (seq #)	D	:DS: . . .	S
:Dy:	year delta created	D	<i>nn</i>	S
:F:	s.file name	N/A	<i>text</i>	S
:FB:	floor boundary	F	:R:	S
:FD:	file descriptive text	C	<i>text</i>	M
:FL:	flag list	F	<i>text</i>	M
:GB:	gotten body	B	<i>text</i>	M
:I:	SCCS delta ID (SID)	D	:R::L::B::S:	S
:J:	joint edit flag	F	yes or no	S
:KF:	keyword error/warning flag	F	yes or no	S
:L:	level number	D	<i>nnnn</i>	S
:Ld:	lines deleted by delta	D	<i>nnnnn</i>	S
:Li:	lines inserted by delta	D	<i>nnnnn</i>	S
:LK:	locked releases	F	:R: . . .	S
:Lu:	lines unchanged by delta	D	<i>nnnnn</i>	S
:M:	module name	F	<i>text</i>	S
:MF:	MR validation flag	F	yes or no	S
:MP:	MR validation program	F	<i>text</i>	S
:MR:	MR numbers for delta	D	<i>text</i>	M
:ND:	null delta flag	F	yes or no	S
:Q:	user defined keyword	F	<i>text</i>	S

:P:	user who created delta	D	<i>username</i>	S
:PN:	s.file's pathname	N/A	<i>text</i>	S
:R:	release number	D	<i>nnnn</i>	S
:S:	sequence number	D	<i>nnnn</i>	S
:T:	time delta created	D	:Th::Tm::Ts:	S
:Th:	hour delta created	D	<i>nn</i>	S
:Tm:	minutes delta created	D	<i>nn</i>	S
:Ts:	seconds delta created	D	<i>nn</i>	S
:UN:	user names	U	<i>text</i>	M
:W:	a form of what string	N/A	:Z::M:\t:I:	S
:Y:	module type flag	F	<i>text</i>	S
:Z:	what string delimiter	N/A	@(#)	S

*B = body, D = delta table, F = flags, U = user names

**S = simple format, M = multi-line format

EXAMPLES

The following command:

```
example% /usr/ccs/bin/prs -e -d":I:\t:P:" program.c
```

produces:

```
1.6 username  
1.5 username  
...
```

FILES

/tmp/pr????? temporary file

SEE ALSO

sccs(1), sccs-cdc(1), sccs-delta(1), sccs-get(1), sccs-help(1), sccs-prt(1), sccs-sact(1), sccs-sccsdiff(1), what(1), sccsfile(4)

Programming Utilities Guide

DIAGNOSTICS

Use the SCCS **help** command for explanations (see **sccs-help(1)**).

NAME	sccs-prt, prt – display delta table information from an SCCS file
SYNOPSIS	<code>/usr/ccs/bin/prt [-abdefistu] [-cdate-time] [-rdate-time] [-ysid] s.filename ...</code>
DESCRIPTION	<p>prt prints selected portions of an SCCS file. By default, it prints the delta table (version log).</p> <p>If a directory name is used in place of the <i>s.filename</i> argument, the prt command applies to all <i>s.files</i> in that directory. Unreadable <i>s.files</i> produce an error; processing continues with the next file (if any). The use of ‘-’ as the <i>s.filename</i> argument indicates that the names of files are to be read from the standard input, one <i>s.file</i> per line.</p>
OPTIONS	<p>If any option other than -y, -c, or -r is supplied, the name of each file being processed (preceded by one NEWLINE and followed by two NEWLINE characters) appears above its contents.</p> <p>If none of the -u, -f, -t, or -b options are used, -d is assumed. -s, -i are mutually exclusive, as are -c and -r.</p> <p>-a Display log entries for all deltas, including those marked as removed.</p> <p>-b Print the body of the <i>s.file</i>.</p> <p>-d Print delta table entries. This is the default.</p> <p>-e Everything. This option implies -d, -i, -u, -f, and -t.</p> <p>-f Print the flags of each named <i>s.file</i>.</p> <p>-i Print the serial numbers of included, excluded, and ignored deltas.</p> <p>-s Print only the first line of the delta table entries; that is, only up to the statistics.</p> <p>-t Print the descriptive text contained in the <i>s.file</i>.</p> <p>-u Print the user-names and/or numerical group IDs of users allowed to make deltas.</p> <p>-cdate-time Exclude delta table entries that are specified cutoff date and time. Each entry is printed as a single line, preceded by the name of the SCCS file. This format (also produced by -r, and -y) makes it easy to sort multiple delta tables in chronological order. When both -y and -c, or -y and -r are supplied, prt stops printing when the first of the two conditions is met.</p> <p>-rdate-time Exclude delta table entries that are newer than the specified cutoff date and time.</p> <p>-ysid Exclude delta table entries made prior to the SID specified. If no delta in the table has the specified SID, the entire table is printed. If no SID is specified, the most recent delta is printed.</p>

USAGE
Output Format

The following format is used to print those portions of the **s**.file that are specified by the various options.

- NEWLINE
- Type of delta (**D** or **R**)
- SPACE
- SCCS delta ID (SID)
- TAB
- Date and time of creation in the form: *yy/mm/dd hh/mm/ss*
- SPACE
- Username the delta's creator
- TAB
- Serial number of the delta
- SPACE
- Predecessor delta's serial number
- TAB
- Line-by-line change statistics in the form: *inserted/deleted/unchanged*
- NEWLINE
- List of included deltas, followed by a NEWLINE (only if there were any such deltas and the **-i** options was used)
- List of excluded deltas, followed by a NEWLINE (only if there were any such deltas and the **-i** options was used)
- List of ignored deltas, followed by a NEWLINE (only if there were any such deltas and the **-i** options was used)
- List of modification requests (MR s), followed by a NEWLINE (only if any MR numbers were supplied).
- Lines of the delta commentary (if any), followed by a NEWLINE.

EXAMPLES

The following command:

```
example% /usr/ccs/bin/prt -y program.c
```

produces a one-line display of the delta table entry for the most recent version:

```
s.program.c: D 1.6 88/07/06 21:39:39 username 5 4 00159/00080/00636...
```

SEE ALSO

sccs(1), **sccs-cdc(1)**, **sccs-delta(1)**, **sccs-get(1)**, **sccs-help(1)**, **sccs-prs(1)**, **sccs-sact(1)**, **sccs-sccsdiff(1)**, **what(1)**, **sccsfile(4)**

DIAGNOSTICS

Use the SCCS **help** command for explanations (see **sccs-help(1)**).

NAME	sccs-rmdel, rmdel – remove a delta from an SCCS file						
SYNOPSIS	<code>/usr/ccs/bin/rmdel -rsid s.filename . . .</code>						
DESCRIPTION	<p>rmdel removes the delta specified by the SCCS delta ID (SID) supplied with -r. The delta to be removed must be the most recent (leaf) delta in its branch. In addition, the SID must <i>not</i> be that of a version checked out for editing: it must not appear in any entry of the version lock file (p.file).</p> <p>If you created the delta, or, if you own the file and directory and have write permission, you can remove it with rmdel.</p> <p>If a directory name is used in place of the <i>s.filename</i> argument, the rmdel command applies to all s.files in that directory. Unreadable s.files produce an error; processing continues with the next file (if any). The use of '-' as the <i>s.filename</i> argument indicates that the names of files are to be read from the standard input, one s.file per line.</p>						
OPTIONS	-rsid Remove the version corresponding to the indicated SID (delta).						
FILES	<table border="0"> <tr> <td>p.file</td> <td>permissions file</td> </tr> <tr> <td>s.file</td> <td>history file</td> </tr> <tr> <td>z.file</td> <td>temporary copy of the s.file</td> </tr> </table>	p.file	permissions file	s.file	history file	z.file	temporary copy of the s.file
p.file	permissions file						
s.file	history file						
z.file	temporary copy of the s.file						
SEE ALSO	<p>sccs(1), sccs-admin(1), sccs-cdc(1), sccs-comb(1), sccs-delta(1), sccs-help(1), sccs-prs(1), sccs-prt(1), sccs-sccsdiff(1), sccs-unget(1), what(1), sccsfile(4)</p> <p><i>Programming Utilities Guide</i></p>						
DIAGNOSTICS	Use the SCCS help command for explanations (see sccs-help(1)).						

NAME	sccs-sact, sact – show editing activity status of an SCCS file
SYNOPSIS	<code>/usr/ccs/bin/sact s.filename ...</code>
DESCRIPTION	<p>sact informs the user of any SCCS files that are checked out for editing.</p> <p>The output for each named file consists of five fields separated by SPACE characters.</p> <ul style="list-style-type: none">• SID of a delta that currently exists in the SCCS file, to which changes will be made to make the new delta• SID for the new delta to be created• Username of the person who has the file checked out for editing.• Date that the version was checked out.• Time that the version was checked out. <p>If a directory name is used in place of the <i>s.filename</i> argument, the sact command applies to all s.files in that directory. Unreadable s.files produce an error; processing continues with the next file (if any). The use of '-' as the <i>s.filename</i> argument indicates that the names of files are to be read from the standard input, one s.file per line.</p>
SEE ALSO	sccs(1) , sccs-delta(1) , sccs-get(1) , sccs-help(1) , sccs-prs(1) , sccs-prt(1) , what(1) , sccsfile(4) <i>Programming Utilities Guide</i>
DIAGNOSTICS	Use the SCCS help command for explanations (see sccs-help(1)).
BUGS	sact is not recognized as a subcommand of sccs(1) .

NAME	sccs-sccsdiff, sccsdiff – compare two versions of an SCCS file
SYNOPSIS	<i>/usr/ccs/bin/sccsdiff</i> [-p] -rsid <i>-rsid</i> [<i>diff-options</i>] <i>s.filename</i>
DESCRIPTION	sccsdiff compares two versions of an SCCS file and displays the differences between the two versions. Any number of SCCS files may be specified; the options specified apply to all named <i>s.files</i> .
OPTIONS	-p Pipe output for each file through pr(1) . -rsid Specify a version corresponding to the indicated SCCS delta ID (SID) for comparison. Versions are passed to diff(1) in the order given. <i>diff-options</i> Pass options to diff(1) , including: -c , -e , -f , -h , -b and -D .
FILES	<i>/tmp/get?????</i> temporary files
SEE ALSO	diff(1) , sccs(1) , sccs-delta(1) , sccs-get(1) , sccs-help(1) , sccs-prs(1) , sccs-prt(1) , what(1) , sccsfile(4) <i>Programming Utilities Guide</i>
DIAGNOSTICS	<i>filename</i> : No differences If the two versions are the same. Use the SCCS help command for explanations of other messages (see sccs-help(1)).

NAME	sccs-unget, unget – undo a previous get of an SCCS file
SYNOPSIS	<code>/usr/ccs/bin/unget [-ns] [-rsid] s.filename ...</code>
DESCRIPTION	<p>unget undoes the effect of a get -e done prior to the creation of the pending delta. If a directory name is used in place of the <i>s.filename</i> argument, the unget command applies to all <i>s.files</i> in that directory. Unreadable <i>s.files</i> produce an error; processing continues with the next file (if any). The use of - as the <i>s.filename</i> argument indicates that the names of files are to be read from the standard input, one <i>s.file</i> per line.</p>
OPTIONS	<p>-n Retain the retrieved version, which is otherwise removed.</p> <p>-s Suppress display of the SCCS delta ID (SID).</p> <p>-rsid When multiple versions are checked out, specify which pending delta to abort. A diagnostic results if the specified SID is ambiguous, or if it is necessary but omitted from the command line.</p>
SEE ALSO	sccs(1) , sccs-delta(1) , sccs-get(1) , sccs-help(1) , sccs-prs(1) , sccs-prt(1) , sccs-rmdel(1) , sccs-sact(1) , sccs-sccsdiff(1) , what(1) , sccsfile(4)
DIAGNOSTICS	Use the SCCS help command for explanations (see sccs-help(1)).

NAME	sccs-val, val – validate an SCCS file
SYNOPSIS	<code>/usr/ccs/bin/val –</code> <code>/usr/ccs/bin/val [–s] [–m name] [–rsid] [–y type] s.filename ...</code>
DESCRIPTION	<p>val determines if the specified s.files files meet the characteristics specified by the indicated arguments. val can process up to 50 files on a single command line.</p> <p>val has a special argument, ‘–’, which reads the standard input until the end-of-file condition is detected. Each line read is independently processed as if it were a command line argument list.</p> <p>val generates diagnostic messages on the standard output for each command line and file processed and also returns a single 8-bit code upon exit as described below.</p> <p>The 8-bit code returned by val is a disjunction of the possible errors, that is, it can be interpreted as a bit string where (moving from left to right) the bits set are interpreted as follows:</p> <ul style="list-style-type: none"> bit 0 = missing file argument bit 1 = unknown or duplicate option bit 2 = corrupted s.file bit 3 = can not open file or file not in s.file format bit 4 = the SCCS delta ID (SID) is invalid or ambiguous bit 5 = the SID does not exist bit 6 = mismatch between %Y% and –y argument bit 7 = mismatch between %M% –m argument <p>val can process two or more files on a given command line, and in turn can process multiple command lines (when reading the standard input). In these cases, an aggregate code is returned which is the logical OR of the codes generated for each command line and file processed.</p>
OPTIONS	<p>–s Silent. Suppress the normal error or warning messages.</p> <p>–m name Compare <i>name</i> with the %M% ID keyword in the s.file.</p> <p>–rsid Check to see if the indicated SID is ambiguous, invalid, or absent from the s.file.</p> <p>–y type Compare <i>type</i> with the %Y% ID keyword.</p>
SEE ALSO	sccs(1) , sccs-admin(1) , sccs-delta(1) , sccs-get(1) , sccs-help(1) , what(1) , sccsfile(4) <i>Programming Utilities Guide</i>
DIAGNOSTICS	Use the SCCS help command for explanations (see sccs-help(1)).

NAME	sccs – front end for the Source Code Control System (SCCS)
SYNOPSIS	sccs [-r] [-drootprefix] [-psubdir] <i>subcommand</i> [<i>option ...</i>] [<i>filename ...</i>]
DESCRIPTION	<p>The sccs command is a comprehensive, straightforward front end to the various utility programs of the Source Code Control System (SCCS).</p> <p>sccs applies the indicated <i>subcommand</i> to the history file associated with each of the indicated files.</p> <p>The name of an SCCS history file is derived by prepending the ‘s.’ prefix to the filename of a working copy. The sccs command normally expects these ‘s.files’ to reside in an SCCS subdirectory. Thus, when you supply sccs with a <i>filename</i> argument, it normally applies the subcommand to a file named s.filename in the SCCS subdirectory. If <i>filename</i> is a pathname, sccs looks for the history file in the SCCS subdirectory of that file’s parent directory. If <i>filename</i> is a directory, however, sccs applies the subcommand to every s.file file it contains. Thus, the command:</p> <p style="padding-left: 40px;">example% sccs get program.c</p> <p>would apply the get subcommand to a history file named:</p> <p style="padding-left: 40px;">SCCS/s.program.c</p> <p>while the command:</p> <p style="padding-left: 40px;">example% sccs get SCCS</p> <p>would apply it to every s.file in the SCCS subdirectory.</p> <p>Options for the sccs command itself must appear before the <i>subcommand</i> argument. Options for a given subcommand must appear after the <i>subcommand</i> argument. These options are specific to each subcommand, and are described along with the subcommands themselves (see Subcommands, below).</p>
Running Setuid	The sccs command also includes the capability to run “setuid” to provide additional protection. However this does not apply to subcommands such as sccs-admin (1), since this would allow anyone to change the authorizations of the history file. Commands that would do so always run as the real user.
OPTIONS	<p>-r Run sccs with the real user ID, rather than set to the effective user ID.</p> <p>-drootprefix Define the root portion of the pathname for SCCS history files. The default root portion is the current directory. Note: <i>rootprefix</i> is prepended to the entire <i>filename</i> argument, even if <i>filename</i> is an absolute pathname. -d overrides any directory specified by the PROJECTDIR environment variable (see ENVIRONMENT, below).</p> <p>-psubdir Define the (sub)directory within which a history file is expected to reside. SCCS is the default. (See EXAMPLES, below).</p>

USAGE

Subcommands

Many of the following **sccs** subcommands invoke programs that reside in **/usr/ccs/bin**. Many of these subcommands accept additional arguments that are documented in the reference page for the utility program the subcommand invokes.

admin Modify the flags or checksum of an SCCS history file. Refer to **sccs-admin(1)** for more information about the **admin** utility. While **admin** can be used to initialize a history file, you may find that the **create** subcommand is simpler to use for this purpose.

cdc -rsid [-y[comment]]

Annotate (change) the delta commentary. Refer to **sccs-cdc(1)**. Note: The **fix** subcommand can be used to replace the delta, rather than merely annotating the existing commentary.

-rsid Specify the SCCS delta ID (SID) to which the change notation is to be added. The SID for a given delta is a number, in Dewey decimal format, composed of two or four fields: the *release* and *level* fields, and for branch deltas, the *branch* and *sequence* fields. For instance, the SID for the initial delta is normally **1.1**.

-y[comment]

Specify the comment with which to annotate the delta commentary. If **-y** is omitted, **sccs** prompts for a comment. A null *comment* results in an empty annotation.

check [-b] [-u[username]]

Check for files currently being edited. Like **info** and **tell**, but returns an exit code, rather than producing a listing of files. **check** returns a non-zero exit status if anything is being edited.

-b Ignore branches.

-u[username]

Only check files being edited by you. When *username* is specified, only check files being edited by that user.

clean [-b]

Remove everything in the current directory that can be retrieved from an SCCS history. Does not remove files that are being edited.

-b Do not check branches to see if they are being edited. '**clean -b**' is dangerous when branch versions are kept in the same directory.

comb Generate scripts to combine deltas. Refer to **sccs-comb(1)**.

create Create (initialize) history files. **create** performs the following steps:

- Renames the original source file to **,program.c** in the current directory.
- Create the history file called **s.program.c** in the SCCS subdirectory.
- Performs an '**sccs get**' on **program.c** to retrieve a read-only copy of the initial version.

deledit [-s] [-y[*comment*]]

Equivalent to an '**sccs delta**' and then an '**sccs edit**'. **deledit** checks in a delta, and checks the file back out again, but leaves the current working copy of the file intact.

-s Silent. Do not report delta numbers or statistics.

-y[*comment*]

Supply a comment for the delta commentary. If -y is omitted, **delta** prompts for a comment. A NULL *comment* results in an empty comment field for the delta.

delget [-s] [-y[*comment*]]

Perform an '**sccs delta**' and then an '**sccs get**' to check in a delta and retrieve read-only copies of the resulting new version. See the **deledit** subcommand for a description of -s and -y. **sccs** performs a **delta** on all the files specified in the argument list, and then a **get** on all the files. If an error occurs during the **delta**, the **get** is not performed.

delta [-s] [-y[*comment*]]

Check in pending changes. Records the line-by-line changes introduced while the file was checked out. The effective user ID must be the same as the ID of the person who has the file checked out. Refer to **sccs-delta(1)**. See the **deledit** subcommand for a description of -s and -y.

diffs [-C] [-c*date-time*] [-rsid] *diff-options*

Compare (in **diff(1)** format) the working copy of a file that is checked out for editing, with a version from the SCCS history. Use the most recent checked-in version by default. The **diffs** subcommand accepts the same options as **diff**, with the exception that the -c option to **diff** must be specified as -C.

-C Pass the -c option to **diff**.

-c*date-time*

Use the most recent version checked in before the indicated date and time for comparison. *date-time* takes the form: *yy[mm[dd[hh[mm[ss]]]]]*. Omitted units default to their maximum possible values; that is -c7502 is equivalent to -c750228235959.

-rsid Use the version corresponding to the indicated delta for comparison.

edit Retrieve a version of the file for editing. '**sccs edit**' extracts a version of the file that is writable by you, and creates a **p**.file in the SCCS subdirectory as lock on the history, so that no one else can check that version in or out. ID keywords are retrieved in unexpanded form. **edit** accepts the same options as **get**, below.

enter Similar to **create**, but omits the final '**sccs get**'. This may be used if an '**sccs edit**' is to be performed immediately after the history file is initialized.

fix -rsid

Revise a (leaf) delta. Remove the indicated delta from the SCCS history, but leave a working copy of the current version in the directory. This is useful for incorporating trivial updates for which no audit record is needed, or for revising the delta commentary. **fix** must be followed by a **-r** option, to specify the SID of the delta to remove. The indicated delta must be the most recent (leaf) delta in its branch. Use **fix** with caution since it does not leave an audit trail of differences (although the previous commentary is retained within the history file).

get [-ekmps] [-cdate-time] [-rsid]

Retrieve a version from the SCCS history. By default, this is a read-only working copy of the most recent version; ID keywords are in expanded form. Refer to **sccs-get(1)**.

- e** Retrieve a version for editing. Same as **sccs edit**.
- k** Retrieve a writable copy but do not check out the file. ID keywords are unexpanded.
- m** Precede each line with the SID of the delta in which it was added.
- p** Produce the retrieved version on the standard output. Reports that would normally go to the standard output (delta ID's and statistics) are directed to the standard error.
- s** Silent. Do not report version numbers or statistics.
- cdate-time**
Retrieve the latest version checked in prior to the date and time indicated by the *date-time* argument. *date-time* takes the form: *yy[mm[dd[hh[mm[ss]]]]]*.
- rsid** Retrieve the version corresponding to the indicated SID.

help message-code | sccs-command**help stuck**

Supply more information about SCCS diagnostics. **help** displays a brief explanation of the error when you supply the code displayed by an SCCS diagnostic message. If you supply the name of an SCCS command, it prints a usage line. **help** also recognizes the keyword **stuck**. Refer to **sccs-help(1)**.

info [-b] [-u[username]]

Display a list of files being edited, including the version number checked out, the version to be checked in, the name of the user who holds the lock, and the date and time the file was checked out.

- b** Ignore branches.
- u[username]**
Only list files checked out by you. When *username* is specified, only list files checked out by that user.

print Print the entire history of each named file. Equivalent to an **'sccs prs -e'** followed by an **'sccs get -p -m'**.

prs [-el] [-cdate-time] [-rsid]

Peruse (display) the delta table, or other portion of an s.file. Refer to **sccs-prs**(1).

-e Display delta table information for all deltas earlier than the one specified with -r (or all deltas if none is specified).

-l Display information for all deltas later than, and including, that specified by -c or -r.

-cdate-time

Specify the latest delta checked in before the indicated date and time. The *date-time* argument takes the form:
yy[mm[dd[hh[mm[ss]]]]].

-rsid Specify a given delta by SID.

prt [-y] Display the delta table, but omit the MR field (see **sccsfile**(4) for more information on this field). Refer to **sccs-prt**(1).

-y Display the most recent delta table entry. The format is a single output line for each filename argument, which is convenient for use in a pipeline with **awk**(1) or **sed**(1).

rmdel -rsid

Remove the indicated delta from the history file. That delta must be the most recent (leaf) delta in its branch. Refer to **sccs-rmdel**(1).

sccsdiff -rold-sid -rnew-sid diff-options

Compare two versions corresponding to the indicated SIDs (deltas) using **diff**. Refer to **sccs-sccsdiff**(1).

tell [-b] [-u[username]]

Display the list of files that are currently checked out, one filename per line.

-b Ignore branches.

-u[username]

Only list files checked out to you. When *username* is specified, only list files check out to that user.

unedit “Undo” the last **edit** or ‘**get -e**’, and return the working copy to its previous condition. **unedit** backs out all pending changes made since the file was checked out.

unget Same as **unedit**. Refer to **sccs-unget**(1).

val Validate the history file. Refer to **sccs-val**(1).

what Display any expanded ID keyword strings contained in a binary (object) or text file. Refer to **what**(1) for more information.

EXAMPLES

sccs converts the command:

```
example% sccs -d/usr/src/include get stdio.h
```

to:

```
/usr/ccs/bin/get /usr/src/include/SCCS/s.stdio.h
```

The command:

```
example% sccs -pprivate get include/stdio.h
```

becomes:

```
/usr/ccs/bin/get include/private/s.stdio.h
```

To initialize the history file for a source file named **program.c**: make the SCCS subdirectory, and then use '**sccs create**':

```
example% mkdir SCCS  
example% sccs create program.c  
program.c:  
1.1  
14 lines
```

After verifying the working copy, you can remove the backup file that starts with a comma:

```
example% diff program.c ,program.c  
example% rm ,program.c
```

To check out a copy of **program.c** for editing, edit it, and then check it back in:

```
example% sccs edit program.c  
1.1  
new delta 1.2  
14 lines  
example% vi program.c  
your editing session  
example% sccs delget program.c  
comments? clarified cryptic diagnostic  
1.2  
3 inserted  
2 deleted  
12 unchanged  
1.2  
15 lines
```

To retrieve a file from another directory into the current directory:

example% sccs get /usr/src/sccs/cc.c

or:

example% sccs -p/usr/src/sccs/ get cc.c

To check out all files under SCCS in the current directory:

example% sccs edit SCCS

To check in all files currently checked out to you:

example% sccs delta `sccs tell -u`

ENVIRONMENT

If the environment variable **PROJECTDIR** is set to contain an absolute pathname (beginning with a slash), **sccs** searches for SCCS history files in the directory given by that variable. If **PROJECTDIR** does not begin with a slash, it is taken as the name of a user, and **sccs** searches the **src** or **source** subdirectory of that user's home directory for history files.

FILES

SCCS	SCCS subdirectory
SCCS/d.file	temporary file of differences
SCCS/p.file	lock (permissions) file for checked-out versions
SCCS/q.file	temporary file
SCCS/s.file	SCCS history file
SCCS/x.file	temporary copy of the s.file
SCCS/z.file	temporary lock file
/usr/ccs/bin/*	SCCS utility programs

SEE ALSO

awk(1), **diff(1)**, **sccs-admin(1)**, **sccs-cdc(1)**, **sccs-comb(1)**, **sccs-delta(1)**, **sccs-get(1)**, **sccs-help(1)**, **sccs-prs(1)**, **sccs-rmdel(1)**, **sccs-sact(1)**, **sccs-sccsdiff(1)**, **sccs-unget(1)**, **sccs-val(1)**, **sed(1)**, **what(1)**, **sccsfile(4)**

Programming Utilities Guide

BUGS

There is no **sact** subcommand to invoke **/usr/ccs/bin/sact** (see **sccs-sact(1)**). However, the **info** subcommand performs an equivalent function.

NAME	script – make record of a terminal session
SYNOPSIS	script [-a] [<i>filename</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	script makes a record of everything printed on your screen. The record is written to <i>filename</i> . If no file name is given, the record is saved in the file typescript . The script command forks and creates a sub-shell, according to the value of \$SHELL , and records the text from this session. The script ends when the forked shell exits or when CTRL-D is typed.
OPTIONS	-a Append the session record to <i>filename</i> , rather than overwrite it.
NOTES	script places <i>everything</i> that appears on the screen in <i>filename</i> , including prompts.

NAME	sdiff – print differences between two files side-by-side																		
SYNOPSIS	sdiff [-l] [-s] [-o output] [-w n] <i>filename1 filename2</i>																		
AVAILABILITY	SUNWesu																		
DESCRIPTION	sdiff uses the output of the diff command to produce a side-by-side listing of two files indicating lines that are different. Lines of the two files are printed with a blank gutter between them if the lines are identical, a < in the gutter if the line appears only in <i>filename1</i> , a > in the gutter if the line appears only in <i>filename2</i> , and a for lines that are different. (See the EXAMPLES section below.)																		
OPTIONS	<p>-l Print only the left side of any lines that are identical.</p> <p>-s Do not print identical lines.</p> <p>-o output Use the argument <i>output</i> as the name of a third file that is created as a user-controlled merge of <i>filename1</i> and <i>filename2</i>. Identical lines of <i>filename1</i> and <i>filename2</i> are copied to <i>output</i>. Sets of differences, as produced by diff, are printed; where a set of differences share a common gutter character. After printing each set of differences, sdiff prompts the user with a % and waits for one of the following user-typed commands:</p> <table border="0"> <tr><td>l</td><td>Append the left column to the output file.</td></tr> <tr><td>r</td><td>Append the right column to the output file.</td></tr> <tr><td>s</td><td>Turn on silent mode; do not print identical lines.</td></tr> <tr><td>v</td><td>Turn off silent mode.</td></tr> <tr><td>e l</td><td>Call the editor with the left column.</td></tr> <tr><td>e r</td><td>Call the editor with the right column.</td></tr> <tr><td>e b</td><td>Call the editor with the concatenation of left and right.</td></tr> <tr><td>e</td><td>Call the editor with a zero length file.</td></tr> <tr><td>q</td><td>Exit from the program.</td></tr> </table> <p>On exit from the editor, the resulting file is concatenated to the end of the <i>output</i> file.</p> <p>-w n Use the argument <i>n</i> as the width of the output line. The default line length is 130 characters.</p>	l	Append the left column to the output file.	r	Append the right column to the output file.	s	Turn on silent mode; do not print identical lines.	v	Turn off silent mode.	e l	Call the editor with the left column.	e r	Call the editor with the right column.	e b	Call the editor with the concatenation of left and right.	e	Call the editor with a zero length file.	q	Exit from the program.
l	Append the left column to the output file.																		
r	Append the right column to the output file.																		
s	Turn on silent mode; do not print identical lines.																		
v	Turn off silent mode.																		
e l	Call the editor with the left column.																		
e r	Call the editor with the right column.																		
e b	Call the editor with the concatenation of left and right.																		
e	Call the editor with a zero length file.																		
q	Exit from the program.																		
EXAMPLES	<p>A sample output of sdiff follows.</p> <table border="0"> <tr><td>x</td><td> </td><td>y</td></tr> <tr><td>a</td><td></td><td>a</td></tr> <tr><td>b</td><td><</td><td></td></tr> <tr><td>c</td><td><</td><td></td></tr> <tr><td>d</td><td></td><td>d</td></tr> <tr><td></td><td>></td><td>c</td></tr> </table>	x		y	a		a	b	<		c	<		d		d		>	c
x		y																	
a		a																	
b	<																		
c	<																		
d		d																	
	>	c																	

ENVIRONMENT

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

LC_CTYPE

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

SEE ALSO

diff(1), **ed(1)**, **environ(5)**

NAME	sed – stream editor
SYNOPSIS	sed [-n] [-e <i>script</i>] [-f <i>sfilename</i>] [<i>filename</i> ...]
AVAILABILITY	SUNWcsu
DESCRIPTION	sed copies the named <i>filename</i> (standard input default) to the standard output, edited according to a script of commands.
OPTIONS	<p>-n Suppress the default output.</p> <p>-e <i>script</i> <i>script</i> is an edit command for sed. See USAGE below for more information on the format of <i>script</i>. If there is just one -e option and no -f options, the flag -e may be omitted.</p> <p>-f <i>sfilename</i> Take the script from <i>sfilename</i>.</p>
USAGE	<p>A script consists of editing commands, one per line, of the following form:</p> <p style="padding-left: 40px;">[<i>address</i> [, <i>address</i>]] <i>function</i> [<i>arguments</i>]</p> <p>In normal operation, sed cyclically copies a line of input into a <i>pattern space</i> (unless there is something left after a D command), applies in sequence all commands whose <i>addresses</i> select that pattern space, and copies the resulting pattern space to the standard output (except under -n).</p> <p>Some of the commands use a <i>hold space</i> to save all or part of the <i>pattern space</i> for subsequent retrieval.</p> <p>An <i>address</i> is either a decimal number that counts input lines cumulatively across files, a \$ that addresses the last line of input, or a context address, that is, a <i>/regular expression/</i> in the style of ed(1) modified thus:</p> <ul style="list-style-type: none"> • In a context address, the construction <i>\?regular expression?</i>, where <i>?</i> is any character, is identical to <i>/regular expression/</i>. Note: In the context address <i>\xabc\xdefx</i>, the second <i>x</i> stands for itself, so that the regular expression is abcxdef. • The escape sequence <i>\n</i> matches a new-line <i>embedded</i> in the pattern space. • A period (.) matches any character except the <i>terminal</i> new-line of the pattern space. • A command line with no addresses selects every line. • A command line with one address selects each line that matches the address. • A command line with two addresses selects the inclusive range from the first pattern space that matches the first address through the next pattern space that matches the second address. Thereafter the process is repeated, looking again for the first address. (If the second address is a number less than or equal to the line number selected by the first address, only the line corresponding to the first address is selected.)

Editing commands can be applied only to non-selected pattern spaces by use of the negation function ! (described below).

In the following list of functions the maximum number of permissible addresses for each function is indicated.

The *text* argument consists of one or more lines, all but the last of which end with \ to hide the new-line. Backslashes in text are treated like backslashes in the replacement string of an *s* command, and may be used to protect initial blanks and tabs against the stripping that is done on every script line. The *rfile* or *wfile* argument must terminate the command line and must be preceded by exactly one blank. Each *wfile* is created before processing begins. There can be at most 10 distinct *wfile* arguments.

The following table lists the functions.

Maximum Number of Addresses	Command	Description
1	a \ <i>text</i>	Append. Place <i>text</i> on the output before reading the next input line.
2	b <i>label</i>	Branch to the : command bearing the <i>label</i> . If <i>label</i> is empty, branch to the end of the script.
2	c \ <i>text</i>	Change. Delete the pattern space. Place <i>text</i> on the output. Start the next cycle.
2	d	Delete the pattern space. Start the next cycle.
2	D	Delete the initial segment of the pattern space through the first new-line. Start the next cycle. (See the N command below.)
2	g	Replace the contents of the pattern space by the contents of the hold space.
2	G	Append the contents of the hold space to the pattern space.
2	h	Replace the contents of the hold space by the contents of the pattern space.
2	H	Append the contents of the pattern space to the hold space.
1	i \ <i>text</i>	Insert. Place <i>text</i> on the standard output.
2	l	List the pattern space on the standard output in an unambiguous form. Non-printable characters are displayed in octal notation and long lines are folded.
2	n	Copy the pattern space to the standard output. Replace the pattern space with the next line of input.
2	N	Append the next line of input to the pattern space with an embedded new-line. (The current line number changes.)
2	p	Print. Copy the pattern space to the standard output.
2	P	Copy the initial segment of the pattern space through the first new-line to the standard output.
1	q	Quit. Branch to the end of the script. Do not start a new cycle.
2	r <i>rfile</i>	Read the contents of <i>rfile</i> . Place them on the output before reading the next input line.

Maximum Number of Addresses	Command	Description
2	t <i>label</i>	Test. Branch to the : command bearing the <i>label</i> if any substitutions have been made since the most recent reading of an input line or execution of a t . If <i>label</i> is empty, branch to the end of the script.
2	w <i>wfile</i>	Write. Append the pattern space to <i>wfile</i> . The first occurrence of w will cause <i>wfile</i> to be cleared. Subsequent invocations of w will append. Each time the sed command is used, <i>wfile</i> is overwritten.
2	x	Exchange the contents of the pattern and hold spaces.
2	! <i>function</i>	Don't. Apply the <i>function</i> (or group, if <i>function</i> is {) only to lines <i>not</i> selected by the address(es).
0	: <i>label</i>	This command does nothing; it bears a <i>label</i> for b and t commands to branch to.
1	=	Place the current line number on the standard output as a line.
2	{	Execute the following commands through a matching } only when the pattern space is selected.
0		An empty command is ignored.
0	#	If a # appears as the first character on a line of a script file, then that entire line is treated as a comment, with one exception: if a # appears on the first line and the character after the # is an n , then the default output will be suppressed. The rest of the line after #n is also ignored. A script file must contain at least one non-comment line.

Maximum Number of Addresses	Command (Using <i>strings</i>) and Description
2	<p>s/ <i>regular expression/ replacement/ flags</i></p> <p>Substitute the <i>replacement</i> string for instances of the <i>regular expression</i> in the pattern space. Any character may be used instead of /. For a fuller description see ed (1). <i>flags</i> is zero or more of:</p> <p><i>n n= 1 - 512</i>. Substitute for just the <i>nth</i> occurrence of the <i>regular expression</i>.</p> <p>g Global. Substitute for all nonoverlapping instances of the <i>regular expression</i> rather than just the first one.</p> <p>p Print the pattern space if a replacement was made.</p> <p>w <i>wfile</i> Write. Append the pattern space to <i>wfile</i> if a replacement was made.</p>
2	<p>y/ <i>string1 / string2 /</i></p> <p>Transform. Replace all occurrences of characters in <i>string1</i> with the corresponding characters in <i>string2</i>. <i>string1</i> and <i>string2</i> must have the same number of characters. For example, y/abc/ABC/ replaces a with A, b with B, and c with C.</p>

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

awk(1), **ed(1)**, **grep(1)**, **environ(5)**

NAME	set, unset, setenv, unsetenv, export – shell built-in functions to determine the characteristics for environmental variables of the current shell and its descendents
SYNOPSIS	
sh	set [—aefhkntuvx [<i>argument ...</i>]] unset [<i>name ...</i>] export [<i>name ...</i>]
csh	set [<i>var</i> [= <i>value</i>]] set <i>var</i> [<i>n</i>] = <i>word</i> unset <i>pattern</i> setenv [<i>VAR</i> [<i>word</i>]] unsetenv <i>variable</i>
ksh	set [±aefhkmpstuvx] [±o option]... [±A name] [<i>arg ...</i>] unset [–f] <i>name ...</i> †† export [<i>name</i> [= <i>value</i>]] ...
DESCRIPTION	
sh	<p>The set built-in command has the following options:</p> <ul style="list-style-type: none"> –a Mark variables which are modified or created for export. –e Exit immediately if a command exits with a non-zero exit status. –f Disable file name generation. –h Locate and remember function commands as functions are defined (function commands are normally located when the function is executed). –k All keyword arguments are placed in the environment for a command, not just those that precede the command name. –n Read commands but do not execute them. –t Exit after reading and executing one command. –u Treat unset variables as an error when substituting. –v Print shell input lines as they are read. –x Print commands and their arguments as they are executed. –– Do not change any of the flags; useful in setting \$1 to –. <p>Using + rather than – causes these flags to be turned off. These flags can also be used upon invocation of the shell. The current set of flags may be found in \$–. The remaining arguments are positional parameters and are assigned, in order, to \$1, \$2, ... If no arguments are given the values of all names are printed.</p>

For each *name*, **unset** removes the corresponding variable or function value. The variables **PATH**, **PS1**, **PS2**, **MAILCHECK**, and **IFS** cannot be unset.

With the **export** built-in, the given *names* are marked for automatic export to the *environment* of subsequently executed commands. If no arguments are given, variable names that have been marked for export during the current shell's execution are listed. (Variable names exported from a parent shell are listed only if they have been exported again during the current shell's execution.) Function names are *not* exported.

csh With no arguments, **set** displays the values of all shell variables. Multiword values are displayed as a parenthesized list. With the *var* argument alone, **set** assigns an empty (null) value to the variable *var*. With arguments of the form *var = value* **set** assigns *value* to *var*, where *value* is one of:

<i>word</i>	A single word (or quoted string).
(<i>wordlist</i>)	A space-separated list of words enclosed in parentheses.

Values are command and filename expanded before being assigned. The form **set var[n]** = *word* replaces the *n*'th word in a multiword value with *word*.

unset removes variables whose names match (filename substitution) *pattern*. All variables are removed by '**unset ***'; this has noticeably distasteful side effects.

With no arguments, **setenv** displays all environment variables. With the *VAR* argument, **setenv** sets the environment variable *VAR* to have an empty (null) value. (By convention, environment variables are normally given upper-case names.) With both *VAR* and *word* arguments, **setenv** sets the environment variable **NAME** to the value *word*, which must be either a single word or a quoted string. The most commonly used environment variables, **USER**, **TERM**, and **PATH**, are automatically imported to and exported from the **cs**h variables **user**, **term**, and **path**; there is no need to use **setenv** for these. In addition, the shell sets the **PWD** environment variable from the **cs**h variable **cwd** whenever the latter changes.

The environment variables **LC_CTYPE**, **LC_MESSAGES**, **LC_TIME**, **LC_COLLATE**, **LC_NUMERIC**, and **LC_MONETARY** take immediate effect when changed within the C shell.

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S./English).

LC_NUMERIC

Determines the value of the radix character (decimal point (".") in the "C" locale) and thousand separator (empty string ("") in the "C" locale).

unsetenv removes *variable* from the environment. As with **unset**, pattern matching is not performed.

ksh The flags for the **set** built-in have meaning as follows:

- A** Array assignment. Unset the variable *name* and assign values sequentially from the list *arg*. If **+A** is used, the variable *name* is not unset first.
- a** All subsequent variables that are defined are automatically exported.
- e** If a command has a non-zero exit status, execute the **ERR** trap, if set, and exit. This mode is disabled while reading profiles.
- f** Disables file name generation.
- h** Each command becomes a tracked alias when first encountered.
- k** All variable assignment arguments are placed in the environment for a command, not just those that precede the command name.
- m** Background jobs will run in a separate process group and a line will print upon completion. The exit status of background jobs is reported in a completion message. On systems with job control, this flag is turned on automatically for interactive shells.
- n** Read commands and check them for syntax errors, but do not execute them. Ignored for interactive shells.
- o** The following argument can be one of the following option names:
 - allexport** Same as **-a**.
 - errexit** Same as **-e**.
 - bgnice** All background jobs are run at a lower priority. This is the default mode. **emacs** Puts you in an **emacs** style in-line editor for command entry.
 - gmacs** Puts you in a **gmacs** style in-line editor for command entry.
 - ignoreeof** The shell will not exit on end-of-file. The command **exit** must be used.
 - keyword** Same as **-k**.
 - markdirs**

All directory names resulting from file name generation have a trailing / appended.

monitor Same as **-m**.

noclobber

Prevents redirection > from truncating existing files. Require >| to truncate a file when turned on.

noexec Same as **-n**.

noglob Same as **-f**.

nolog Do not save function definitions in history file.

nounset Same as **-u**.

privileged

Same as **-p**.

verbose Same as **-v**.

trackall Same as **-h**.

vi Puts you in insert mode of a **vi** style in-line editor until you hit escape character **033**. This puts you in control mode. A return sends the line.

viraw Each character is processed as it is typed in **vi** mode.

xtrace Same as **-x**.

If no option name is supplied then the current option settings are printed.

- p** Disables processing of the `$HOME/.profile` file and uses the file `/etc/suid_profile` instead of the `ENV` file. This mode is on whenever the effective uid is not equal to the real uid, or when the effective gid is not equal to the real gid. Turning this off causes the effective uid and gid to be set to the real uid and gid.
- s** Sort the positional parameters lexicographically.
- t** Exit after reading and executing one command.
- u** Treat unset parameters as an error when substituting.
- v** Print shell input lines as they are read.
- x** Print commands and their arguments as they are executed.
- Turns off **-x** and **-v** flags and stops examining arguments for flags.
- Do not change any of the flags; useful in setting `$1` to a value beginning with **-**. If no arguments follow this flag then the positional parameters are unset.

Using **+** rather than **-** causes these flags to be turned off. These flags can also be used upon invocation of the shell. The current set of flags may be found in `$-`. Unless **-A** is specified, the remaining arguments are positional parameters and are assigned, in order, to `$1 $2 ...`. If no arguments are given then the names and values of all variables are printed on the standard output.

The variables given by the list of *names* are unassigned, i.e., their values and attributes are erased. **readonly** variables cannot be unset. If the **-f** flag is set, then the names refer to *function* names. Unsetting `ERRNO`, `LINENO`, `MAILCHECK`, `OPTARG`, `OPTIND`, `RANDOM`,

SECONDS, **TMOUT**, and **_** removes their special meaning even if they are subsequently assigned.

When using **unset**, the variables given by the list of *names* are unassigned, i.e., their values and attributes are erased. **readonly** variables cannot be unset. If the **-f** flag is set, then the names refer to *function* names. Unsetting **ERRNO**, **LINENO**, **MAILCHECK**, **OPTARG**, **OPTIND**, **RANDOM**, **SECONDS**, **TMOUT**, and **_** removes their special meaning even if they are subsequently assigned.

With the **export** built-in, the given *names* are marked for automatic export to the **environment** of subsequently-executed commands.

On this man page, **ksh(1)** commands that are preceded by one or two † (daggers) are treated specially in the following ways:

1. Variable assignment lists preceding the command remain in effect when the command completes.
2. I/O redirections are processed after variable assignments.
3. Errors cause a script that contains them to abort.
4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.

SEE ALSO **csh(1)**, **ksh(1)**, **read(1)**, **sh(1)**, **typeset(1)**

NAME	set, unset – set and unset local or global environment variables
SYNOPSIS	<pre>set [-l variable[=value]] ... set [-e variable[=value]] ... set [-f file variable[=value]] ... unset -l variable ... unset -f file variable ...</pre>
DESCRIPTION	<p>The set command sets <i>variable</i> in the environment, or adds <i>variable=value</i> to <i>file</i>. If <i>variable</i> is not equated to a value, set expects the value to be on <i>stdin</i>. The unset command removes <i>variable</i>. Note that the FMLI predefined, read-only variables (such as ARG1), may not be set or unset.</p> <p>Note that at least one of the above options must be used for each variable being set or unset. If you set a variable with the <i>-filename</i> option, you must thereafter include <i>filename</i> in references to that variable. For example, <code>\${(file) VARIABLE}</code>.</p> <p>FMLI inherits the UNIX environment when invoked.</p>
OPTIONS	<p>-l Sets or unsets the specified variable in the local environment. Variables set with -l will not be inherited by processes invoked from FMLI.</p> <p>-e Sets the specified variable in the UNIX environment. Variables set with -e will be inherited by any processes started from FMLI. Note that these variables cannot be unset.</p> <p>-f file Sets or unsets the specified variable in the global environment. The argument <i>file</i> is the name, or pathname, of a file containing lines of the form <i>variable=value</i>. <i>file</i> will be created if it does not already exist. Note that no space intervenes between -f and <i>file</i>.</p>
EXAMPLE	<p>Storing a selection made in a menu:</p> <pre>name=Selection 2 action=`set -l SELECTION=2` close</pre>
NOTES	<p>Variables set to be available to the UNIX environment (those set using the -e option) can only be set for the current <i>fmli</i> process and the processes it calls.</p> <p>When using the -f option, unless <i>file</i> is unique to the process, other users of FMLI on the same machine will be able to expand these variables, depending on the read/write permissions on <i>file</i>.</p> <p>A variable set in one frame may be referenced or unset in any other frame. This includes local variables.</p>
SEE ALSO	env(1) , sh(1)

NAME	setcolor – redefine or create a color
SYNOPSIS	setcolor <i>color red_level green_level blue_level</i>
DESCRIPTION	The setcolor command takes four arguments: <i>color</i> , which must be a string naming the color; and the arguments <i>red_level</i> , <i>green_level</i> , and <i>blue_level</i> , which must be integer values defining, respectively, the intensity of the red, green, and blue components of <i>color</i> . Intensities must be in the range of 0 to 1000. If you are redefining an existing color, you must use its current name (default color names are: black , blue , green , cyan , red , magenta , yellow , and white). setcolor returns the color's name string.
EXAMPLES	The following is an example of the arguments that setcolor takes: <pre>`setcolor blue 100 24 300`</pre>

NAME	sh, jsh – shell: the standard shell, and job control shell -- command interpreters
SYNOPSIS	sh [-acefhiknprstuvx] [<i>argument...</i>] jsh [-acefhiknprstuvx] [<i>argument...</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	sh is a command programming language that executes commands read from a terminal or a file. The command jsh is an interface to the shell which provides all of the functionality of sh and enables Job Control (see “Job Control,” below). See “Invocation,” below for the meaning of arguments to the shell.
Definitions	A <i>blank</i> is a tab or a space. A <i>name</i> is a sequence of ASCII letters, digits, or underscores, beginning with a letter or an underscore. A <i>parameter</i> is a name, a digit, or any of the characters *, @, #, ?, -, \$, and !\^.
USAGE Commands	<p>A <i>simple-command</i> is a sequence of non-blank <i>words</i> separated by <i>blanks</i>. The first <i>word</i> specifies the name of the command to be executed. Except as specified below, the remaining <i>words</i> are passed as arguments to the invoked command. The command name is passed as argument 0 (see exec(2)). The <i>value</i> of a <i>simple-command</i> is its exit status if it terminates normally, or (octal) 200+status if it terminates abnormally; see signal(5) for a list of status values.</p> <p>A <i>pipeline</i> is a sequence of one or more <i>commands</i> separated by . The standard output of each <i>command</i> but the last is connected by a pipe(2) to the standard input of the next <i>command</i>. Each <i>command</i> is run as a separate process; the shell waits for the last <i>command</i> to terminate. The exit status of a <i>pipeline</i> is the exit status of the last command in the <i>pipeline</i>.</p> <p>A <i>list</i> is a sequence of one or more <i>pipelines</i> separated by ;, &, &&, or , and optionally terminated by ; or &. Of these four symbols, ; and & have equal precedence, which is lower than that of && and . The symbols && and also have equal precedence. A semicolon (;) causes sequential execution of the preceding <i>pipeline</i> (that is, the shell waits for the <i>pipeline</i> to finish before executing any commands following the semicolon); an ampersand (&) causes asynchronous execution of the preceding pipeline (that is, the shell does <i>not</i> wait for that pipeline to finish). The symbol && () causes the <i>list</i> following it to be executed only if the preceding pipeline returns a zero (non-zero) exit status. An arbitrary number of newlines may appear in a <i>list</i>, instead of semicolons, to delimit commands.</p> <p>A <i>command</i> is either a <i>simple-command</i> or one of the following. Unless otherwise stated, the value returned by a command is that of the last <i>simple-command</i> executed in the command.</p>

for *name* [**in** *word* ...] **do** *list* **done**

Each time a **for** command is executed, *name* is set to the next *word* taken from the **in** *word* list. If **in** *word* ... is omitted, then the **for** command executes the **do** *list* once for each positional parameter that is set (see “Parameter Substitution,” below). Execution ends when there are no more words in the list.

case *word* **in** [*pattern* [| *pattern*] ...) *list* ;;] ... **esac**

A **case** command executes the *list* associated with the first *pattern* that matches *word*. The form of the patterns is the same as that used for file-name generation (see “File Name Generation”) except that a slash, a leading dot, or a dot immediately following a slash need not be matched explicitly.

if *list*; **then** *list*; [**elif** *list*; **then** *list*] ... [**else** *list*]; **fi**

The *list* following **if** is executed and, if it returns a zero exit status, the *list* following the first **then** is executed. Otherwise, the *list* following **elif** is executed and, if its value is zero, the *list* following the next **then** is executed. Failing that, the **else** *list* is executed. If no **else** *list* or **then** *list* is executed, then the **if** command returns a zero exit status.

while *list* **do** *list* **done**

A **while** command repeatedly executes the **while** *list* and, if the exit status of the last command in the *list* is zero, executes the **do** *list*; otherwise the loop terminates. If no commands in the **do** *list* are executed, then the **while** command returns a zero exit status; **until** may be used in place of **while** to negate the loop termination test.

(list)

Execute *list* in a sub-shell.

{ *list*;

list is executed in the current (that is, parent) shell. The { must be followed by a space.

name () { *list*;

Define a function which is referenced by *name*. The body of the function is the *list* of commands between { and }. The { must be followed by a space. Execution of functions is described below (see “Execution”). The { and } are unnecessary if the body of the function is a *command* as defined above, under “Commands.”

The following words are only recognized as the first word of a command and when not quoted:

if then else elif fi case esac for while until do done { }

Comments Lines

A word beginning with # causes that word and all the following characters up to a new-line to be ignored.

**Command
Substitution**

The shell reads commands from the string between two grave accents (` `) and the standard output from these commands may be used as all or part of a word. Trailing new-lines from the standard output are removed.

No interpretation is done on the string before the string is read, except to remove backslashes (\) used to escape other characters. Backslashes may be used to escape a grave accent (`) or another backslash (\) and are removed before the command string is read. Escaping grave accents allows nested command substitution. If the command substitution lies within a pair of double quotes ("... '... '..."), a backslash used to escape a double quote (\") will be removed; otherwise, it will be left intact.

If a backslash is used to escape a newline character (**newline**), both the backslash and the newline are removed (see the later section on "Quoting"). In addition, backslashes used to escape dollar signs (\\$) are removed. Since no parameter substitution is done on the command string before it is read, inserting a backslash to escape a dollar sign has no effect. Backslashes that precede characters other than \, ', ", **newline**, and \$ are left intact when the command string is read.

Parameter Substitution

The character \$ is used to introduce substitutable *parameters*. There are two types of parameters, positional and keyword. If *parameter* is a digit, it is a positional parameter. Positional parameters may be assigned values by **set**. Keyword parameters (also known as variables) may be assigned values by writing:

```
name=value [ name=value ] ...
```

Pattern-matching is not performed on *value*. There cannot be a function and a variable with the same *name*.

\${parameter}

The value, if any, of the parameter is substituted. The braces are required only when *parameter* is followed by a letter, digit, or underscore that is not to be interpreted as part of its name. If *parameter* is * or @, all the positional parameters, starting with \$1, are substituted (separated by spaces). Parameter \$0 is set from argument zero when the shell is invoked.

\${parameter:-word}

If *parameter* is set and is non-null, substitute its value; otherwise substitute *word*.

\${parameter:=word}

If *parameter* is not set or is null set it to *word*; the value of the parameter is substituted. Positional parameters may not be assigned in this way.

\${parameter:?word}

If *parameter* is set and is non-null, substitute its value; otherwise, print *word* and exit from the shell. If *word* is omitted, the message "parameter null or not set" is printed.

\${parameter:+word}

If *parameter* is set and is non-null, substitute *word*; otherwise substitute nothing.

In the above, *word* is not evaluated unless it is to be used as the substituted string, so that, in the following example, **pwd** is executed only if **d** is not set or is null:

```
echo ${d:-'pwd'}
```

If the colon (:) is omitted from the above expressions, the shell only checks whether *parameter* is set or not.

The following parameters are automatically set by the shell.

- # The number of positional parameters in decimal.
- Flags supplied to the shell on invocation or by the **set** command.
- ? The decimal value returned by the last synchronously executed command.
- \$ The process number of this shell.
- ! The process number of the last background command invoked.

The following parameters are used by the shell. The parameters in this section are also referred to as environment variables.

HOME The default argument (home directory) for the **cd** command, set to the user's login directory by **login(1)** from the password file (see **passwd(4)**).

PATH The search path for commands (see "Execution," below).

CDPATH

The search path for the **cd** command.

MAIL If this parameter is set to the name of a mail file *and* the **MAILPATH** parameter is not set, the shell informs the user of the arrival of mail in the specified file.

MAILCHECK

This parameter specifies how often (in seconds) the shell will check for the arrival of mail in the files specified by the **MAILPATH** or **MAIL** parameters. The default value is **600** seconds (10 minutes). If set to **0**, the shell will check before each prompt.

MAILPATH

A colon (:) separated list of file names. If this parameter is set, the shell informs the user of the arrival of mail in any of the specified files. Each file name can be followed by % and a message that will be printed when the modification time changes. The default message is **you have mail**.

PS1 Primary prompt string, by default "\$ ".

PS2 Secondary prompt string, by default "> ".

IFS Internal field separators, normally **space**, **tab**, and **newline** (see "Blank Interpretation").

SHACCT

If this parameter is set to the name of a file writable by the user, the shell will write an accounting record in the file for each shell procedure executed.

SHELL When the shell is invoked, it scans the environment (see "Environment," below) for this name.

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

If **LC_CTYPE** and **LC_MESSAGES** (see **environ(5)**) are not set in the environment, the operational behavior of the shell for each corresponding locale category is determined by the value of the **LANG** environment variable. If **LC_ALL** is set, its contents are used to override both the **LANG** and the other **LC_*** variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how the shell behaves.

The shell gives default values to **PATH**, **PS1**, **PS2**, **MAILCHECK**, and **IFS**. **HOME** and **MAIL** are set by **login(1)**.

Blank Interpretation

After parameter and command substitution, the results of substitution are scanned for internal field separator characters (those found in **IFS**) and split into distinct arguments where such characters are found. Explicit null arguments (" " or '') are retained. Implicit null arguments (those resulting from *parameters* that have no values) are removed.

Input/Output Redirection

A command's input and output may be redirected using a special notation interpreted by the shell. The following may appear anywhere in a *simple-command* or may precede or follow a *command* and are *not* passed on as arguments to the invoked command. Note: Parameter and command substitution occurs before *word* or *digit* is used.

<*word* Use file *word* as standard input (file descriptor 0).

>*word* Use file *word* as standard output (file descriptor 1). If the file does not exist, it is created; otherwise, it is truncated to zero length.

<code>>>word</code>	Use file <i>word</i> as standard output. If the file exists, output is appended to it (by first seeking to the EOF); otherwise, the file is created.
<code><<[-]word</code>	After parameter and command substitution is done on <i>word</i> , the shell input is read up to the first line that literally matches the resulting <i>word</i> , or to an EOF. If, however, <code>-</code> is appended to <code><<</code> : <ol style="list-style-type: none"> 1) leading tabs are stripped from <i>word</i> before the shell input is read (but after parameter and command substitution is done on <i>word</i>), 2) leading tabs are stripped from the shell input as it is read and before each line is compared with <i>word</i>, and 3) shell input is read up to the first line that literally matches the resulting <i>word</i>, or to an EOF. <p>If any character of <i>word</i> is quoted (see “Quoting,” later), no additional processing is done to the shell input. If no characters of <i>word</i> are quoted:</p> <ol style="list-style-type: none"> 1) parameter and command substitution occurs, 2) (escaped) <code>\newlines</code> are removed, and 3) <code>\</code> must be used to quote the characters <code>\</code>, <code>\$</code>, and <code>'</code>. <p>The resulting document becomes the standard input.</p>
<code><&digit</code>	Use the file associated with file descriptor <i>digit</i> as standard input. Similarly for the standard output using <code>>&digit</code> .
<code><&-</code>	The standard input is closed. Similarly for the standard output using <code>>&-</code> .

If any of the above is preceded by a digit, the file descriptor which will be associated with the file is that specified by the digit (instead of the default 0 or 1). For example:

```
... 2>&1
```

associates file descriptor 2 with the file currently associated with file descriptor 1.

The order in which redirections are specified is significant. The shell evaluates redirections left-to-right. For example:

```
... 1>xxx 2>&1
```

first associates file descriptor 1 with file *xxx*. It associates file descriptor 2 with the file associated with file descriptor 1 (that is, *xxx*). If the order of redirections were reversed, file descriptor 2 would be associated with the terminal (assuming file descriptor 1 had been) and file descriptor 1 would be associated with file *xxx*.

Using the terminology introduced on the first page, under “Commands,” if a *command* is composed of several *simple commands*, redirection will be evaluated for the entire *command* before it is evaluated for each *simple command*. That is, the shell evaluates redirection for the entire *list*, then each *pipeline* within the *list*, then each *command* within each *pipeline*, then each *list* within each *command*.

If a command is followed by `&` the default standard input for the command is the empty file `/dev/null`. Otherwise, the environment for the execution of a command contains the file descriptors of the invoking shell as modified by input/output specifications.

File Name Generation	<p>Before a command is executed, each command <i>word</i> is scanned for the characters *, ?, and [. If one of these characters appears the word is regarded as a <i>pattern</i>. The word is replaced with alphabetically sorted file names that match the pattern. If no file name is found that matches the pattern, the word is left unchanged. The character . at the start of a file name or immediately following a /, as well as the character / itself, must be matched explicitly.</p> <ul style="list-style-type: none"> * Matches any string, including the null string. ? Matches any single character. [...] Matches any one of the enclosed characters. A pair of characters separated by – matches any character lexically between the pair, inclusive. If the first character following the opening [is a !, any character not enclosed is matched. <p>Note: All quoted characters (see below) must be matched explicitly in a filename.</p>
Quoting	<p>The following characters have a special meaning to the shell and cause termination of a word unless quoted:</p> <p style="text-align: center;">; & () ^ < > newline space tab</p> <p>A character may be <i>quoted</i> (that is, made to stand for itself) by preceding it with a backslash (\) or inserting it between a pair of quote marks (‘ ’ or " "). During processing, the shell may quote certain characters to prevent them from taking on a special meaning. Backslashes used to quote a single character are removed from the word before the command is executed. The pair \bnewline is removed from a word before command and parameter substitution.</p> <p>All characters enclosed between a pair of single quote marks (‘ ’), except a single quote, are quoted by the shell. Backslash has no special meaning inside a pair of single quotes. A single quote may be quoted inside a pair of double quote marks (for example, " ’ "), but a single quote can not be quoted inside a pair of single quotes.</p> <p>Inside a pair of double quote marks (" "), parameter and command substitution occurs and the shell quotes the results to avoid blank interpretation and file name generation. If \$* is within a pair of double quotes, the positional parameters are substituted and quoted, separated by quoted spaces (" \$1 \$2 ..."); however, if \$@ is within a pair of double quotes, the positional parameters are substituted and quoted, separated by unquoted spaces (" \$1" "\$2" ...). \ quotes the characters \, ‘, , and \$. The pair \bnewline is removed before parameter and command substitution. If a backslash precedes characters other than \, ‘, , \$, and newline, then the backslash itself is quoted by the shell.</p>
Prompting	<p>When used interactively, the shell prompts with the value of PS1 before reading a command. If at any time a newline is typed and further input is needed to complete a command, the secondary prompt (that is, the value of PS2) is issued.</p>
Environment	<p>The <i>environment</i> (see environ(5)) is a list of name-value pairs that is passed to an executed program in the same way as a normal argument list. The shell interacts with the environment in several ways. On invocation, the shell scans the environment and creates a</p>

parameter for each name found, giving it the corresponding value. If the user modifies the value of any of these parameters or creates new parameters, none of these affects the environment unless the **export** command is used to bind the shell's parameter to the environment (see also **set -a**). A parameter may be removed from the environment with the **unset** command. The environment seen by any executed command is thus composed of any unmodified name-value pairs originally inherited by the shell, minus any pairs removed by **unset**, plus any modifications or additions, all of which must be noted in **export** commands.

The environment for any *simple-command* may be augmented by prefixing it with one or more assignments to parameters. Thus:

```
TERM=450 command
and
(export TERM; TERM=450; command)
```

are equivalent as far as the execution of *command* is concerned if *command* is not a Special Command. If *command* is a Special Command, then

```
TERM=450 command
```

will modify the **TERM** variable in the current shell.

If the **-k** flag is set, *all* keyword arguments are placed in the environment, even if they occur after the command name. The following example first prints **a=b c** and **c**:

```
echo a=b c
a=b c
set -k
echo a=b c
c
```

Signals The **INTERRUPT** and **QUIT** signals for an invoked command are ignored if the command is followed by **&**; otherwise signals have the values inherited by the shell from its parent, with the exception of signal 11 (but see also the **trap** command below).

Execution Each time a command is executed, the command substitution, parameter substitution, blank interpretation, input/output redirection, and filename generation listed above are carried out. If the command name matches the name of a defined function, the function is executed in the shell process (note how this differs from the execution of shell procedures). If the command name does not match the name of a defined function, but matches one of the *Special Commands* listed below, it is executed in the shell process. The positional parameters **\$1**, **\$2**, . . . are set to the arguments of the function. If the command name matches neither a *Special Command* nor the name of a defined function, a new process is created and an attempt is made to execute the command via **exec(2)**.

The shell parameter **PATH** defines the search path for the directory containing the command. Alternative directory names are separated by a colon (:). The default path is **/usr/bin**. The current directory is specified by a null path name, which can appear immediately after the equal sign, between two colon delimiters anywhere in the path list, or at the end of the path list. If the command name contains a / the search path is not

used. Otherwise, each directory in the path is searched for an executable file. If the file has execute permission but is not an **a.out** file, it is assumed to be a file containing shell commands. A sub-shell is spawned to read it. A parenthesized command is also executed in a sub-shell.

The location in the search path where a command was found is remembered by the shell (to help avoid unnecessary *execs* later). If the command was found in a relative directory, its location must be re-determined whenever the current directory changes. The shell forgets all remembered locations whenever the **PATH** variable is changed or the **hash -r** command is executed (see below).

Special Commands

Input/output redirection is now permitted for these commands. File descriptor 1 is the default output location. When Job Control is enabled, additional *Special Commands* are added to the shell's environment (see "Job Control").

: No effect; the command does nothing. A zero exit code is returned.

. filename

Read and execute commands from *filename* and return. The search path specified by **PATH** is used to find the directory containing *filename*.

break [n]

Exit from the enclosing **for** or **while** loop, if any. If *n* is specified, break *n* levels.

continue [n]

Resume the next iteration of the enclosing **for** or **while** loop. If *n* is specified, resume at the *n*-th enclosing loop.

cd [argument]

Change the current directory to *argument*. The shell parameter **HOME** is the default *argument*. The shell parameter **CDPATH** defines the search path for the directory containing *argument*. Alternative directory names are separated by a colon (:). The default path is **<null>** (specifying the current directory). Note: The current directory is specified by a null path name, which can appear immediately after the equal sign or between the colon delimiters anywhere else in the path list. If *argument* begins with a / the search path is not used. Otherwise, each directory in the path is searched for *argument*.

echo [argument ...]

Echo arguments. See **echo(1)** for usage and description.

eval [argument ...]

The arguments are read as input to the shell and the resulting command(s) executed.

exec [argument ...]

The command specified by the arguments is executed in place of this shell without creating a new process. Input/output arguments may appear and, if no other arguments are given, cause the shell input/output to be modified.

exit [n]

Causes a shell to exit with the exit status specified by *n*. If *n* is omitted the exit status is that of the last command executed (an EOF will also cause the shell to

exit.)

export [*name* ...]

The given *names* are marked for automatic export to the *environment* of subsequently executed commands. If no arguments are given, variable names that have been marked for export during the current shell's execution are listed. (Variable names exported from a parent shell are listed only if they have been exported again during the current shell's execution.) Function names are *not* exported.

getopts

Use in shell scripts to support command syntax standards (see **intro(1)**); it parses positional parameters and checks for legal options. See **getoptcv(1)** for usage and description.

hash [**-r**] [*name* ...]

For each *name*, the location in the search path of the command specified by *name* is determined and remembered by the shell. The **-r** option causes the shell to forget all remembered locations. If no arguments are given, information about remembered commands is presented. *Hits* is the number of times a command has been invoked by the shell process. *Cost* is a measure of the work required to locate a command in the search path. If a command is found in a "relative" directory in the search path, after changing to that directory, the stored location of that command is recalculated. Commands for which this will be done are indicated by an asterisk (*) adjacent to the *hits* information. *Cost* will be incremented when the recalculation is done.

newgrp [*argument*]

Equivalent to **exec newgrp** *argument*. See **newgrp(1M)** for usage and description.

pwd Print the current working directory. See **pwd(1)** for usage and description.

read *name* ...

One line is read from the standard input and, using the internal field separator, **IFS** (normally space or tab), to delimit word boundaries, the first word is assigned to the first *name*, the second word to the second *name*, etc., with leftover words assigned to the last *name*. Lines can be continued using **\newline**. Characters other than **newline** can be quoted by preceding them with a backslash. These backslashes are removed before words are assigned to *names*, and no interpretation is done on the character that follows the backslash. The return code is **0**, unless an EOF is encountered.

readonly [*name* ...]

The given *names* are marked *readonly* and the values of these *names* may not be changed by subsequent assignment. If no arguments are given, a list of all *readonly* names is printed.

return [*n*]

Causes a function to exit with the return value specified by *n*. If *n* is omitted, the return status is that of the last command executed.

set [**—aefhkntuvx** [*argument* ...]]

- a** Mark variables which are modified or created for export.
- e** Exit immediately if a command exits with a non-zero exit status.
- f** Disable file name generation.
- h** Locate and remember function commands as functions are defined (function commands are normally located when the function is executed).
- k** All keyword arguments are placed in the environment for a command, not just those that precede the command name.
- n** Read commands but do not execute them.
- t** Exit after reading and executing one command.
- u** Treat unset variables as an error when substituting.
- v** Print shell input lines as they are read.
- x** Print commands and their arguments as they are executed.
- Do not change any of the flags; useful in setting **\$1** to **-**.

Using **+** rather than **-** causes these flags to be turned off. These flags can also be used upon invocation of the shell. The current set of flags may be found in **\$-**. The remaining arguments are positional parameters and are assigned, in order, to **\$1**, **\$2**, If no arguments are given the values of all names are printed.

shift [*n*]

The positional parameters from **\$n+1** ... are renamed **\$1** If *n* is not given, it is assumed to be 1.

stop *pid* ...

Halt execution of the process number *pid*. (see **ps(1)**).

test

Evaluate conditional expressions. See **test(1)** for usage and description.

times

Print the accumulated user and system times for processes run from the shell.

trap [*argument*] [*n*] ...

The command *argument* is to be read and executed when the shell receives numeric or symbolic signal(s) (*n*). (Note: *argument* is scanned once when the trap is set and once when the trap is taken.) Trap commands are executed in order of signal number or corresponding symbolic names. Any attempt to set a trap on a signal that was ignored on entry to the current shell is ineffective. An attempt to trap on signal 11 (memory fault) produces an error. If *argument* is absent all trap(s) *n* are reset to their original values. If *argument* is the null string this signal is ignored by the shell and by the commands it invokes. If *n* is 0 the command *argument* is executed on exit from the shell. The **trap** command with no arguments prints a list of commands associated with each signal number.

type [*name* ...]

For each *name*, indicate how it would be interpreted if used as a command name.

ulimit [*-[HS][a|cdfnstv]*]

ulimit [*-[HS][c|d|f|n|s|t|v]*] *limit*

ulimit prints or sets hard or soft resource limits. These limits are described in **getrlimit(2)**.

If *limit* is not present, **ulimit** prints the specified limits. Any number of limits may be printed at one time. The *-a* option prints all limits.

If *limit* is present, **ulimit** sets the specified limit to *limit*. The string **unlimited** requests the largest valid limit. Limits may be set for only one resource at a time. Any user may set a soft limit to any value below the hard limit. Any user may lower a hard limit. Only a super-user may raise a hard limit; see **su(1M)**.

The *-H* option specifies a hard limit. The *-S* option specifies a soft limit. If neither option is specified, **ulimit** will set both limits and print the soft limit.

The following options specify the resource whose limits are to be printed or set. If no option is specified, the file size limit is printed or set.

- c* maximum core file size (in 512-byte blocks)
- d* maximum size of data segment or heap (in kbytes)
- f* maximum file size (in 512-byte blocks)
- n* maximum file descriptor plus 1
- s* maximum size of stack segment (in kbytes)
- t* maximum CPU time (in seconds)
- v* maximum size of virtual memory (in kbytes)

umask [*nnn*]

The user file-creation mask is set to *nnn* (see **umask(1)**). If *nnn* is omitted, the current value of the mask is printed.

unset [*name ...*]

For each *name*, remove the corresponding variable or function value. The variables **PATH**, **PS1**, **PS2**, **MAILCHECK**, and **IFS** cannot be unset.

wait [*n*]

Wait for your background process whose process id is *n* and report its termination status. If *n* is omitted, all your shell's currently active background processes are waited for and the return code will be zero.

Invocation

If the shell is invoked through **exec(2)** and the first character of argument zero is *-*, commands are initially read from **/etc/profile** and from **\$HOME/.profile**, if such files exist. Thereafter, commands are read as described below, which is also the case when the shell is invoked as **/usr/bin/sh**. The flags below are interpreted by the shell on invocation only. Note: Unless the *-c* or *-s* flag is specified, the first argument is assumed to be the name of a file containing commands, and the remaining arguments are passed as positional parameters to that command file:

- c** *string* If the **-c** flag is present commands are read from *string*.
- i** If the **-i** flag is present or if the shell input and output are attached to a terminal, this shell is *interactive*. In this case TERMINATE is ignored (so that **kill 0** does not kill an interactive shell) and INTERRUPT is caught and ignored (so that **wait** is interruptible). In all cases, QUIT is ignored by the shell.
- p** If the **-p** flag is present, the shell will not set the effective user and group IDs to the real user and group IDs.
- r** If the **-r** flag is present the shell is a restricted shell (see **rsh(1M)**).
- s** If the **-s** flag is present or if no arguments remain, commands are read from the standard input. Any remaining arguments specify the positional parameters. Shell output (except for *Special Commands*) is written to file descriptor 2.

The remaining flags and arguments are described under the **set** command above.

Job Control (jsh)

When the shell is invoked as **jsh**, Job Control is enabled in addition to all of the functionality described previously for **sh**. Typically Job Control is enabled for the interactive shell only. Non-interactive shells typically do not benefit from the added functionality of Job Control.

With Job Control enabled every command or pipeline the user enters at the terminal is called a *job*. All jobs exist in one of the following states: foreground, background or stopped. These terms are defined as follows: 1) a job in the foreground has read and write access to the controlling terminal; 2) a job in the background is denied read access and has conditional write access to the controlling terminal (see **stty(1)**); 3) a stopped job is a job that has been placed in a suspended state, usually as a result of a **SIGTSTP** signal (see **signal(5)**).

Every job that the shell starts is assigned a positive integer, called a *job number* which is tracked by the shell and will be used as an identifier to indicate a specific job. Additionally the shell keeps track of the *current* and *previous* jobs. The *current job* is the most recent job to be started or restarted. The *previous job* is the first non-current job.

The acceptable syntax for a Job Identifier is of the form:

%jobid

where, *jobid* may be specified in any of the following formats:

- %** or **+** for the current job
- for the previous job
- ?<string>** specify the job for which the command line uniquely contains *string*.
- n** for job number *n*, where *n* is a job number
- pref** where *pref* is a unique prefix of the command name (for example, if the command **ls -l name** were running in the background, it could be referred to as **%ls**); *pref* cannot contain blanks unless it is quoted.

When Job Control is enabled, the following commands are added to the user's environment to manipulate jobs:

bg [%*jobid* ...]

Resumes the execution of a stopped job in the background. If %*jobid* is omitted the current job is assumed.

fg [%*jobid* ...]

Resumes the execution of a stopped job in the foreground, also moves an executing background job into the foreground. If %*jobid* is omitted the current job is assumed.

jobs [-p | -l] [%*jobid* ...]

jobs -x *command* [*arguments*]

Reports all jobs that are stopped or executing in the background. If %*jobid* is omitted, all jobs that are stopped or running in the background will be reported. The following options will modify/enhance the output of **jobs**:

-l Report the process group ID and working directory of the jobs.

-p Report only the process group ID of the jobs.

-x Replace any *jobid* found in *command* or *arguments* with the corresponding process group ID, and then execute *command* passing it *arguments*.

kill [-signal] %*jobid*

Builtin version of **kill** to provide the functionality of the **kill** command for processes identified with a *jobid*.

stop %*jobid* ...

Stops the execution of a background job(s).

suspend

Stops the execution of the current shell (but not if it is the login shell).

wait [%*jobid* ...]

wait builtin accepts a job identifier. If %*jobid* is omitted **wait** behaves as described above under **Special Commands**.

EXIT CODES

Errors detected by the shell, such as syntax errors, cause the shell to return a non-zero exit status. If the shell is being used non-interactively execution of the shell file is abandoned. Otherwise, the shell returns the exit status of the last command executed (see also the **exit** command above).

jsh Only

If the shell is invoked as **jsh** and an attempt is made to exit the shell while there are stopped jobs, the shell issues one warning:

There are stopped jobs.

This is the only message. If another exit attempt is made, and there are still stopped jobs they will be sent a **SIGHUP** signal from the kernel and the shell is exited.

FILES	\$HOME/.profile /dev/null /etc/profile /tmp/sh*
SEE ALSO	intro(1), echo(1), getoptcv(1) login(1), pwd(1), ps(1), shell_builtins(1), stty(1), rsh(1M), newgrp(1M), dup(2), exec(2), fork(2), getrlimit(2), pipe(2), ulimit(2), setlocale(3C), profile(4), passwd(4), environ(5), signal(5)
NOTES	<p>Words used for filenames in input/output redirection are not interpreted for filename generation (see “File Name Generation,” above). For example, cat file1 >a* will create a file named a*.</p> <p>Because commands in pipelines are run as separate processes, variables set in a pipeline have no effect on the parent shell.</p> <p>If you get the error message <i>cannot fork, too many processes</i>, try using the wait(1) command to clean up your background processes. If this doesn't help, the system process table is probably full or you have too many active foreground processes. (There is a limit to the number of process ids associated with your login, and to the number the system can keep track of.)</p> <p>Only the last process in a pipeline can be waited for.</p> <p>If a command is executed, and a command with the same name is installed in a directory in the search path before the directory where the original command was found, the shell will continue to exec the original command. Use the hash command to correct this situation.</p>

NAME	shell – run a command using shell
SYNOPSIS	shell <i>command</i> [<i>command</i>] . . .
DESCRIPTION	The shell function concatenate its arguments, separating each by a space, and passes this string to the shell (\$SHELL if set, otherwise /usr/bin/sh).
EXAMPLES	<p>Since the Form and Menu Language does not directly support background processing, the shell function can be used instead.</p> <pre>shell "build prog > /dev/null &"</pre> <p>If you want the user to continue to be able to interact with the application while the background job is running, the output of an executable run by shell in the background must be redirected: to a file if you want to save the output, or to /dev/null if you don't want to save it (or if there is no output), otherwise your application may appear to be hung until the background job finishes processing.</p> <p>shell can also be used to execute a command that has the same name as an FMLI built-in function.</p>
NOTES	The arguments to shell will be concatenate using spaces, which may or may not do what is expected. The variables set in local environments will not be expanded by the shell because "local" means "local to the current process."
SEE ALSO	sh(1)

NAME	shell_builtins – shell command interpreter built-in functions																																																																
DESCRIPTION	<p>The shell command interpreters (sh(1), csh(1), and ksh(1)), have special built-in functions which are interpreted by the shell as commands. Many of these built-in commands are implemented by more than one of the shells, and some are unique to a particular shell. These are:</p> <table border="0"> <tr> <td>alias</td> <td>bg</td> <td>break</td> <td>case</td> </tr> <tr> <td>cd</td> <td>chdir</td> <td>continue</td> <td>dirs</td> </tr> <tr> <td>eval</td> <td>exec</td> <td>exit</td> <td>export</td> </tr> <tr> <td>fc</td> <td>fg</td> <td>for</td> <td>foreach</td> </tr> <tr> <td>function</td> <td>getopts</td> <td>glob</td> <td>goto</td> </tr> <tr> <td>hash</td> <td>hashstat</td> <td>history</td> <td>if</td> </tr> <tr> <td>jobs</td> <td>kill</td> <td>let</td> <td>limit</td> </tr> <tr> <td>logout</td> <td>newgrp</td> <td>notify</td> <td>onintr</td> </tr> <tr> <td>popd</td> <td>print</td> <td>pushd</td> <td>read</td> </tr> <tr> <td>readonly</td> <td>rehash</td> <td>repeat</td> <td>return</td> </tr> <tr> <td>select</td> <td>set</td> <td>setenv</td> <td>shift</td> </tr> <tr> <td>source</td> <td>stop</td> <td>suspend</td> <td>switch</td> </tr> <tr> <td>test</td> <td>times</td> <td>trap</td> <td>type</td> </tr> <tr> <td>typeset</td> <td>ulimit</td> <td>umask</td> <td>unalias</td> </tr> <tr> <td>unhash</td> <td>unlimit</td> <td>unset</td> <td>unsetenv</td> </tr> <tr> <td>until</td> <td>wait</td> <td>whence</td> <td>while</td> </tr> </table>	alias	bg	break	case	cd	chdir	continue	dirs	eval	exec	exit	export	fc	fg	for	foreach	function	getopts	glob	goto	hash	hashstat	history	if	jobs	kill	let	limit	logout	newgrp	notify	onintr	popd	print	pushd	read	readonly	rehash	repeat	return	select	set	setenv	shift	source	stop	suspend	switch	test	times	trap	type	typeset	ulimit	umask	unalias	unhash	unlimit	unset	unsetenv	until	wait	whence	while
alias	bg	break	case																																																														
cd	chdir	continue	dirs																																																														
eval	exec	exit	export																																																														
fc	fg	for	foreach																																																														
function	getopts	glob	goto																																																														
hash	hashstat	history	if																																																														
jobs	kill	let	limit																																																														
logout	newgrp	notify	onintr																																																														
popd	print	pushd	read																																																														
readonly	rehash	repeat	return																																																														
select	set	setenv	shift																																																														
source	stop	suspend	switch																																																														
test	times	trap	type																																																														
typeset	ulimit	umask	unalias																																																														
unhash	unlimit	unset	unsetenv																																																														
until	wait	whence	while																																																														
Bourne Shell, sh, Special Commands	<p>Input/output redirection is now permitted for these commands. File descriptor 1 is the default output location. When Job Control is enabled, additional <i>Special Commands</i> are added to the shell's environment.</p> <p>Additional to these built-in reserved command words, sh also uses:</p> <p>: No effect; the command does nothing. A zero exit code is returned.</p> <p>. filename Read and execute commands from <i>filename</i> and return. The search path specified by PATH is used to find the directory containing <i>filename</i>.</p>																																																																
C shell, csh	<p>Built-in commands are executed within the C shell. If a built-in command occurs as any component of a pipeline except the last, it is executed in a subshell. Additional to these built-in reserved command words, csh also uses:</p> <p>: Null command. This command is interpreted, but performs no action.</p>																																																																
Korn Shell, ksh, Special Commands	<p>Input/Output redirection is permitted. Unless otherwise indicated, the output is written on file descriptor 1 and the exit status, when there is no syntax error, is zero.</p> <p>Commands that are preceded by one or two † (daggers) are treated specially in the following ways:</p> <ol style="list-style-type: none"> 1. Variable assignment lists preceding the command remain in effect when the command completes. 																																																																

2. I/O redirections are processed after variable assignments.
3. Errors cause a script that contains them to abort.
4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.

Additional to these built-in reserved command words, **ksh** also uses:

† : [*arg* ...]

The command only expands parameters.

† . *file* [*arg* ...]

Read the complete *file* then execute the commands. The commands are executed in the current shell environment. The search path specified by **PATH** is used to find the directory containing *file*. If any arguments *arg* are given, they become the positional parameters. Otherwise the positional parameters are unchanged. The exit status is the exit status of the last command executed. the loop termination test.

SEE ALSO

alias(1), break(1), case(1), cd(1), chmod(1), csh(1), echo(1), exec(1), exit(1), for(1), find(1), function(1), getoptcv(1) getopt(1), glob(1), hash(1), history(1), if(1), intro(1), jobs(1), kill(1), ksh(1), let(1), limit(1), login(1), logout(1), newgrp(1), nice(1), nohup(1), print(1), pwd(1), read(1), readonly(1), repeat(1), set(1), shift(1), sh(1), suspend(1), time(1), times(1), trap(1), typeset(1), umask(1), wait(1), while(1), test(1B), chdir(2), chmod(2), creat(2), umask(2), getopt(3C), profile(4), environ(5)

NAME	shift – shell built-in function to traverse either a shell's argument list or a list of field-separated words
SYNOPSIS	
sh	shift [<i>n</i>]
csh	shift [<i>variable</i>]
ksh	† shift [<i>n</i>]
DESCRIPTION	
sh	The positional parameters from $\$n+1 \dots$ are renamed $\$1 \dots$. If <i>n</i> is not given, it is assumed to be 1.
csh	The components of argv , or <i>variable</i> , if supplied, are shifted to the left, discarding the first component. It is an error for the variable not to be set or to have a null value.
ksh	The positional parameters from $\$n+1 \dots$ are renamed $\$1 \dots$, default <i>n</i> is 1. The parameter <i>n</i> can be any arithmetic expression that evaluates to a non-negative number less than or equal to $\$#$. On this man page, ksh (1) commands that are preceded by one or two † (daggers) are treated specially in the following ways: <ol style="list-style-type: none"> 1. Variable assignment lists preceding the command remain in effect when the command completes. 2. I/O redirections are processed after variable assignments. 3. Errors cause a script that contains them to abort. 4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.
SEE ALSO	csh (1), ksh (1), sh (1)

NAME	shutdown – close down the system at a given time
SYNOPSIS	<code>/usr/ucb/shutdown [-fhknr] time [warning-message ...]</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>shutdown provides an automated procedure to notify users when the system is to be shut down. <i>time</i> specifies when shutdown will bring the system down; it may be the word now (indicating an immediate shutdown), or it may specify a future time in one of two formats: <i>+number</i> and <i>hour:min</i>. The first form brings the system down in <i>number</i> minutes, and the second brings the system down at the time of day indicated in 24-hour notation.</p> <p>At intervals that get closer as the apocalypse approaches, warning messages are displayed at terminals of all logged-in users, and of users who have remote mounts on that machine.</p> <p>At shutdown time a message is written to the system log daemon, syslogd(1M), containing the time of shutdown, the instigator of the shutdown, and the reason. Then a terminate signal is sent to init, which brings the system down to single-user mode.</p>
OPTIONS	<p>As an alternative to the above procedure, these options can be specified:</p> <ul style="list-style-type: none">-f Arrange, in the manner of fastboot(1B), that when the system is rebooted, the file systems will not be checked.-h Execute halt(1M).-k Simulate shutdown of the system. Do not actually shut down the system.-n Prevent the normal sync(2) before stopping.-r Execute reboot(1M).
FILES	<code>/etc/xtab</code> list of remote hosts that have mounted this host
SEE ALSO	login (1), fastboot (1B), halt (1M), reboot (1M), syslogd (1M), sync (2)
NOTES	Only allows you to bring the system down between now and 23:59 if you use the absolute time for shutdown.

NAME	size – print section sizes in bytes of object files
SYNOPSIS	size [-f] [-F] [-n] [-o] [-V] [-x] <i>filename</i> ...
DESCRIPTION	<p>The size command produces segment or section size information in bytes for each loaded section in ELF or COFF object files. size prints out the size of the text, data, and bss (uninitialized data) segments (or sections) and their total.</p> <p>size processes ELF and COFF object files entered on the command line. If an archive file is input to the size command, the information for each object file in the archive is displayed.</p> <p>When calculating segment information, the size command prints out the total file size of the non-writable segments, the total file size of the writable segments, and the total memory size of the writable segments minus the total file size of the writable segments.</p> <p>If it cannot calculate segment information, size calculates section information. When calculating section information, it prints out the total size of sections that are allocatable, non-writable, and not NOBITS, the total size of the sections that are allocatable, writable, and not NOBITS, and the total size of the writable sections of type NOBITS. (NOBITS sections do not actually take up space in the <i>filename</i>.)</p> <p>If size cannot calculate either segment or section information, it prints an error message and stops processing the file.</p>
OPTIONS	<p>-f Print out the size of each allocatable section, the name of the section, and the total of the section sizes. If there is no section data, size prints out an error message and stops processing the file.</p> <p>-F Print out the size of each loadable segment, the permission flags of the segment, then the total of the loadable segment sizes. If there is no segment data, size prints an error message and stops processing the file.</p> <p>-n Print out non-loadable segment or non-allocatable section sizes. If segment data exists, size prints out the memory size of each loadable segment or file size of each non-loadable segment, the permission flags, and the total size of the segments. If there is no segment data, size prints out, for each allocatable and non-allocatable section, the memory size, the section name, and the total size of the sections. If there is no segment or section data, size prints an error message and stops processing.</p> <p>-o Print numbers in octal, not decimal.</p> <p>-V Print the version information for the size command on the standard error output.</p> <p>-x Print numbers in hexadecimal; not decimal.</p>

EXAMPLES The examples below are typical **size** output.

```
example% size filename  
2724 + 88 + 0 = 2812
```

```
example% size -f filename  
26(.text) + 5(.init) + 5(.fini) = 36
```

```
example% size -F filename  
2724(r-x) + 88(rwx) + 0(rwx) = 2812 (If statically linked)
```

SEE ALSO **as(1)**, **cc(1B)**, **ld(1)**, **a.out(4)**, **ar(4)**

NOTES Since the size of bss sections is not known until link-edit time, the **size** command will not give the true total size of pre-linked objects.

NAME	sleep – suspend execution for an interval
SYNOPSIS	sleep <i>time</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>sleep suspends execution for <i>time</i> seconds. It is used to execute a command after a certain amount of time, as in:</p> <pre style="margin-left: 40px;">(sleep 105; command)&</pre> <p>or to execute a command every so often, as in:</p> <pre style="margin-left: 40px;">while true do command sleep 37 done</pre>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of sleep for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how sleep behaves.</p> <p>LC_CTYPE</p> <p>Determines how sleep handles characters. When LC_CTYPE is set to a valid value, sleep can display and handle text and filenames containing valid characters for that locale. sleep can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. sleep can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p> <p>LC_MESSAGES</p> <p>Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).</p>
SEE ALSO	alarm(2) , sleep(3C) , environ(5)

NAME	soelim – resolve and eliminate .so requests from nroff or troff input
SYNOPSIS	soelim [<i>filename ...</i>]
AVAILABILITY	SUNWdoc
DESCRIPTION	<p>soelim reads the specified files or the standard input and performs the textual inclusion implied by the nroff(1) directives of the form</p> <p style="padding-left: 40px;">.so <i>somefile</i></p> <p>when they appear at the beginning of input lines. This is useful since programs such as tbl(1) do not normally do this; it allows the placement of individual tables in separate files to be run as a part of a large document.</p> <p>An argument consisting of '-' is taken to be a file name corresponding to the standard input.</p> <p>Note: Inclusion can be suppressed by using '``' instead of '.', that is,</p> <p style="padding-left: 40px;">`so /usr/share/lib/tmac/tmac.s</p>
EXAMPLES	A sample usage of soelim would be
	example% soelim exum?.n tbl nroff -ms col lpr
SEE ALSO	more (1), nroff (1), tbl (1)

NAME	sort – sort and/or merge files
SYNOPSIS	sort [-cmu] [-ooutput] [-T directory] [-ykmem] [-dfiMnr] [-btx] [+pos1 [-pos2]] [<i>filename</i> ...]
AVAILABILITY	SUNWdoc
DESCRIPTION	<p>The sort command sorts lines of all the named files together and writes the result on the standard output. The standard input is read if '-' is used as a file name or no input-files are named.</p> <p>Comparisons are based on one or more sort keys extracted from each line of input. By default, there is one sort key, the entire input line, and ordering is lexicographic by bytes in machine collating sequence.</p>
OPTIONS	<p>The following options alter the default behavior:</p> <ul style="list-style-type: none"> -c Check that the input-file is sorted according to the ordering rules; give no output unless the file is out of sort. -m Merge only, the input-files are already sorted. -u Unique: suppress all but one in each set of lines having equal keys. -ooutput The argument given is the name of an output-file to use instead of the standard output. This file may be the same as one of the inputs. There may be optional blanks between -o and <i>output</i>. -T directory The <i>directory</i> argument is the name of a directory in which to place temporary files. -ykmem The amount of main memory used by sort has a large impact on its performance. Sorting a small file in a large amount of memory is a waste. If this option is omitted, sort begins using a system default memory size, and continues to use more space as needed. If this option is presented with a value <i>kmem</i>, sort will start using that number of kilobytes of memory, unless the administrative minimum or maximum is violated, in which case the corresponding extremum will be used. Thus, -y0 is guaranteed to start with minimum memory. By convention, -y (with no argument) starts with maximum memory. <p>The following options override the default ordering rules:</p> <ul style="list-style-type: none"> -d “Dictionary” order: only letters, digits, and blanks (spaces and tabs) are significant in comparisons. -f Fold lower-case letters into upper case. -i Ignore non-printable characters.

- M** Compare as months. The first three non-blank characters of the field are folded to upper case and compared. For example, in English the sorting order is "JAN" < "FEB" < ... < "DEC". Invalid fields compare low to "JAN". The **-M** option implies the **-b** option (see below).
- n** An initial numeric string, consisting of optional blanks, optional minus sign, and zero or more digits with optional decimal point, is sorted by arithmetic value. The **-n** option implies the **-b** option (see below). Note: The **-b** option is only effective when restricted sort key specifications are in effect.
- r** Reverse the sense of comparisons.

When ordering options appear before restricted sort key specifications, the requested ordering rules are applied globally to all sort keys. When attached to a specific sort key (described below), the specified ordering options override all global ordering options for that key.

The notation *+pos1 -pos2* restricts a sort key to one beginning at *pos1* and ending just before *pos2*. The characters at position *pos1* and just before *pos2* are included in the sort key (provided that *pos2* does not precede *pos1*). A missing *-pos2* means the end of the line.

Specifying *pos1* and *pos2* involves the notion of a field, a minimal sequence of characters followed by a field separator or a new-line. By default, the first blank (space or tab) of a sequence of blanks acts as the field separator. All blanks in a sequence of blanks are considered to be part of the next field; for example, all blanks at the beginning of a line are considered to be part of the first field. The treatment of field separators can be altered using the options:

- b** Ignore leading blanks when determining the starting and ending positions of a restricted sort key. If the **-b** option is specified before the first *+pos1* argument, it will be applied to all *+pos1* arguments. Otherwise, the **b** flag may be attached independently to each *+pos1* or *-pos2* argument (see below).
- tx** Use *x* as the field separator character; *x* is not considered to be part of a field (although it may be included in a sort key). Each occurrence of *x* is significant (for example, *xx* delimits an empty field).

pos1 and *pos2* each have the form *m.n* optionally followed by one or more of the flags **bdfinr**. A starting position specified by *+m.n* is interpreted to mean the *n*+1st character in the *m*+1st field. A missing *.n* means *.0*, indicating the first character of the *m*+1st field. If the **b** flag is in effect *n* is counted from the first non-blank in the *m*+1st field; *+m.0b* refers to the first non-blank character in the *m*+1st field.

A last position specified by *-m.n* is interpreted to mean the *n*th character (including separators) after the last character of the *m*th field. A missing *.n* means *.0*, indicating the last character of the *m*th field. If the **b** flag is in effect *n* is counted from the last leading blank in the *m*+1st field; *-m.1b* refers to the first non-blank in the *m*+1st field.

When there are multiple sort keys, later keys are compared only after all earlier keys compare equal. Lines that otherwise compare equal are ordered with all bytes significant.

EXAMPLES

Sort the contents of *input-file* with the second field as the sort key:

```
example% sort +1 -2 input-file
```

Sort, in reverse order, the contents of *input-file1* and *input-file2*, placing the output in *output-file* and using the first character of the second field as the sort key:

```
example% sort -r -o output-file +1.0 -1.2 input-file1 input-file2
```

Sort, in reverse order, the contents of *input-file1* and *input-file2* using the first non-blank character of the second field as the sort key:

```
example% sort -r +1.0b -1.1b input-file1 input-file2
```

Print the password file, **passwd(4)**, sorted by the numeric user ID (the third colon-separated field):

```
example% sort -t: +2n -3 /etc/passwd
```

Print the lines of the already sorted file *input-file*, suppressing all but the first occurrence of lines having the same third field (the options **-um** with just one input-file make the choice of a unique representative from a set of equal lines predictable):

```
example% sort -um +2 -3 input-file
```

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

LC_TIME

Determines how **sort** handles date and time formats. In the "C" locale, date and time handling follows the U.S. rules.

FILES	<code>/var/tmp/stm???</code> temporary files
SEE ALSO	<code>comm(1)</code> , <code>join(1)</code> , <code>uniq(1)</code> , <code>environ(5)</code>
DIAGNOSTICS	Comments and exits with non-zero status for various trouble conditions (for example, when input lines are too long), and for disorders discovered under the <code>-c</code> option.
NOTES	When the last line of an input-file is missing a new-line character, sort appends one, prints a warning message, and continues. sort does not guarantee preservation of relative line ordering on equal keys.

NAME	sortbib – sort a bibliographic database
SYNOPSIS	sortbib [<i>-sKEYS</i>] <i>database</i> . . .
AVAILABILITY	SUNWdoc
DESCRIPTION	<p>sortbib sorts files of records containing refer key-letters by user-specified keys. Records may be separated by blank lines, or by '[' and ']' delimiters, but the two styles may not be mixed together. This program reads through each <i>database</i> and pulls out key fields, which are sorted separately. The sorted key fields contain the file pointer, byte offset, and length of corresponding records. These records are delivered using disk seeks and reads, so sortbib may not be used in a pipeline to read standard input.</p> <p>The most common key-letters and their meanings are given below.</p> <ul style="list-style-type: none"> %A Author's name %B Book containing article referenced %C City (place of publication) %D Date of publication %E Editor of book containing article referenced %F Footnote number or label (supplied by refer) %G Government order number %H Header commentary, printed before reference %I Issuer (publisher) %J Journal containing article %K Keywords to use in locating reference %L Label field used by -k option of refer %M Bell Labs Memorandum (undefined) %N Number within volume %O Other commentary, printed at end of reference %P Page number(s) %Q Corporate or Foreign Author (unreversed) %R Report, paper, or thesis (unpublished) %S Series title %T Title of article or book %V Volume number %X Abstract — used by roffbib, not by refer %Y,Z Ignored by refer

By default, **sortbib** alphabetizes by the first %A and the %D fields, which contain the senior author and date.

sortbib sorts on the last word on the %A line, which is assumed to be the author's last name. A word in the final position, such as 'jr.' or 'ed.', will be ignored if the name beforehand ends with a comma. Authors with two-word last names or unusual constructions can be sorted correctly by using the **nroff** convention '\0' in place of a blank. A %Q field is considered to be the same as %A, except sorting begins with the first, not the last, word. **sortbib** sorts on the last word of the %D line, usually the year. It also ignores leading articles (like 'A' or 'The') when sorting by titles in the %T or %J fields; it will ignore articles of any modern European language. If a sort-significant field is absent from a record, **sortbib** places that record before other records containing that field.

No more than 16 databases may be sorted together at one time. Records longer than 4096 characters will be truncated.

- OPTIONS** -sKEYS Specify new *KEYS*. For instance, -sATD will sort by author, title, and date, while -sA+D will sort by all authors, and date. Sort keys past the fourth are not meaningful.
- SEE ALSO** **addbib(1)**, **indxbib(1)**, **lookbib(1)**, **refer(1)**, **roffbib(1)**
- BUGS** Records with missing author fields should probably be sorted by title.

NAME	spell – find spelling errors	
SYNOPSIS	spell [-bvx] [+local_file] [<i>filename</i>] . . .	
AVAILABILITY	SUNWesu	
DESCRIPTION	spell collects words from the named <i>filenames</i> and looks them up in a spelling list. Words that neither occur among nor are derivable (by applying certain inflections, prefixes, and/or suffixes) from words in the spelling list are printed on the standard output. If no <i>filenames</i> are named, words are collected from the standard input. Copies of all output are accumulated in the spellhist file.	
OPTIONS	-b	British spelling is checked. Besides preferring "centre," "colour," "programme," "speciality," "travelled," and so forth, this option insists upon <i>-ise</i> in words like "standardise."
	-v	All words not literally in the spelling list are printed, and plausible derivations from the words in the spelling list are indicated.
	-x	Every plausible stem is displayed, one per line, with = preceding each word.
	+local_file	<i>local_file</i> is the name of a user-provided file that contains a sorted list of words, one per line. With this option, the user can specify a set of words that are correct spellings (in addition to spell 's own spelling list) for each job. Words found in <i>local_file</i> are removed from spell 's output. Use sort(1) to order <i>local_file</i> in ASCII collating sequence. If this ordering is not followed, some entries in <i>local_file</i> may be ignored.
FILES	D_SPELL=/usr/lib/spell/hlist[ab]	hashed spelling lists, American & British
	S_SPELL=/usr/lib/spell/hstop	hashed stop list
	H_SPELL=/var/adm/spellhist	history file
	/usr/share/lib/dict/words	master dictionary
SEE ALSO	sort(1)	
NOTES	Because copies of all output are accumulated in the spellhist file, spellhist may grow quite large and require purging.	

NAME	spline – interpolate smooth curve
SYNOPSIS	spline [-aknpx] . . .
AVAILABILITY	SUNWesu
DESCRIPTION	spline takes pairs of numbers from the standard input as abscissas and ordinates of a function. It produces a similar set, which is approximately equally spaced and includes the input set, on the standard output. The cubic spline output (R. W. Hamming, <i>Numerical Methods for Scientists and Engineers</i> , 2nd ed., 349ff) has two continuous derivatives, and sufficiently many points to look smooth when plotted, for example by graph (1).
OPTIONS	<p>-a Supply abscissas automatically (they are missing from the input); spacing is given by the next argument, or is assumed to be 1 if next argument is not a number.</p> <p>-k The constant k used in the boundary value computation</p> $y_0'' = ky_1'', \quad y_n'' = ky_{n-1}''$ <p>is set by the next argument. By default $k = 0$.</p> <p>-n Space output points so that approximately n intervals occur between the lower and upper x limits. (Default $n = 100$.)</p> <p>-p Make output periodic, that is, match derivatives at ends. First and last input values should normally agree.</p> <p>-x Next 1 (or 2) arguments are lower (and upper) x limits. Normally these limits are calculated from the data. Automatic abscissas start at lower limit (default 0).</p>
SEE ALSO	graph (1) R. W. Hamming, <i>Numerical Methods for Scientists and Engineers</i> , 2nd ed.
DIAGNOSTICS	When data is not strictly monotonic in x , spline reproduces the input without interpolating extra points.
BUGS	A limit of 1000 input points is enforced silently.

NAME	split – split a file into pieces
SYNOPSIS	split [<i>-n</i>] [<i>filename</i> [<i>name</i>]]
AVAILABILITY	SUNWesu
DESCRIPTION	<p>split reads <i>filename</i> and writes it in <i>n</i>-line pieces into a set of output-files. The name of the first output-file is <i>name</i> with aa appended, and so on lexicographically, up to zz (a maximum of 676 files). The maximum length of <i>name</i> is 2 characters less than the maximum filename length allowed by the filesystem. See statvfs(2). If no output name is given, x is used as the default (output-files will be called xaa, xab, and so forth).</p> <p>If no input-file is given, or if - is given in its stead, then the standard input-file is used.</p>
OPTIONS	-n Number of lines in each piece. Defaults to 1000 lines.
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of split for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how split behaves.</p> <p>LC_CTYPE</p> <p>Determines how split handles characters. When LC_CTYPE is set to a valid value, split can display and handle text and filenames containing valid characters for that locale. split can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. split can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p> <p>LC_MESSAGES</p> <p>Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).</p>
SEE ALSO	csplit(1) , statvfs(2) , environ(5)

NAME	srchtxt – display contents of, or search for a text string in, message data bases
SYNOPSIS	srchtxt [-s] [-l <i>locale</i>] [-m <i>msgfile</i> , ...] [<i>text</i>]
AVAILABILITY	SUNWloc
DESCRIPTION	<p>The srchtxt utility is used to display all the text strings in message data bases, or to search for a text string in message data bases (see mkmsgs(1)). These data bases are files in the directory /usr/lib/locale/locale/LC_MESSAGES (see setlocale(3C)), unless a file name given with the -m option contains a /. The directory <i>locale</i> can be viewed as the name of the language in which the text strings are written. If the -l option is not specified, the files accessed will be determined by the value of the environment variable LC_MESSAGES. If LC_MESSAGES is not set, the files accessed will be determined by the value of the environment variable LANG. If LANG is not set, the files accessed will be in the directory /usr/lib/locale/C/LC_MESSAGES, which contains default strings.</p> <p>If no <i>text</i> argument is present, then all the text strings in the files accessed will be displayed.</p> <p>If the -s option is not specified, the displayed text is prefixed by message sequence numbers. The message sequence numbers are enclosed in angle brackets: <msgfile:msgnum>.</p> <p><i>msgfile</i> name of the file where the displayed text occurred</p> <p><i>msgnum</i> sequence number in <i>msgfile</i> where the displayed text occurred</p> <p>This display is in the format used by gettext(1) and gettext(3C).</p>
OPTIONS	<p>-s Suppress printing of the message sequence numbers of the messages being displayed.</p> <p>-l <i>locale</i> Access files in the directory /usr/lib/locale/locale/LC_MESSAGES. If -m <i>msgfile</i> is also supplied, <i>locale</i> is ignored for <i>msgfiles</i> containing a /.</p> <p>-m <i>msgfile</i> Access files specified by one or more <i>msgfiles</i>. If <i>msgfile</i> contains a / character, then <i>msgfile</i> is interpreted as a pathname; otherwise, it will be assumed to be in the directory determined as described above. To specify more than one <i>msgfile</i>, separate the file names using commas.</p> <p><i>text</i> Search for the text string specified by <i>text</i> and display each one that matches. <i>text</i> can take the form of a regular expression (see ed(1)).</p>

EXAMPLES

The following examples show uses of **srchtxt**.

Example 1:

If message files have been installed in a locale named **french** by using **mkmsgs(1)**, then you could display the entire set of text strings in the **french** locale (**/usr/lib/locale/french/LC_MESSAGES/***) by typing:

```
example% srchtxt -l french
```

Example 2:

If a set of error messages associated with the operating system have been installed in the file **UX** in the **french** locale (**/usr/lib/locale/french/LC_MESSAGES/UX**), then, using the value of the **LANG** environment variable to determine the locale to be searched, you could search that file in that locale for all error messages dealing with files by typing:

```
example% setenv LANG=french; export LANG  
example% srchtxt -m UX "[Ff]ichier"
```

If **/usr/lib/locale/french/LC_MESSAGES/UX** contained the following strings:

```
Erreur E/S\n  
Liste d'arguments trop longue\n  
Fichier inexistant\n  
Argument invalide\n  
Trop de fichiers ouverts\n  
Fichier trop long\n  
Trop de liens\n  
Argument hors du domaine\n  
Identificateur supprim\n  
Etreinte fatale\n  
.  
.  
.
```

then the following strings would be displayed:

```
<UX:3>Fichier inexistant\n  
<UX:5>Trop de fichiers ouverts\n  
<UX:6>Fichier trop long\n
```

Example 3:

If a set of error messages associated with the operating system have been installed in the file **UX** and a set of error messages associated with the **INGRESS** data base product have been installed in the file **ingress**, both in the **german** locale, then you could search for the pattern **[Dd]atei** in both the files **UX** and **ingress** in the **german** locale by typing:

```
example% srchtxt -l german -m UX,ingress "[Dd]atei"
```

ENVIRONMENT

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

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

FILES

`/usr/lib/locale/C/LC_MESSAGES/*` default files created by `mkmsgs(1)`
`/usr/lib/locale/locale/LC_MESSAGES/*` message files created by `mkmsgs(1)`

SEE ALSO

`ed(1)`, `exstr(1)`, `gettext(1)`, `mkmsgs(1)`, `gettext(3C)`, `setlocale(3C)`, `environ(5)`

DIAGNOSTICS

The error messages produced by `srchtxt` are intended to be self-explanatory. They indicate an error in the command line or errors encountered while searching for a particular locale and/or message file.

NAME	strchg, strconf – change or query stream configuration
SYNOPSIS	<pre> strchg -h <i>module1</i> [, <i>module2</i> . . .] strchg -p [-a -u <i>module</i>] strchg -f <i>filename</i> strconf [-m -t <i>module</i>] </pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>These commands are used to alter or query the configuration of the stream associated with the user's standard input. The strchg command pushes modules on and/or pops modules off the stream. The strconf command queries the configuration of the stream. Only the super-user or owner of a STREAMS device may alter the configuration of that stream.</p> <p>Invoked without any arguments, strconf prints a list of all the modules in the stream as well as the topmost driver. The list is printed with one name per line where the first name printed is the topmost module on the stream (if one exists) and the last item printed is the name of the driver.</p>
OPTIONS	<p>The following options apply to strchg and, -h, -f, and -p are mutually exclusive.</p> <p>-h <i>module1</i> [, <i>module2</i> . . .] Mnemonic for <i>push</i>, pushes modules onto a stream. It takes as arguments the names of one or more pushable streams modules. These modules are pushed in order; that is, <i>module1</i> is pushed first, <i>module2</i> is pushed second, etc.</p> <p>-p Mnemonic for <i>pop</i>, pops modules off the stream. With the -p option alone, strchg pops the topmost module from the stream.</p> <p>-a <i>module</i> Pop all the modules above the topmost driver off the stream. This option requires the -p option.</p> <p>-u <i>module</i> All modules above, but not including <i>module</i> are popped off the stream. This option requires the -p option.</p> <p>-f <i>filename</i> Specify a <i>filename</i> that contains a list of modules representing the desired configuration of the stream. Each module name must appear on a separate line where the first name represents the topmost module and the last name represents the module that should be closest to the driver. strchg will determine the current configuration of the stream and pop and push the necessary modules in order to end up with the desired configuration.</p> <p>The following options apply to strconf and, -m and -t are mutually exclusive.</p> <p>-m <i>module</i> Determine if the named <i>module</i> is present on a stream. If it is, strconf prints the message yes and returns zero. If not, strconf prints the message no and returns a non-zero value. The -t and -m options are mutually exclusive.</p>

-t module

Print only the topmost module (if one exists). The **-t** and **-m** options are mutually exclusive.

EXAMPLES

The following command pushes the module **ldterm** on the stream associated with the user's standard input:

```
example% strchg -h ldterm
```

The following command pops the topmost module from the stream associated with **/dev/term/24**. The user must be the owner of this device or the super-user.

```
example% strchg -p < /dev/term/24
```

If the file **fileconf** contains the following:

```
ttcompat
ldterm
ptem
```

then the command

```
example% strchg -f fileconf
```

will configure the user's standard input stream so that the module **ptem** is pushed over the driver, followed by **ldterm** and **ttcompat** closest to the stream head.

The **strconf** command with no arguments lists the modules and topmost driver on the stream; for a stream that has only the module **ldterm** pushed above the **zs** driver, it would produce the following output:

```
ldterm
zs
```

The following command asks if **ldterm** is on the stream

```
example% strconf -m ldterm
```

and produces the following output while returning an exit status of 0:

```
yes
```

SEE ALSO

streamio(7)

DIAGNOSTICS

strchg returns zero on success. It prints an error message and returns non-zero status for various error conditions, including usage error, bad module name, too many modules to push, failure of an **ioctl** on the stream, or failure to open *filename* from the **-f** option.

strconf returns zero on success (for the **-m** or **-t** option, "success" means the named or topmost module is present). It returns a non-zero status if invoked with the **-m** or **-t** option and the module is not present. It prints an error message and returns non-zero status for various error conditions, including usage error or failure of an **ioctl** on the stream.

NOTES

If the user is neither the owner of the stream nor the super-user, the **strchg** command will fail. If the user does not have read permissions on the stream and is not the super-user, the **strconf** command will fail.

If modules are pushed in the wrong order, one could end up with a stream that does not function as expected. For ttys, if the line discipline module is not pushed in the correct place, one could have a terminal that does not respond to any commands.

NAME	strings – find printable strings in an object or binary file
SYNOPSIS	strings [-a] [-o] [<i>-number</i>] [<i>filename...</i>]
DESCRIPTION	The strings command looks for ASCII strings in a binary file. A string is any sequence of 4 or more printing characters ending with a newline or a null character. strings is useful for identifying random object files and many other things.
OPTIONS	-a Look everywhere in the file for strings. If this flag is omitted, strings only looks in the initialized data space of object files. -o Precede each string by its offset in the file. -number Use a <i>number</i> as the minimum string length rather than the default, which is 4.
ENVIRONMENT	LC_CTYPE determines how strings handles characters. When LC_CTYPE is set to a valid value, strings can display and handle text and filenames containing valid characters for that locale. strings can display and handle Extended Unix Code (EUC) characters where any character can be 1, 2, or 3 bytes wide. strings can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.
SEE ALSO	od (1)
NOTES	The algorithm for identifying strings is extremely primitive. For backwards compatibility, the options -a or just the - are interchangeable.

NAME	strip – strip symbol table, debugging and line number information from an object file
SYNOPSIS	strip [-blrVx] <i>filename</i> ...
DESCRIPTION	<p>The strip command strips the symbol table, debugging information, and line number information from ELF object files. Once this stripping process has been done, no symbolic debugging access will be available for that file; therefore, this command is normally run only on production modules that have been debugged and tested.</p> <p>If strip is executed on a common archive file (see ar(4)) in addition to processing the members, strip will remove the archive symbol table. The archive symbol table must be restored by executing the ar(1) command with the -s option before the archive can be linked by the ld(1) command. strip will produce appropriate warning messages when this situation arises.</p> <p>strip is used to reduce the file storage overhead taken by the object file.</p>
OPTIONS	<p>The amount of information stripped from the ELF object file can be controlled by using any of the following options:</p> <ul style="list-style-type: none">-b Same effect as the default behavior. This option is obsolete and will be removed in the next release.-l Strip line number information only; do not strip the symbol table or debugging information.-r Same effect as the default behavior. This option is obsolete and will be removed in the next release.-V Print, on standard error, the version number of strip.-x Do not strip the symbol table; debugging and line number information may be stripped.
FILES	<i>tmp/strip*</i> temporary files
SEE ALSO	ar(1) , as(1) , ld(1) , tmpnam(3S) , a.out(4) , ar(4)
NOTES	<p>The symbol table section will not be removed if it is contained within a segment, or the file is either a relocatable or dynamic shared object.</p> <p>The line number and debugging sections will not be removed if they are contained within a segment, or their associated relocation section is contained within a segment.</p>

NAME	stty – set the options for a terminal
SYNOPSIS	stty [-a] [-g] [<i>modes</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>stty sets certain terminal I/O options for the device that is the current standard input; without arguments, it reports the settings of certain options.</p> <p>In this report, if a character is preceded by a caret (^), then the value of that option is the corresponding control character (for example, “^h” is CTRL-H; in this case, recall that CTRL-H is the same as the “back-space” key.) The sequence “^^” means that an option has a null value.</p> <p>For detailed information about the modes listed from Control Modes through Local Modes, below, see termio(7). For detailed information about the modes listed under Hardware Flow Control Modes and Clock Modes, below, see termiox(7).</p> <p>Options described in the Combination Modes section are implemented using options in the earlier sections. Note: Many combinations of options make no sense, but no sanity checking is performed. Hardware flow control and clock modes options may not be supported by all hardware interfaces.</p>
OPTIONS	<p>-a Report all of the option settings.</p> <p>-g Report current settings in a form that can be used as an argument to another stty command.</p>
Control Modes	<p>parenb (-parenb) Enable (disable) parity generation and detection.</p> <p>parext (-parext) Enable (disable) extended parity generation and detection for mark and space parity.</p> <p>parodd (-parodd) Select odd (even) parity, or mark (space) parity if parext is enabled.</p> <p>cs5 cs6 cs7 cs8 Select character size (see termio(7)).</p> <p>0 Hang up line immediately.</p> <p>110 300 600 1200 1800 2400 4800 9600 19200 38400 Set terminal baud rate to the number given, if possible. (All speeds are not supported by all hardware interfaces.)</p> <p>ispeed 0 110 300 600 1200 1800 2400 4800 9600 19200 38400 Set terminal input baud rate to the number given, if possible. (Not all hardware supports split baud rates.) If the input baud rate is set to zero, the input baud rate will be specified by the value of the output baud rate.</p>

	ospeed 0 110 300 600 1200 1800 2400 4800 9600 19200 38400	Set terminal output baud rate to the number given, if possible. (Not all hardware supports split baud rates.) If the output baud rate is set to zero, the line will be hung up immediately.
	hupcl (-hupcl)	Hang up (do not hang up) connection on last close.
	hup (-hup)	Same as hupcl (-hupcl) .
	cstopb (-cstopb)	Use two (one) stop bits per character.
	cread (-cread)	Enable (disable) the receiver.
	crtsets (-crtsets)	Enable hardware flow control. Raise the RTS (Request to Send) modem control line. Suspends output until the CTS (Clear to Send) line is raised.
	clocal (-clocal)	Assume a line without (with) modem control.
	loblk (-loblk)	Block (do not block) output from a non-current layer.
	defeucw	Set the widths of multibyte Extended Unix Code (EUC) characters in struct <code>eucioc</code> to default values for the current locale specified by <code>LC_CTYPE</code> ; width is expressed in terms of bytes per character, and screen or display columns per character (see getwidth(3I) and ldterm(7)).
Input Modes	ignbrk (-ignbrk)	Ignore (do not ignore) break on input.
	brkint (-brkint)	Signal (do not signal) INTR on break.
	ignpar (-ignpar)	Ignore (do not ignore) parity errors.
	parmrk (-parmrk)	Mark (do not mark) parity errors (see termio(7)).
	inpck (-inpck)	Enable (disable) input parity checking.
	istrip (-istrip)	Strip (do not strip) input characters to seven bits.
	inlcr (-inlcr)	Map (do not map) NL to CR on input.
	igncr (-igncr)	Ignore (do not ignore) CR on input.
	icrnl (-icrnl)	Map (do not map) CR to NL on input.
	iuclc (-iuclc)	Map (do not map) upper-case alphabets to lower case on input.
	ixon (-ixon)	Enable (disable) START/STOP output control. Output is stopped by sending STOP control character and started by sending the START control character.
	ixany (-ixany)	Allow any character (only DC1) to restart output.
	ixoff (-ixoff)	Request that the system send (not send) START/STOP characters when the input queue is nearly empty/full.
	imaxbel (-imaxbel)	Echo (do not echo) BEL when the input line is too long.

Output Modes	opost (-opost)	Post-process output (do not post-process output; ignore all other output modes).
	olcuc (-olcuc)	Map (do not map) lower-case alphabetic to upper case on output.
	onlcr (-onlcr)	Map (do not map) NL to CR-NL on output.
	ocrnl (-ocrnl)	Map (do not map) CR to NL on output.
	onocr (-onocr)	Do not (do) output CRs at column zero.
	onlret (-onlret)	On the terminal NL performs (does not perform) the CR function.
	ofill (-ofill)	Use fill characters (use timing) for delays.
	ofdel (-ofdel)	Fill characters are DELs (NULs).
	cr0 cr1 cr2 cr3	Select style of delay for carriage returns (see termio(7)).
	nl0 nl1	Select style of delay for line-feeds (see termio(7)).
	tab0 tab1 tab2 tab3	Select style of delay for horizontal tabs (see termio(7)).
	bs0 bs1	Select style of delay for backspaces (see termio(7)).
	ff0 ff1	Select style of delay for form-feeds (see termio(7)).
	vt0 vt1	Select style of delay for vertical tabs (see termio(7)).
Local Modes	isig (-isig)	Enable (disable) the checking of characters against the special control characters INTR, QUIT, and SWTCH.
	icanon (-icanon)	Enable (disable) canonical input (ERASE and KILL processing). Does not set MIN or TIME.
	xcase (-xcase)	Canonical (unprocessed) upper/lower-case presentation.
	echo (-echo)	Echo back (do not echo back) every character typed.
	echoe (-echoe)	Echo (do not echo) ERASE character as a backspace-space-backspace string. Note: This mode will erase the ERASEed character on many CRT terminals; however, it does not keep track of column position and, as a result, may be confusing on escaped characters, tabs, and backspaces.
	echok (-echok)	Echo (do not echo) NL after KILL character.
	lfkc (-lfkc)	The same as echok (-echok); obsolete.
	echonl (-echonl)	Echo (do not echo) NL.
	noflsh (-noflsh)	Disable (enable) flush after INTR, QUIT, or SWTCH.
	stwrap (-stwrap)	Disable (enable) truncation of lines longer than 79 characters on a synchronous line.
	tostop (-tostop)	Send (do not send) SIGTTOU when background processes write to the terminal.
	echoctl (-echoctl)	Echo (do not echo) control characters as <i>^char</i> , delete as <i>^?</i>
	echoprt (-echoprt)	Echo (do not echo) erase character as character is "erased".
	echoke (-echoke)	BS-SP-BS erase (do not BS-SP-BS erase) entire line on line kill.

	flusho (-flusho)	Output is (is not) being flushed.
	pendin (-pendin)	Retype (do not retype) pending input at next read or input character.
	iexten (-iexten)	Enable (disable) extended (implementation-defined) functions for input data.
	stflush (-stflush)	Enable (disable) flush on a synchronous line after every write (2).
	stappl (-stappl)	Use application mode (use line mode) on a synchronous line.
Hardware Flow Control Modes	rtsxoff (-rtsxoff)	Enable (disable) RTS hardware flow control on input.
	ctsxon (-ctsxon)	Enable (disable) CTS hardware flow control on output.
	dtrxoff (-dtrxoff)	Enable (disable) DTR hardware flow control on input.
	cdxon (-cdxon)	Enable (disable) CD hardware flow control on output.
	isxoff (-isxoff)	Enable (disable) isochronous hardware flow control on input.
Clock Modes	xcibrg	Get transmit clock from internal baud rate generator.
	xctset	Get the transmit clock from transmitter signal element timing (DCE source) lead, CCITT V.24 circuit 114, EIA-232-D pin 15.
	xcrset	Get transmit clock from receiver signal element timing (DCE source) lead, CCITT V.24 circuit 115, EIA-232-D pin 17.
	rcibrg	Get receive clock from internal baud rate generator.
	rctset	Get receive clock from transmitter signal element timing (DCE source) lead, CCITT V.24 circuit 114, EIA-232-D pin 15.
	rcrset	Get receive clock from receiver signal element timing (DCE source) lead, CCITT V.24 circuit 115, EIA-232-D pin 17.
	tsetcoff	Transmitter signal element timing clock not provided.
	tsetcrbrg	Output receive baud rate generator on transmitter signal element timing (DTE source) lead, CCITT V.24 circuit 113, EIA-232-D pin 24.
	tsetctbrg	Output transmit baud rate generator on transmitter signal element timing (DTE source) lead, CCITT V.24 circuit 113, EIA-232-D pin 24.
	tsetctset	Output transmitter signal element timing (DCE source) on transmitter signal element timing (DTE source) lead, CCITT V.24 circuit 113, EIA-232-D pin 24.
	tsetcrset	Output receiver signal element timing (DCE source) on transmitter signal element timing (DTE source) lead, CCITT V.24 circuit 113, EIA-232-D pin 24.
	rsetcoff	Receiver signal element timing clock not provided.
	rsetcrbrg	Output receive baud rate generator on receiver signal element timing (DTE source) lead, CCITT V.24 circuit 128, no EIA-232-D pin.

	rsetctbrg	Output transmit baud rate generator on receiver signal element timing (DTE source) lead, CCITT V.24 circuit 128, no EIA-232-D pin.
	rsetctset	Output transmitter signal element timing (DCE source) on receiver signal element timing (DTE source) lead, CCITT V.24 circuit 128, no EIA-232-D pin.
	rsetcrset	Output receiver signal element timing (DCE source) on receiver signal element timing (DTE source) lead, CCITT V.24 circuit 128, no EIA-232-D pin.
Control Assignments	<i>control-character c</i>	Set <i>control-character</i> to <i>c</i> , where <i>control-character</i> is ctab , discard , dsusp , eof , eol , eol2 , erase , intr , kill , lnext , quit , reprint , start , stop , susp , swtch , or werase . (ctab is used with -stappl , see termio(7)). If <i>c</i> is preceded by a caret (^) indicating an escape from the shell, then the value used is the corresponding control character (for example, “^d” is a CTRL-D). “^?” is interpreted as DEL and “^_” is interpreted as undefined.
	min, time number	Set the value of min or time to <i>number</i> . MIN and TIME are used in Non-Canonical mode input processing (-icanon).
	line i	Set line discipline to <i>i</i> ($0 < i < 127$).
Combination Modes	evenp or parity	Enable parenb and cs7 .
	oddp	Enable parenb , cs7 , and parodd .
	spacep	Enable parenb , cs7 , and parext .
	markp	Enable parenb , cs7 , parodd , and parext .
	-parity , or -evenp	Disable parenb , and set cs8 .
	-oddp	Disable parenb and parodd , and set cs8 .
	-spacep	Disable parenb and parext , and set cs8 .
	-markp	Disable parenb , parodd , and parext , and set cs8 .
	raw (-raw or cooked)	Enable (disable) raw input and output (no ERASE, KILL, INTR, QUIT, SWTCH, EOT, or output post processing).
	nl (-nl)	Unset (set) icrnl , onlcr . In addition -nl unsets inlcr , igncr , ocrnl , and onlret .
	lcase (-lcase)	Set (unset) xcase , iuclc , and olcuc .
	LCASE (-LCASE)	Same as lcase (-lcase).
	tabs (-tabs or tab3)	Preserve (expand to spaces) tabs when printing.
	ek	Reset ERASE and KILL characters back to normal # and @.
	sane	Resets all modes to some reasonable values.
	<i>term</i>	Set all modes suitable for the terminal type <i>term</i> , where <i>term</i> is one of tty33 , tty37 , vt05 , tn300 , ti700 , or tek .

	async	Set normal asynchronous communications where clock settings are xcibrg , rcibrg , tsetcoff and rsetcoff .
Window Size	rows <i>n</i>	Set window size to <i>n</i> rows.
	columns <i>n</i>	Set window size to <i>n</i> columns.
	cols <i>n</i>	Set window size to <i>n</i> columns. Note: cols is a shorthand alias for columns
	ypixels <i>n</i>	Set vertical window size to <i>n</i> pixels.
	xpixels <i>n</i>	Set horizontal window size to <i>n</i> pixels.
SEE ALSO	tabs(1) , ioctl(2) , getwidth(3I) , ldterm(7) , termio(7) , termiox(7)	

NAME	stty – set the options for a terminal	
SYNOPSIS	/usr/ucb/stty [-a] [-g] [-h] [modes]	
AVAILABILITY	SUNWscpu	
DESCRIPTION	stty sets certain terminal I/O options for the device that is the current standard output; without arguments, it reports the settings of certain options.	
OPTIONS	<p>In this report, if a character is preceded by a caret (^), then the value of that option is the corresponding CTRL character (for example, “^h” is CTRL-h; in this case, recall that CTRL-h is the same as the “back-space” key.) The sequence “^^” means that an option has a null value.</p> <p>-a Report all of the option settings.</p> <p>-g Report current settings in a form that can be used as an argument to another stty command.</p> <p>-h Report all the option settings with the control characters in an easy to read column format.</p> <p>Options in the last group are implemented using options in the previous groups. Note: Many combinations of options make no sense, but no sanity checking is performed. Hardware flow control and clock modes options may not be supported by all hardware interfaces. The options are selected from the following:</p>	
Special Requests	all	Reports the same option settings as stty without arguments, but with the control characters in column format.
	everything	Everything stty knows about is printed. Same as -h option.
	speed	The terminal speed alone is reported on the standard output.
	size	The terminal (window) sizes are printed on the standard output, first rows and then columns. This option is only appropriate if currently running a window system.
		size and speed always report on the settings of /dev/tty, and always report the settings to the standard output.
Control Modes	parenb (-parenb)	Enable (disable) parity generation and detection.
	parext (-parext)	Enable (disable) extended parity generation and detection for mark and space parity.
	parodd (-parodd)	Select odd (even) parity, or mark (space) parity if parext is enabled.
	cs5 cs6 cs7 cs8	Select character size (see termio(7)).
	0	Hang up line immediately.
	110 300 600 1200 1800 2400 4800 9600 19200 exta 38400 extb	Set terminal baud rate to the number given, if possible. (All speeds are not supported by all hardware interfaces.)
	ispeed 0 110 300 600 1200 1800 2400 4800 9600 19200 exta 38400 extb	Set terminal input baud rate to the number given, if possible. (Not

		all hardware supports split baud rates.) If the input baud rate is set to zero, the input baud rate will be specified by the value of the output baud rate.
	ospeed 0 110 300 600 1200 1800 2400 4800 9600 19200 exta 38400 extb	Set terminal output baud rate to the number given, if possible. (Not all hardware supports split baud rates.) If the baud rate is set to zero, the line will be hung up immediately.
	hupcl (-hupcl)	Hang up (do not hang up) connection on last close.
	hup (-hup)	Same as hupcl (-hupcl) .
	cstopb (-cstopb)	Use two (one) stop bits per character.
	cread (-cread)	Enable (disable) the receiver.
	clocal (-clocal)	Assume a line without (with) modem control.
	crtcts (-crtcts)	Enable hardware flow control. Raise the RTS (Request to Send) modem control line. Suspends output until the CTS (Clear to Send) line is raised.
	loblk (-loblk)	Block (do not block) output from a non-current layer.
Input Modes	ignbrk (-ignbrk)	Ignore (do not ignore) break on input.
	brkint (-brkint)	Signal (do not signal) INTR on break.
	ignpar (-ignpar)	Ignore (do not ignore) parity errors.
	parmrk (-parmrk)	Mark (do not mark) parity errors (see termio(7)).
	inpck (-inpck)	Enable (disable) input parity checking.
	istrip (-istrip)	Strip (do not strip) input characters to seven bits.
	inlcr (-inlcr)	Map (do not map) NL to CR on input.
	igncr (-igncr)	Ignore (do not ignore) CR on input.
	icrnl (-icrnl)	Map (do not map) CR to NL on input.
	iuclc (-iuclc)	Map (do not map) upper-case alphabets to lower case on input.
	ixon (-ixon)	Enable (disable) START/STOP output control. Output is stopped by sending an STOP and started by sending an START.
	ixany (-ixany)	Allow any character (only START) to restart output.
	decctlq (-decctlq)	Same as -ixany .
	ixoff (-ixoff)	Request that the system send (not send) START/STOP characters when the input queue is nearly empty/full.
	tandem (-tandem)	Same as ixoff .
	imaxbel (-imaxbel)	Echo (do not echo) BEL when the input line is too long.
	iexten (-iexten)	Enable (disable) extended (implementation-defined) functions for input data.
Output Modes	opost (-opost)	Post-process output (do not post-process output; ignore all other output modes).
	olcuc (-olcuc)	Map (do not map) lower-case alphabets to upper case on output.
	onlcr (-onlcr)	Map (do not map) NL to CR-NL on output.
	ocrnl (-ocrnl)	Map (do not map) CR to NL on output.
	onocr (-onocr)	Do not (do) output CRs at column zero.
	onlret (-onlret)	On the terminal NL performs (does not perform) the CR function.

	ofill (-ofill)	Use fill characters (use timing) for delays.
	ofdel (-ofdel)	Fill characters are DELs (NULs).
	cr0 cr1 cr2 cr3	Select style of delay for carriage returns (see termio(7)).
	nl0 nl1	Select style of delay for line-feeds (see termio(7)).
	tab0 tab1 tab2 tab3	Select style of delay for horizontal tabs (see termio(7)).
	bs0 bs1	Select style of delay for backspaces (see termio(7)).
	ff0 ff1	Select style of delay for form-feeds (see termio(7)).
	vt0 vt1	Select style of delay for vertical tabs (see termio(7)).
Local Modes	isig (-isig)	Enable (disable) the checking of characters against the special control characters INTR, QUIT, and SWTCH.
	icanon (-icanon)	Enable (disable) canonical input (ERASE and KILL processing). Does not set MIN or TIME.
	cbreak (-cbreak)	Equivalent to -icanon min 1 time 0 .
	xcase (-xcase)	Canonical (unprocessed) upper/lower-case presentation.
	echo (-echo)	Echo back (do not echo back) every character typed.
	echoe (-echoe)	Echo (do not echo) ERASE character as a backspace-space-backspace string. Note: This mode will erase the ERASEed character on many CRT terminals; however, it does <i>not</i> keep track of column position and, as a result, may be confusing on escaped characters, tabs, and backspaces.
	crterase (-crterase)	Same as echoe .
	echok (-echok)	Echo (do not echo) NL after KILL character.
	lfkc (-lfkc)	The same as echok (-echok); obsolete.
	echonl (-echonl)	Echo (do not echo) NL.
	noflsh (-noflsh)	Disable (enable) flush after INTR, QUIT, or SWTCH.
	stwrap (-stwrap)	Disable (enable) truncation of lines longer than 79 characters on a synchronous line. (Does not apply to the 3B2.)
	tostop (-tostop)	Send (do not send) SIGTTOU for background processes.
	echoctl (-echoctl)	Echo (do not echo) control characters as <i>^char</i> , delete as <i>^?</i>
	ctlecho (-ctlecho)	Same as echoctl .
	echoprt (-echoprt)	Echo (do not echo) erase character as character is "erased".
	prterase (-prterase)	Same as echoprt .
	echoke (-echoke)	BS-SP-BS erase (do not BS-SP-BS erase) entire line on line kill.
	crtkill (-crtkill)	Same as echoke .
	flusho (-flusho)	Output is (is not) being flushed.
	pendin (-pendin)	Retype (do not retype) pending input at next read or input character.
	stflush (-stflush)	Enable (disable) flush on a synchronous line after every write(2) . (Does not apply to the 3B2.)
	stappl (-stappl)	Use application mode (use line mode) on a synchronous line. (Does not apply to the 3B2.)

Hardware Flow Control Modes	rtsxoff (-rtsxoff)	Enable (disable) RTS hardware flow control on input.
	ctsxon (-ctsxon)	Enable (disable) CTS hardware flow control on output.
	dterxoff (-dterxoff)	Enable (disable) DTER hardware flow control on input.
	rlsdxon (-rlsdxon)	Enable (disable) RLSD hardware flow control on output.
	isxoff (-isxoff)	Enable (disable) isochronous hardware flow control on input.
Clock Modes	xcibrg	Get transmit clock from internal baud rate generator.
	xctset	Get the transmit clock from transmitter signal element timing (DCE source) lead, CCITT V.24 circuit 114, EIA-232-D pin 15.
	xcrset	Get transmit clock from receiver signal element timing (DCE source) lead, CCITT V.24 circuit 115, EIA-232-D pin 17.
	rcibrg	Get receive clock from internal baud rate generator.
	rctset	Get receive clock from transmitter signal element timing (DCE source) lead, CCITT V.24 circuit 114, EIA-232-D pin 15.
	rcrset	Get receive clock from receiver signal element timing (DCE source) lead, CCITT V.24 circuit 115, EIA-232-D pin 17.
	tsetcoff	Transmitter signal element timing clock not provided.
	tsetcr	Output receive clock on transmitter signal element timing (DTE source) lead, CCITT V.24 circuit 113, EIA-232-D pin 24, clock source.
	tsetcxc	Output transmit clock on transmitter signal element timing (DTE source) lead, CCITT V.24 circuit 113, EIA-232-D pin 24, clock source.
	rsetcoff	Receiver signal element timing clock not provided.
	rsetcr	Output receive clock on receiver signal element timing (DTE source) lead, CCITT V.24 circuit 128, no EIA-232-D pin, clock source.
rsetcxc	Output transmit clock on receiver signal element timing (DTE source) lead, CCITT V.24 circuit 128, no EIA-232-D pin, clock source.	
Control Assignments	control-character <i>c</i>	Set <i>control-character</i> to <i>c</i> , where <i>control-character</i> is intr , quit , erase , kill , eof , eol , eol2 , swtch , start , stop , susp , dsusp , rpnt , flush , werase , lnext , min , ctab , time , or brk (ctab is used with -stappl ; min and time are used with -icanon ; see termio(7)). If <i>c</i> is preceded by an (escaped from the shell) caret (^), then the value used is the corresponding CTRL character (for example, “^d” is a CTRL-d); “^?” is interpreted as DEL and “^_” is interpreted as undefined.
	line <i>i</i>	Set line discipline to <i>i</i> ($0 < i < 127$).
Combination Modes	evenp or parity	Enable parenb and cs7 .
	-evenp , or -parity	Disable parenb , and set cs8 .
	even (-even)	Same as evenp (-evenp).
	oddp	Enable parenb , cs7 , and parodd .

	-oddp	Disable parenb and parodd , and set cs8 .
	odd (-odd)	Same as oddp (-oddp).
	spacep	Enable parenb , cs7 , and parext .
	-spacep	Disable parenb and parext , and set cs8 .
	markp	Enable parenb , cs7 , parodd , and parext .
	-markp	Disable parenb , parodd , and parext , and set cs8 .
	raw (-raw or cooked)	Enable (disable) raw input and output (no ERASE, KILL, INTR, QUIT, SWTCH, EOT, or output post processing).
	nl (-nl)	Unset (set) icrnl , onlcr . In addition -nl unsets inlcr , igncr , ocrnl , and onlret .
	lcase (-lcase)	Set (unset) xcase , iucl , and olcuc .
	LCASE (-LCASE)	Same as lcase (-lcase).
	tabs (-tabs or tab3)	Preserve (expand to spaces) tabs when printing.
	ek	Reset ERASE and KILL characters back to normal # and @.
	sane	Resets all modes to some reasonable values.
	term	Set all modes suitable for the terminal type <i>term</i> , where <i>term</i> is one of tty33 , tty37 , vt05 , tn300 , ti700 , or tek .
	async	Set normal asynchronous communications where clock settings are xcibrg , rcibrg , tsetcoff and rsetcoff .
	litout (-litout)	Disable (enable) parenb , istrip , and opost , and set cs8 (cs7).
	pass8 (-pass8)	Disable (enable) parenb and istrip , and set cs8 (cs7).
	crt	Set options for a CRT (echoe , echoctl , and, if ≥ 1200 baud, echoke .)
	dec	Set all modes suitable for Digital Equipment Corp. operating systems users (ERASE , KILL , and INTR characters to ^? , ^U , and ^C , decctlq , and crt .)
Window Size	rows <i>n</i>	Set window size to <i>n</i> rows.
	columns <i>n</i>	Set window size to <i>n</i> columns.
	cols <i>n</i>	An alias for columns <i>n</i> .
	ypixels <i>n</i>	Set vertical window size to <i>n</i> pixels.
	xpixels <i>n</i>	Set horizontal window size to <i>n</i> pixels.
SEE ALSO	tabs(1) , ioctl(2) , termio(7) , termiox(7)	

NAME	sum – print checksum and block count for a file
SYNOPSIS	sum [-r] <i>filename</i>
AVAILABILITY	SUNWesu
DESCRIPTION	sum calculates and prints a 16-bit checksum for the named file, and also prints the number of 512 byte blocks in the file. It is typically used to look for bad spots, or to validate a file communicated over some transmission line.
OPTIONS	-r Use an alternate (machine-dependent) algorithm in computing the checksum.
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of sum for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how sum behaves.</p> <p>LC_CTYPE Determines how sum handles characters. When LC_CTYPE is set to a valid value, sum can display and handle text and filenames containing valid characters for that locale. sum can display and handle Extended Unix code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. sum can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p> <p>LC_MESSAGES Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).</p>
SEE ALSO	wc(1) , environ(5)
DIAGNOSTICS	“Read error” is indistinguishable from end of file on most devices; check the block count.

NAME	sum – calculate a checksum for a file
SYNOPSIS	<i>/usr/ucb/sum filename</i>
AVAILABILITY	SUNWscpu
DESCRIPTION	sum calculates and displays a 16-bit checksum for the named file, and also displays the size of the file in kilobytes. It is typically used to look for bad spots, or to validate a file communicated over some transmission line. The checksum is calculated by an algorithm which may yield different results on machines with 16-bit ints and machines with 32-bit ints , so it cannot always be used to validate that a file has been transferred between machines with different-sized ints .
SEE ALSO	sum(1) , wc(1)
DIAGNOSTICS	Read error is indistinguishable from EOF on most devices; check the block count.
NOTES	Obsolete.

NAME	suspend – shell built-in function to halt the current shell
SYNOPSIS	
sh	suspend
csch	suspend
ksh	suspend
DESCRIPTION	
sh	Stops the execution of the current shell (but not if it is the login shell).
csch	Stop the shell in its tracks, much as if it had been sent a stop signal with ^Z . This is most often used to stop shells started by su .
ksh	Stops the execution of the current shell (but not if it is the login shell).
SEE ALSO	csch(1) , ksh(1) , kill(1) , sh(1) , su(1)

NAME	symorder – rearrange a list of symbols
SYNOPSIS	symorder [-s] <i>objectfile</i> <i>symbolfile</i>
DESCRIPTION	<p>This program was specifically designed to cut down on the overhead of getting symbols from /kernel/unix.</p> <p><i>objectfile</i> is updated in place to put the requested symbols first in the symbol table, in the order specified. This is done by swapping the old symbols in the required spots with the new ones. If all of the order symbols are not found, an error is generated.</p> <p><i>symbolfile</i> is a file containing symbols to be found in <i>objectfile</i>, one symbol per line.</p>
OPTIONS	-s Work silently, that is, display nothing except error messages. This is useful for checking the error status.
FILES	/kernel/unix
SEE ALSO	nlist(3E)

NAME	sysV-make – maintain, update, and regenerate groups of programs
SYNOPSIS	<code>/usr/ccs/lib/svr4.make [-f <i>makefile</i>] [-eiknpqrst] [<i>names</i>]</code>
DESCRIPTION	<p>This is the “vanilla” System V version of make. If the environment variable <code>USE_SVR4_MAKE</code> is set, then the command make will invoke this version of make. (See also the <code>ENVIRONMENT</code> section.)</p> <p>make allows the programmer to maintain, update, and regenerate groups of computer programs. make executes commands in <i>makefile</i> to update one or more target <i>names</i> (<i>names</i> are typically programs). If the <code>-f</code> option is not present, then makefile, Makefile, and the Source Code Control System (SCCS) files s.makefile, and s.Makefile are tried in order. If <i>makefile</i> is <code>-</code>, the standard input is taken. More than one <code>-f <i>makefile</i></code> argument pair may appear.</p> <p>make updates a target only if its dependents are newer than the target. All prerequisite files of a target are added recursively to the list of targets. Missing files are deemed to be outdated.</p> <p>The following list of four directives can be included in <i>makefile</i> to extend the options provided by make. They are used in <i>makefile</i> as if they were targets:</p> <ul style="list-style-type: none"> .DEFAULT: If a file must be made but there are no explicit commands or relevant built-in rules, the commands associated with the name .DEFAULT are used if it exists. .IGNORE: Same effect as the <code>-i</code> option. .PRECIOUS: Dependents of the .PRECIOUS entry will not be removed when quit or interrupt are hit. .SILENT: Same effect as the <code>-s</code> option. <p>The options for make are listed below:</p> <ul style="list-style-type: none"> <code>-e</code> Environment variables override assignments within makefiles. <code>-f <i>makefile</i></code> Description filename (<i>makefile</i> is assumed to be the name of a description file). <code>-i</code> Ignore error codes returned by invoked commands. <code>-k</code> Abandon work on the current entry if it fails, but continue on other branches that do not depend on that entry. <code>-n</code> No execute mode. Print commands, but do not execute them. Even command lines beginning with an <code>@</code> are printed. <code>-p</code> Print out the complete set of macro definitions and target descriptions. <code>-q</code> Question. make returns a zero or non-zero status code depending on whether or not the target file has been updated. <code>-r</code> Do not use the built-in rules. <code>-s</code> Silent mode. Do not print command lines before executing.

-t Touch the target files (causing them to be updated) rather than issue the usual commands.

Creating the makefile

The makefile invoked with the **-f** option is a carefully structured file of explicit instructions for updating and regenerating programs, and contains a sequence of entries that specify dependencies. The first line of an entry is a blank-separated, non-null list of targets, then a **:**, then a (possibly null) list of prerequisite files or dependencies. Text following a **;** and all following lines that begin with a tab are shell commands to be executed to update the target. The first non-empty line that does not begin with a tab or **#** begins a new dependency or macro definition. Shell commands may be continued across lines with a backslash-new-line (**\ new-line**) sequence. Everything printed by **make** (except the initial tab) is passed directly to the shell as is. Thus,

```
echo a\  
b
```

will produce

```
ab
```

exactly the same as the shell would.

Sharp (**#**) and new-line surround comments including contained **\ new-line** sequences.

The following makefile says that **pgm** depends on two files **a.o** and **b.o**, and that they in turn depend on their corresponding source files (**a.c** and **b.c**) and a common file **incl.h**:

```
pgm: a.o b.o  
    cc a.o b.o -o pgm  
a.o: incl.h a.c  
    cc -c a.c  
b.o: incl.h b.c  
    cc -c b.c
```

Command lines are executed one at a time, each by its own shell. The **SHELL** environment variable can be used to specify which shell **make** should use to execute commands. The default is **/usr/bin/sh**. The first one or two characters in a command can be the following: **@**, **-**, **@-**, or **-@**. If **@** is present, printing of the command is suppressed. If **-** is present, **make** ignores an error. A line is printed when it is executed unless the **-s** option is present, or the entry **.SILENT:** is included in *makefile*, or unless the initial character sequence contains a **@**. The **-n** option specifies printing without execution; however, if the command line has the string **\$(MAKE)** in it, the line is always executed (see the discussion of the **MAKEFLAGS** macro in the “make Environment” section below). The **-t** (touch) option updates the modified date of a file without executing any commands.

Commands returning non-zero status normally terminate **make**. If the **-i** option is present, if the entry **.IGNORE:** is included in *makefile*, or if the initial character sequence of the command contains **-**, the error is ignored. If the **-k** option is present, work is abandoned on the current entry, but continues on other branches that do not depend on that entry.

	<p>Interrupt and quit cause the target to be deleted unless the target is a dependent of the directive .PRECIOUS.</p>
make Environment	<p>The environment is read by make. All variables are assumed to be macro definitions and are processed as such. The environment variables are processed before any makefile and after the internal rules; thus, macro assignments in a makefile override environment variables. The -e option causes the environment to override the macro assignments in a makefile. Suffixes and their associated rules in the makefile will override any identical suffixes in the built-in rules.</p> <p>The MAKEFLAGS environment variable is processed by make as containing any legal input option (except -f and -p) defined for the command line. Further, upon invocation, make “invents” the variable if it is not in the environment, puts the current options into it, and passes it on to invocations of commands. Thus, MAKEFLAGS always contains the current input options. This feature proves very useful for “super-makes”. In fact, as noted above, when the -n option is used, the command \$(MAKE) is executed anyway; hence, one can perform a make -n recursively on a whole software system to see what would have been executed. This result is possible because the -n is put in MAKEFLAGS and passed to further invocations of \$(MAKE). This usage is one way of debugging all of the makefiles for a software project without actually doing anything.</p>
Include Files	<p>If the string include appears as the first seven letters of a line in a <i>makefile</i>, and is followed by a blank or a tab, the rest of the line is assumed to be a filename and will be read by the current invocation, after substituting for any macros.</p>
Macros	<p>Entries of the form <i>string1</i> = <i>string2</i> are macro definitions. <i>string2</i> is defined as all characters up to a comment character or an unescaped new-line. Subsequent appearances of \$(string1[:subst1=[subst2]]) are replaced by <i>string2</i>. The parentheses are optional if a single-character macro name is used and there is no substitute sequence. The optional :subst1=subst2 is a substitute sequence. If it is specified, all non-overlapping occurrences of <i>subst1</i> in the named macro are replaced by <i>subst2</i>. Strings (for the purposes of this type of substitution) are delimited by blanks, tabs, new-line characters, and beginnings of lines. An example of the use of the substitute sequence is shown in the “Libraries” section below.</p>
Internal Macros	<p>There are five internally maintained macros that are useful for writing rules for building targets.</p> <p>\$* The macro \$* stands for the filename part of the current dependent with the suffix deleted. It is evaluated only for inference rules.</p> <p>\$@ The \$@ macro stands for the full target name of the current target. It is evaluated only for explicitly named dependencies.</p>

\$< The **\$<** macro is only evaluated for inference rules or the **.DEFAULT** rule. It is the module that is outdated with respect to the target (the “manufactured” dependent file name). Thus, in the **.c.o** rule, the **\$<** macro would evaluate to the **.c** file. An example for making optimized **.o** files from **.c** files is:

```
.c.o:
    cc -c -O $*.c
```

or:

```
.c.o:
    cc -c -O $<
```

\$? The **\$?** macro is evaluated when explicit rules from the makefile are evaluated. It is the list of prerequisites that are outdated with respect to the target, and essentially those modules that must be rebuilt.

\$% The **\$%** macro is only evaluated when the target is an archive library member of the form **lib(file.o)**. In this case, **\$@** evaluates to **lib** and **\$%** evaluates to the library member, **file.o**.

Four of the five macros can have alternative forms. When an upper case **D** or **F** is appended to any of the four macros, the meaning is changed to “directory part” for **D** and “file part” for **F**. Thus, **\$(@D)** refers to the directory part of the string **\$@**. If there is no directory part, **/** is generated. The only macro excluded from this alternative form is **\$?**.

Suffixes

Certain names (for instance, those ending with **.o**) have inferable prerequisites such as **.c**, **.s**, etc. If no update commands for such a file appear in *makefile*, and if an inferable prerequisite exists, that prerequisite is compiled to make the target. In this case, **make** has inference rules that allow building files from other files by examining the suffixes and determining an appropriate inference rule to use. The current default inference rules are:

```
.c      .c~    .f      .f~    .s      .s~    .sh     .sh~   .C      .C~
.c.a    .c.o   .c~.a   .c~.c   .c~.o   .f.a    .f.o    .f.a    .f.f    .f~.o
.h~.h   .l.c    .l.o    .l~.c   .l~.l   .l~.o   .s.a    .s.o    .s~.a   .s~.o
.s~.s   .sh~.sh .y.c    .y.o    .y~.c   .y~.o   .y~.y   .C.a    .C.o    .C~.a
.C~.C   .C~.o   .L.C    .L.o    .L~.C   .L~.L   .L~.o   .Y.C    .Y.o    .Y~.C
.Y~.o   .Y.Y
```

The internal rules for **make** are contained in the source file **rules.c** for the **make** program. These rules can be locally modified. To print out the rules compiled into the **make** on any machine in a form suitable for recompilation, the following command is used:

```
make -pf - 2>/dev/null </dev/null
```

A tilde in the above rules refers to an SCCS file (see **sccsfile(4)**). Thus, the rule **.c~.o** would transform an SCCS C source file into an object file (**.o**). Because the **s.** of the SCCS files is a prefix, it is incompatible with the **make** suffix point of view. Hence, the tilde is a way of changing any file reference into an SCCS file reference.

A rule with only one suffix (for example, **.c**) is the definition of how to build **x** from **x.c**. In effect, the other suffix is null. This feature is useful for building targets from only one source file, for example, shell procedures and simple C programs.

Additional suffixes are given as the dependency list for **.SUFFIXES**. Order is significant: the first possible name for which both a file and a rule exist is inferred as a prerequisite. The default list is:

```
.SUFFIXES: .o .c .c~ .y .y~ .I .I~ .s .s~ .sh .sh~ .h .h~ .f .f~ .C .C~ .Y .Y~ .L .L~
```

Here again, the above command for printing the internal rules will display the list of suffixes implemented on the current machine. Multiple suffix lists accumulate; **.SUFFIXES:** with no dependencies clears the list of suffixes.

Inference Rules

The first example can be done more briefly.

```
pgm: a.o b.o
cc a.o b.o -o pgm
a.o b.o: incl.h
```

This abbreviation is possible because **make** has a set of internal rules for building files. The user may add rules to this list by simply putting them in the *makefile*.

Certain macros are used by the default inference rules to permit the inclusion of optional matter in any resulting commands. For example, **CFLAGS**, **LFLAGS**, and **YFLAGS** are used for compiler options to **cc(1B)**. Again, the previous method for examining the current rules is recommended.

The inference of prerequisites can be controlled. The rule to create a file with suffix **.o** from a file with suffix **.c** is specified as an entry with **.c.o:** as the target and no dependents. Shell commands associated with the target define the rule for making a **.o** file from a **.c** file. Any target that has no slashes in it and starts with a dot is identified as a rule and not a true target.

Libraries

If a target or dependency name contains parentheses, it is assumed to be an archive library, the string within parentheses referring to a member within the library. Thus, **lib(file.o)** and **\$(LIB)(file.o)** both refer to an archive library that contains **file.o**. (This example assumes the **LIB** macro has been previously defined.) The expression **\$(LIB)(file1.o file2.o)** is not legal. Rules pertaining to archive libraries have the form **.XX.a** where the **XX** is the suffix from which the archive member is to be made. An unfortunate by-product of the current implementation requires the **XX** to be different from the suffix of the archive member. Thus, one cannot have **lib(file.o)** depend upon **file.o** explicitly. The most common use of the archive interface follows. Here, we assume the source files are all C type source:

```
lib: lib(file1.o) lib(file2.o) lib(file3.o)
@echo lib is now up-to-date
.c.a:
$(CC) -c $(CFLAGS) $<
$(AR) $(ARFLAGS) $@ $*.o
rm -f $*.o
```

In fact, the `.c.a` rule listed above is built into **make** and is unnecessary in this example. A more interesting, but more limited example of an archive library maintenance construction follows:

```
lib:    lib(file1.o) lib(file2.o) lib(file3.o)
        $(CC) -c $(CFLAGS) $(?:.o=.c)
        $(AR) $(ARFLAGS) lib $?
        rm $?
        @echo lib is now up-to-date

.c.a;
```

Here the substitution mode of the macro expansions is used. The `?$` list is defined to be the set of object filenames (inside `lib`) whose C source files are outdated. The substitution mode translates the `.o` to `.c`. (Unfortunately, one cannot as yet transform to `.c~`; however, this transformation may become possible in the future.) Also note the disabling of the `.c.a` rule, which would have created each object file, one by one. This particular construct speeds up archive library maintenance considerably. This type of construct becomes very cumbersome if the archive library contains a mix of assembly programs and C programs.

ENVIRONMENT

USE_SVR4_MAKE

If this environment variable is set, then the **make** command will invoke this System V version of **make**. If this variable is not set, then the default version of **make(1S)** is invoked.

`USE_SVR4_MAKE` can be set as follows (Bourne shell):

```
$ USE_SVR4_MAKE="" ; export USE_SVR4_MAKE
```

or (C shell):

```
% setenv USE_SVR4_MAKE
```

FILES

`[Mm]akefile` and `s.[Mm]akefile`
`/usr/bin/sh`

SEE ALSO

`cc(1B)`, `cd(1)`, `make(1S)`, `sh(1)`, `printf(3S)`, `sccsfile(4)`

Programming Utilities Guide

NOTES

Some commands return non-zero status inappropriately; use `-i` or the `-` command line prefix to overcome the difficulty.

Filenames with the characters `= : @` will not work. Commands that are directly executed by the shell, notably `cd(1)`, are ineffectual across new-lines in **make**. The syntax `lib(file1.o file2.o file3.o)` is illegal. You cannot build `lib(file.o)` from `file.o`.

NAME	tabs – set tabs on a terminal
SYNOPSIS	tabs [<i>tabspec</i>] [-T <i>type</i>] [+mn]
AVAILABILITY	SUNWcsu
DESCRIPTION	tabs sets the tab stops on the user's terminal according to the tab specification <i>tabspec</i> , after clearing any previous settings. The user's terminal must have remotely settable hardware tabs.
OPTIONS	<p><i>tabspec</i> Four types of tab specification are accepted for <i>tabspec</i>. They are described below: canned (<i>-code</i>), repetitive (<i>-n</i>), arbitrary (<i>n1,n2,...</i>), and file (<i>—file</i>). If no <i>tabspec</i> is given, the default value is -8, that is, UNIX system "standard" tabs. The lowest column number is 1. Note: For tabs, column 1 always refers to the leftmost column on a terminal, even one whose column markers begin at 0, for example, the DASI 300, DASI 300s, and DASI 450.</p> <p><i>-code</i> Use one of the codes listed below to select a <i>canned</i> set of tabs. The legal codes and their meanings are as follows:</p> <ul style="list-style-type: none"> -a 1,10,16,36,72 Assembler, IBM S/370, first format -a2 1,10,16,40,72 Assembler, IBM S/370, second format -c 1,8,12,16,20,55 COBOL, normal format -c2 1,6,10,14,49 COBOL compact format (columns 1-6 omitted). Using this code, the first typed character corresponds to card column 7, one space gets you to column 8, and a tab reaches column 12. Files using this tab setup should include a format specification as follows (see fspec(4)): <:t-c2 m6 s66 d:> -c3 1,6,10,14,18,22,26,30,34,38,42,46,50,54,58,62,67 COBOL compact format (columns 1-6 omitted), with more tabs than -c2. This is the recommended format for COBOL. The appropriate format specification is (see fspec(4)): <:t-c3 m6 s66 d:> -f 1,7,11,15,19,23 FORTRAN -p 1,5,9,13,17,21,25,29,33,37,41,45,49,53,57,61 PL/I -s 1,10,55 SNOBOL -u 1,12,20,44 UNIVAC 1100 Assembler

- n** A *repetitive* specification requests tabs at columns $1+n$, $1+2*n$, etc. Of particular importance is the value **8**: this represents the UNIX system “standard” tab setting, and is the most likely tab setting to be found at a terminal. Another special case is the value **0**, implying no tabs at all.
- n1,n2,...** The *arbitrary* format permits the user to type any chosen set of numbers, separated by commas, in ascending order. Up to 40 numbers are allowed. If any number (except the first one) is preceded by a plus sign, it is taken as an increment to be added to the previous value. Thus, the formats **1,10,20,30**, and **1,10,+10,+10** are considered identical.
- file** If the name of a *file* is given, **tabs** reads the first line of the file, searching for a format specification (see **fspec(4)**). If it finds one there, it sets the tab stops according to it, otherwise it sets them as **-8**. This type of specification may be used to make sure that a tabbed file is printed with correct tab settings, and would be used with the **pr** command:

example% tabs — file; pr file

Any of the following also may be used; if a given flag occurs more than once, the last value given takes effect:

- Ttype** **tabs** usually needs to know the type of terminal in order to set tabs and always needs to know the type to set margins. *type* is a name listed in **term(5)**. If no **-T** flag is supplied, **tabs** uses the value of the environment variable **TERM**. If **TERM** is not defined in the *environment* (see **environ(5)**), **tabs** tries a sequence that will work for many terminals.
- +mn** The margin argument may be used for some terminals. It causes all tabs to be moved over *n* columns by making column $n+1$ the left margin. If **+m** is given without a value of *n*, the value assumed is **10**. For a TerminoNet, the first value in the tab list should be **1**, or the margin will move even further to the right. The normal (leftmost) margin on most terminals is obtained by **+m0**. The margin for most terminals is reset only when the **+m** flag is given explicitly.

Tab and margin setting is performed via the standard output.

EXAMPLES

The command:

example% tabs -a

is an example using *-code* (*canned* specification) to set tabs to the settings required by the IBM assembler: columns 1, 10, 16, 36, 72.

The next command:

example% tabs -8

is an example of using *-n* (*repetitive* specification), where *n* is **8**, causes tabs to be set every eighth position:

$1+(1*8)$, $1+(2*8)$, ... which evaluate to columns 9, 17, ...

The command:

example% tabs 1,8,36

is an example of using *n1,n2,...* (*arbitrary* specification) to set tabs at columns 1, 8, and 36.

The last command:

example% tabs —\$HOME/fspec.list/att4425

is an example of using *—file* (*file* specification) to indicate that tabs should be set according to the first line of *\$HOME/fspec.list/att4425* (see *fspec*(4)).

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

newform(1), *pr*(1), *tput*(1), *fspec*(4), *terminfo*(4), *environ*(5), *term*(5)

DIAGNOSTICS

- illegal tabs** This occurs when arbitrary tabs are ordered incorrectly.
- illegal increment** This occurs when a zero or missing increment is found in an arbitrary specification.
- unknown tab code** This occurs when a *canned* code cannot be found.
- can't open** The *—file* option was used, and file cannot be opened.
- file indirection** The *—file* option was used and the specification in that file points to yet another file. Indirection of this form is not permitted.

NOTES

There is no consistency among different terminals regarding ways of clearing tabs and setting the left margin.

tabs clears only 20 tabs (on terminals requiring a long sequence), but is willing to set 64. The *tabspec* used with the **tabs** command is different from the one used with the **newform** command. For example, **tabs -8** sets every eighth position; whereas **newform -i-8** indicates that tabs are set every eighth position.

NAME	tail – deliver the last part of a file
SYNOPSIS	<pre>tail [± number [lbc]] [filename] tail [-lbc] [filename] tail [± number [lbc]] [filename] tail [-lbc] [filename]</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>tail copies the named file to the standard output beginning at a designated place. If no file is named, the standard input is used.</p> <p>Copying begins at distance <i>+number</i> from the beginning, or <i>-number</i> from the end of the input (if <i>number</i> is null, the value 10 is assumed). <i>Number</i> is counted in units of lines, blocks, or characters, according to the appended option l, b, or c. When no units are specified, counting is by lines.</p> <p>The r and f options are mutually exclusive. If both are specified on the command line, the f option will be ignored.</p>
OPTIONS	<pre>-l Units of lines. -b Units of blocks. -c Units of characters. -f With the -f (follow) option, if the input-file is not a pipe, the program will not terminate after the line of the input-file has been copied, but will enter an end- less loop, wherein it sleeps for a second and then attempts to read and copy further records from the input-file. Thus it may be used to monitor the growth of a file that is being written by some other process. -r The r option copies lines from the specified starting point in the file in reverse order. The default for r is to print the entire file in reverse order.</pre>
EXAMPLES	<p>For example, the command:</p> <pre>example% tail -f fred</pre> <p>will print the last ten lines of the file fred, followed by any lines that are appended to fred between the time tail is initiated and killed. As another example, the command:</p> <pre>example% tail -15cf fred</pre> <p>will print the last 15 characters of the file fred, followed by any lines that are appended to fred between the time tail is initiated and killed.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of tail for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in</p>

the environment, the "C" (U.S. style) locale determines how **tail** behaves.

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

cat(1), head(1), more(1), pg(1), tail(1), dd(1M), environ(5)

NOTES

Piped tails relative to the end of the file are stored in a buffer, and thus are limited in length. Various kinds of anomalous behavior may happen with character special files.

NAME	talk – talk to another user				
SYNOPSIS	talk <i>username</i> [<i>ttyname</i>]				
AVAILABILITY	SUNWcsu				
DESCRIPTION	<p>talk is a visual communication program that copies lines from your terminal to that of a user on the same or on another host. <i>username</i> is that user's login name.</p> <p>The program is architecture dependent; it works only between machines of the same architecture.</p> <p>If you want to talk to a user who is logged in more than once, the <i>ttyname</i> argument may be used to indicate the appropriate terminal name.</p> <p>When first called, talk sends the message:</p> <p style="padding-left: 40px;">Message from TalkDaemon@ her_machine at time . . . talk: connection requested by your_name@your_machine talk: respond with: talk your_name@your_machine</p> <p>to the user you wish to talk to. At this point, the recipient of the message should reply by typing:</p> <p style="padding-left: 40px;">talk your_name@your_machine</p> <p>It does not matter from which machine the recipient replies, as long as the login name is the same. Once communication is established, the two parties may type simultaneously, with their output appearing in separate windows. Typing Ctrl-L redraws the screen, while your erase, kill, and word kill characters will work in talk as normal. To exit, just type your interrupt character; talk then moves the cursor to the bottom of the screen and restores the terminal.</p> <p>Permission to talk may be denied or granted by use of the mesg(1) command. At the outset talking is allowed. Certain commands, such as pr(1), disallow messages in order to prevent messy output.</p>				
FILES	<table border="0" style="width: 100%;"> <tr> <td style="padding-right: 20px;">/etc/hosts</td> <td>to find the recipient's machine</td> </tr> <tr> <td>/var/adm/utmp</td> <td>to find the recipient's tty</td> </tr> </table>	/etc/hosts	to find the recipient's machine	/var/adm/utmp	to find the recipient's tty
/etc/hosts	to find the recipient's machine				
/var/adm/utmp	to find the recipient's tty				
SEE ALSO	mail(1), mesg(1), pr(1), who(1), write(1)				

NAME	tar – create tape archives, and add or extract files
SYNOPSIS	<pre> /usr/sbin/tar c [bBefFhilvwX [0-7]] [device] [block] [exclude-filename ...] [-I include-filename] filename ... [-C directory filename] /usr/sbin/tar r [bBefFhilvw [0-7]] [device] [block] [-I include-filename] filename ... [-C directory filename] /usr/sbin/tar t [BefFhilvX [0-7]] [device] [exclude-filename ...] [-I include-filename] [filename ...] /usr/sbin/tar u [bBefFhilvw [0-7]] [device] [block] filename ... /usr/sbin/tar x [BefFhilmopvwX [0-7]] [device] [exclude-filename ...] [filename ...] </pre>
DESCRIPTION	<p>tar archives and extracts files to and from a single file called a <i>tarfile</i>. A <i>tarfile</i> is usually a magnetic tape, but it can be any file. tar's actions are controlled by the <i>key</i> argument. The <i>key</i> is a string of characters containing exactly one function letter (c, r, t, u, or x) and one or more function modifiers, depending on the function letter used. Other arguments to the command are <i>filenames</i> (or directory names) specifying which files are to be archived or extracted. In all cases, appearance of a directory name refers to the files and (recursively) subdirectories of that directory.</p>
Function Letters	<p>The function portion of the key is specified by one of the following letters:</p> <p>c Create. Writing begins at the beginning of the <i>tarfile</i>, instead of at the end. This key implies the r key.</p> <p>r Replace. The named <i>filenames</i> are written on the end of the tape. The c and u functions imply this function. See NOTES for more information.</p> <p>t Table of Contents. The names of the specified files are listed each time they occur on the <i>tarfile</i>. If no <i>filenames</i> arguments are given, all the names on the <i>tarfile</i> are listed. With the v function modifier, additional information for the specified files is displayed. The listing is similar to the format produced by the ls -l command.</p> <p>u Update. The named <i>filenames</i> are added to the <i>tarfile</i> if they are not already there, or have been modified since last written on that <i>tarfile</i>. This key implies the r key. See NOTES for more information.</p> <p>x Extract, or restore. The named <i>filenames</i> are extracted from the <i>tarfile</i> and written to the current directory. If a named file matches a directory whose contents had been written onto the <i>tarfile</i>, this directory is (recursively) extracted. Use the file or directory's relative path when appropriate, or tar will not find a match. The owner, modification time, and mode are restored (if possible); otherwise, to restore owner, you must be the superuser. If no <i>filenames</i> argument is given, the entire content of the <i>tarfile</i> is extracted. Note: If several files with the same name are on the <i>tarfile</i>, the last one overwrites all earlier ones. See NOTES for more information.</p>

Function Modifiers

The characters below may be used in addition to the letter that selects the desired function. Use them in the order shown in the **SYNOPSIS**.

- b** Blocking Factor. This causes **tar** to use the *block* argument as the blocking factor for tape records. The default is **1**, the maximum is **20**. This function should not be supplied when operating on regular archives or block special devices. It is mandatory however, when reading archives on raw magnetic tape archives (see **f** below). The block size is determined automatically when reading tapes created on block special devices (key letters **x** and **t**). This determination of the blocking factor may be fooled when reading from a pipe or a socket (see the **B** key letter below). The maximum blocking factor is determined only by the amount of memory available to **tar** when it is run. Larger blocking factors result in better throughput, longer blocks on nine-track tapes, and better media utilization.
- B** Block. Force **tar** to perform multiple reads (if necessary) so as to read exactly enough bytes to fill a block. This option exists so that **tar** can work across the Ethernet, since pipes and sockets return partial blocks even when more data is coming. When reading from standard input, '-', this option is automatically selected to make sure that **tar** can recover from short reads.
- e** Error. If any unexpected errors occur **tar** will exit immediately with a positive exit status.
- f** File. This causes **tar** to use the *device* argument as the name of the tarfile. If **f** is given, **/etc/default/tar** is not searched. If **f** is omitted, **tar** will use the device indicated by the **TAPE** environment variable, if set; otherwise, it will use the default values defined in **/etc/default/tar**. If the name of the tarfile is '-', **tar** writes to the standard output or reads from the standard input, whichever is appropriate. Thus, **tar** can be used as the head or tail of a pipeline. **tar** can also be used to move hierarchies with the command:
 example% cd fromdir; tar cf - . | (cd todir; tar xfBp -)
- F** With one **F** argument, **tar** will exclude all directories named **SCCS** from the tarfile. With two arguments, **FF**, **tar** will exclude all directories named **SCCS**, all files with **.o** as their suffix, and all files named **errs**, **core**, and **a.out**.
- h** Follow symbolic links as if they were normal files or directories. Normally, **tar** does not follow symbolic links.
- i** Ignore. With this option **tar** will ignore directory checksum errors.
- l** Link. This tells **tar** to complain if it cannot resolve all of the links to the files being archived. If **l** is not specified, no error messages are printed.
- m** Modify. This tells **tar** to not extract the modification times from the tarfile. The modification time of the file will be the time of extraction. This option is only valid with the **x** key.
- o** Ownership. This causes extracted files to take on the user and group identifier of the user running the program, rather than those on tape. This happens by default for users other than root. If the 'o' option is not set and the user is root,

- the extracted files will take on the group and user identifiers of the files on tape (see **chown**(1) for more information). The 'o' option is only valid with the **x** key.
- p** Restore the named files to their original modes, ignoring the present **umask**(2). SETUID and sticky information are also extracted if you are the super-user. This option is only useful with the **x** key letter.
- v** Verbose. Normally, **tar** does its work silently. This option causes **tar** to type the name of each file it treats, preceded by the function letter. With the **t** function, **v** gives more information about the tape entries than just the name.
- w** What. This option causes **tar** to print the action to be taken, followed by the name of the file, and then wait for the user's confirmation. If a word beginning with **y** is given, the action is performed. Any other input means no. This is not valid with the **t** key.
- X** Exclude. Use the *exclude-filename* argument as a file containing a list of named files (or directories) to be excluded from the tarfile when using the key letters **c**, **x**, or **t**. Multiple **X** arguments may be used, with one *exclude-filename* per argument. See **NOTES** for more information.
- [0-7]** Select an alternative drive on which the tape is mounted. The default entries are specified in **/etc/default/tar**.

If a file name is preceded by **-I** then the filename is opened. A list filenames, one per line, is treated as if each appeared separately on the command line. Be careful of trailing white space in both include and exclude file lists.

In the case where excluded files (see **X** option) also exist, excluded files take precedence over all included files. So, if a file is specified in both the include and exclude files (or on the command line), it will be excluded.

If a file name is preceded by **-C** in a **c** (create) or **r** (replace) operation, **tar** will perform a **chdir** (see **chsh**(1)) to that file name. This allows multiple directories not related by a close common parent to be archived using short relative path names.

Note: the **-C** option only applies to *one* following directory name and *one* following file name.

If no digit or 'f' is given, the entry in **/etc/default/tar** with digit "0" will be the default.

EXAMPLES

To archive files from **/usr/include** and from **/etc**, onto default tape drive **0** one might use:

```
example% tar c -C /usr include -C /etc .
```

If you get a table of contents from the resulting *tarfile*, you might see something like:

```
include/
include/a.out.h
and all the other files in /usr/include ...
/chown
and all the other files in /etc
```

To extract all files under **include**:

```
example% tar xv include
x include/, 0 bytes, 0 tape blocks
and all files under include. . .
```

Here is a simple example using **tar** to create an archive of your home directory on a tape mounted on drive **/dev/rmt/0**:

```
example% cd
example% tar cvf /dev/rmt/0 .
messages from tar
```

The **c** option means create the archive; the **v** option makes **tar** tell you what it is doing as it works; the **f** option means that you are specifically naming the file onto which the archive should be placed (**/dev/rmt/0** in this example).

Now you can read the table of contents from the archive like this:

```
example% tar tvf /dev/rmt/0
rw-r--r--  1677/40   2123   Nov 7 18:15 1985    ./test.c
...
example%
```

The columns have the following meanings:

- column 1 is the access permissions to **./test.c**
- column 2 is the *user-id/group-id* of **./test.c**
- column 3 is the size of **./test.c** in bytes
- column 4 is the modification date of **./test.c**
- column 5 is the name of **./test.c**

You can extract files from the archive like this:

```
example% tar xvf /dev/rmt/0
messages from tar
example%
```

If there are multiple archive files on a tape, each is separated from the following one by an EOF marker. **tar** does not read the EOF mark on the tape after it finishes reading an archive file because **tar** looks for a special header to decide when it has reached the end of the archive. Now if you try to use **tar** to read the next archive file from the tape, **tar** does not know enough to skip over the EOF mark and tries to read the EOF mark as an archive instead. The result of this is an error message from **tar** to the effect:

```
tar: blocksize=0
```

This means that to read another archive from the tape, you must skip over the EOF marker before starting another **tar** command. You can accomplish this using the **mt(1)** command, as shown in the example below. Assume that you are reading from **/dev/rmt/0n**.

```
example% tar xvfp /dev/rmt/0n    read first archive from tape
messages from tar
example% mt fsf 1                skip over the end-of-file marker
example% tar xvfp /dev/rmt/0n    read second archive from tape
messages from tar
example%
```

Finally, here is an example using **tar** to transfer files across the Ethernet. First, here is how to archive files from the local machine (**example**) to a tape on a remote system (**host**):

```
example% tar cvfb - 20 filenames | rsh host dd of=/dev/rmt/0 obs=20b
messages from tar
example%
```

In the example above, we are *creating* a *tarfile* with the **c** key letter, asking for *verbose* output from **tar** with the **v** option, specifying the name of the output *tarfile* using the **f** option (the standard output is where the *tarfile* appears, as indicated by the '-' sign), and specifying the blocksize (**20**) with the **b** option. If you want to change the blocksize, you must change the blocksize arguments both on the **tar** command *and* on the **dd** command.

Now, here is how to use **tar** to get files from a tape on the remote system back to the local system:

```
example% rsh -n host dd if=/dev/rmt/0 bs=20b | tar xvBfb - 20 filenames
messages from tar
example%
```

In the example above, we are *extracting* from the *tarfile* with the **x** key letter, asking for *verbose output from tar* with the **v** option, telling **tar** it is reading from a pipe with the **B** option, specifying the name of the input *tarfile* using the **f** option (the standard input is where the *tarfile* appears, as indicated by the '-' sign), and specifying the blocksize (**20**) with the **b** option.

ENVIRONMENT

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

LC_CTYPE

Determines how **tar** handles characters. When **LC_CTYPE** is set to a valid value, **tar** can display and handle text and filenames containing valid characters for that locale. **tar** can display and handle Extended Unix code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. **tar** can also handle EUC

characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

LC_TIME

Determines how **tar** handles date and time formats. In the "C" locale, date and time handling follow the U.S. rules.

FILES

/dev/rmt/[0-7][b][n]
/dev/rmt/[0-7]l[b][n]
/dev/rmt/[0-7]m[b][n]
/dev/rmt/[0-7]h[b][n]
/dev/rmt/[0-7]u[b][n]
/dev/rmt/[0-7]c[b][n]
/etc/default/tar

Settings may look like this:

```
archive0=/dev/rmt/0
archive1=/dev/rmt/0n
archive2=/dev/rmt/1
archive3=/dev/rmt/1n
archive4=/dev/rmt/0
archive5=/dev/rmt/0n
archive6=/dev/rmt/1
archive7=/dev/rmt/1n
```

/tmp/tar*

SEE ALSO

ar(1), **chown(1)**, **cpio(1)**, **cs(1)**, **ls(1)**, **mt(1)**, **umask(2)**, **environ(5)**

DIAGNOSTICS

Complaints about bad key characters and tape read/write errors.
 Complaints if enough memory is not available to hold the link tables.

NOTES

There is no way to ask for the *n*-th occurrence of a file.

Tape errors are handled ungracefully.

The **u** option can be slow.

The **b** option should not be used with archives that are going to be updated. The current magnetic tape driver cannot backspace raw magnetic tape. If the archive is on a disk file, the **b** option should not be used at all, because updating an archive stored on disk can destroy it.

Neither the **r** option nor the **u** option can be used with quarter-inch archive tapes, since these tape drives cannot backspace.

When extracting tapes created with the **r** or **u** option, directory modification times may not be set correctly.

When using **r**, **u**, **x**, or **X**, the named files must match exactly to the corresponding files in the *tarfile*. For example, to extract *./filename*, you must specify *./filename*, and not *filename*. The **t** option displays how each file was archived.

The full pathname length cannot exceed 255 characters.

The file name (or leaf) length cannot exceed 100 characters.

The prefix of the pathname cannot exceed 155 characters.

tar does not copy empty directories or special files such as devices.

Filename substitution wildcards do not work for extracting files from the archive. To get around this, use a command of the form:

```
tar xvf... /dev/rmt/0 'tar tf... /dev/rmt/0 | grep 'pattern''
```

NAME	tbl – format tables for nroff or troff
SYNOPSIS	tbl [-ms] [<i>filename</i>] ...
AVAILABILITY	SUNWdoc
DESCRIPTION	<p>tbl is a preprocessor for formatting tables for nroff(1) or troff(1). The input <i>filenames</i> are copied to the standard output, except that lines between .TS and .TE command lines are assumed to describe tables and are reformatted.</p> <p>If no arguments are given, tbl reads the standard input, so tbl may be used as a filter. When tbl is used with eqn(1) or neqn the tbl command should be first, to minimize the volume of data passed through pipes.</p>
OPTIONS	-ms Copy the -ms macro package to the front of the output file.
EXAMPLES	<p>As an example, letting <code>\t</code> represent a TAB (which should be typed as a genuine TAB) the input</p> <pre> .TS c s s c c s c c c l n n. Household Population Town\tHouseholds \tNumber\tSize Bedminster\t789\t3.26 Bernards Twp.\t3087\t3.74 Bernardsville\t2018\t3.30 Bound Brook\t3425\t3.04 Branchburg\t1644\t3.49 Bridgewater\t7897\t3.81 Far Hills\t240\t3.19 .TE </pre>

yields

Town	Household Population	
	Number	Size
Bedminster	789	3.26
Bernards Twp.	3087	3.74
Bernardsville	2018	3.30
Bound Brook	3425	3.04
Branchburg	1644	3.49
Bridgewater	7897	3.81
Far Hills	240	3.19

SEE ALSO `eqn(1)`, `nroff(1)`, `troff(1)`

NAME	tcopy – copy a magnetic tape
SYNOPSIS	tcopy <i>source</i> [<i>destination</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	<p>tcopy copies the magnetic tape mounted on the tape drive specified by the <i>source</i> argument. The only assumption made about the contents of a tape is that there are two tape marks at the end.</p> <p>When only a source drive is specified, tcopy scans the tape, and displays information about the sizes of records and tape files. If a destination is specified, tcopy makes a copies the source tape onto the <i>destination</i> tape, with blocking preserved. As it copies, tcopy produces the same output as it does when only scanning a tape.</p>
SEE ALSO	mt (1), ioctl (2)
NOTES	tcopy will only run on systems supporting an associated set of ioctl (2) requests.

NAME	tee – replicate the standard output
SYNOPSIS	tee [-a] [-i] [<i>filename</i>]...
AVAILABILITY	SUNWcsu
DESCRIPTION	tee transcribes the standard input to the standard output and makes copies in the <i>filenames</i> .
OPTIONS	-a Append the output to the <i>filenames</i> rather than overwriting them. -i Ignore interrupts.
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of tee for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how tee behaves.</p> <p>LC_CTYPE Determines how tee handles characters. When LC_CTYPE is set to a valid value, tee can display and handle text and filenames containing valid characters for that locale. tee can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. tee can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p> <p>LC_MESSAGES Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).</p>
SEE ALSO	environ(5)

NAME	telnet – user interface to a remote system using the TELNET protocol
SYNOPSIS	telnet [<i>host</i> [<i>port</i>]]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>telnet communicates with another host using the TELNET protocol. If telnet is invoked without arguments, it enters command mode, indicated by its prompt telnet>. In this mode, it accepts and executes the commands listed below. If it is invoked with arguments, it performs an open command (see “Telnet Commands” below) with those arguments.</p> <p>Once a connection has been opened, telnet enters input mode. In this mode, text typed is sent to the remote host. The input mode entered will be either “character at a time” or “line by line” depending on what the remote system supports.</p> <p>In “character at a time” mode, most text typed is immediately sent to the remote host for processing.</p> <p>In “line by line” mode, all text is echoed locally, and (normally) only completed lines are sent to the remote host. The “local echo character” (initially ^E) may be used to turn off and on the local echo (this would mostly be used to enter passwords without the password being echoed).</p> <p>In either mode, if the <i>localchars</i> toggle is TRUE (the default in line mode; see below), the user’s quit, intr, and flush characters are trapped locally, and sent as TELNET protocol sequences to the remote side. There are options (see toggle, autoflush, and toggle, autosynch) which cause this action to flush subsequent output to the terminal (until the remote host acknowledges the TELNET sequence) and flush previous terminal input (in the case of quit and intr).</p> <p>While connected to a remote host, telnet command mode may be entered by typing the telnet escape character (initially ^]). When in command mode, the normal terminal editing conventions are available.</p>
USAGE Telnet Commands	<p>The following commands are available. Only enough of each command to uniquely identify it need be typed (this is also true for arguments to the mode, set, toggle, and display commands).</p> <p>open <i>host</i> [<i>port</i>] Open a connection to the named host. If no port number is specified, telnet will attempt to contact a TELNET server at the default port. The host specification may be either a host name (see hosts(4)) or an Internet address specified in the “dot notation” (see inet(7)).</p> <p>close Close any open TELNET session and exit telnet. An EOF (in command mode) will also close a session and exit.</p> <p>quit Same as close, above.</p>

- z** Suspend **telnet**. This command only works when the user is using a shell that supports job control, such as **sh**(1).
- mode** *type*
type is either **line** (for “line by line” mode) or **character** (for “character at a time” mode). The remote host is asked for permission to go into the requested mode. If the remote host is capable of entering that mode, the requested mode will be entered.
- status** Show the current status of **telnet**. This includes the peer one is connected to, as well as the current mode.
- display** [*argument...*]
 Display all, or some, of the **set** and **toggle** values (see **toggle**, *arguments*).
- ? [*command*]
 Get help. With no arguments, **telnet** print a help summary. If a command is specified, **telnet** will print the help information for just that command.
- send** *arguments*
 Send one or more special character sequences to the remote host. The following are the arguments which may be specified (more than one argument may be specified at a time):
- escape** Send the current **telnet** escape character (initially ^]).
 - synch** Send the TELNET **SYNCH** sequence. This sequence discards all previously typed (but not yet read) input on the remote system. This sequence is sent as TCP urgent data (and may not work if the remote system is a 4.2 BSD system — if it does not work, a lower case “r” may be echoed on the terminal).
 - brk** Send the TELNET **BRK** (Break) sequence, which may have significance to the remote system.
 - ip** Send the TELNET **IP** (Interrupt Process) sequence, which aborts the currently running process on the remote system.
 - ao** Send the TELNET **AO** (Abort Output) sequence, which flushes all output **from** the remote system **to** the user’s terminal.
 - ayt** Send the TELNET **AYT** (Are You There) sequence, to which the remote system may or may not respond.
 - ec** Send the TELNET **EC** (Erase Character) sequence, which erases the last character entered.
 - el** Send the TELNET **EL** (Erase Line) sequence, which should cause the remote system to erase the line currently being entered.
 - ga** Send the TELNET **GA** (Go Ahead) sequence, which likely has no significance to the remote system.
 - nop** Send the TELNET **NOP** (No Operation) sequence.
 - ? Print out help information for the **send** command.

set argument value

Set any one of a number of **telnet** variables to a specific value. The special value “off” turns off the function associated with the variable. The values of variables may be interrogated with the **display** command. The variables which may be specified are:

echo This is the value (initially **^E**) which, when in “line by line” mode, toggles between doing local echoing of entered characters (for normal processing), and suppressing echoing of entered characters (for example, entering a password).

escape This is the **telnet** escape character (initially **^I**) which enters **telnet** command mode (when connected to a remote system).

interrupt

If **telnet** is in **localchars** mode (see **toggle localchars**) and the **interrupt** character is typed, a TELNET **IP** sequence (see **send** and **ip**) is sent to the remote host. The initial value for the interrupt character is taken to be the terminal’s **intr** character.

quit If **telnet** is in **localchars** mode (see **toggle localchars**) and the **quit** character is typed, a TELNET **BRK** sequence (see **send**, **brk**) is sent to the remote host. The initial value for the quit character is taken to be the terminal’s **quit** character.

flushoutput

If **telnet** is in **localchars** mode (see **toggle localchars**) and the **flushoutput** character is typed, a TELNET **AO** sequence (see **send**, **ao**) is sent to the remote host. The initial value for the flush character is taken to be the terminal’s **flush** character.

erase If **telnet** is in **localchars** mode (see **toggle localchars**), and if **telnet** is operating in “character at a time” mode, then when this character is typed, a TELNET **EC** sequence (see **send**, **ec**) is sent to the remote system. The initial value for the erase character is taken to be the terminal’s **erase** character.

kill If **telnet** is in **localchars** mode (see **toggle localchars**), and if **telnet** is operating in “character at a time” mode, then when this character is typed, a TELNET **EL** sequence (see **send**, **el**) is sent to the remote system. The initial value for the kill character is taken to be the terminal’s **kill** character.

eof If **telnet** is operating in “line by line” mode, entering this character as the first character on a line sends this character to the remote system. The initial value of the eof character is taken to be the terminal’s **eof** character.

toggle arguments...

Toggle (between TRUE and FALSE) various flags that control how **telnet** responds to events. More than one argument may be specified. The state of these flags may be interrogated with the **display** command. Valid arguments are:

autoflush

If **autoflush** and **localchars** are both TRUE, then when the **ao**, **intr**, or **quit** characters are recognized (and transformed into TELNET sequences; see **set** for details), **telnet** refuses to display any data on the user's terminal until the remote system acknowledges (using a TELNET **Timing Mark** option) that it has processed those TELNET sequences. The initial value for this toggle is TRUE if the terminal user had not done an "stty nofish", otherwise FALSE (see **stty(1)**).

autosynch

If **autosynch** and **localchars** are both TRUE, then when either the **interrupt** or **quit** characters are typed (see **set** for descriptions of the **interrupt** and **quit** characters), the resulting TELNET sequence sent is followed by the TELNET **SYNCH** sequence. This procedure **should** cause the remote system to begin throwing away all previously typed input until both of the TELNET sequences have been read and acted upon. The initial value of this toggle is FALSE.

crmod Toggle RETURN mode. When this mode is enabled, most RETURN characters received from the remote host will be mapped into a RETURN followed by a line feed. This mode does not affect those characters typed by the user, only those received from the remote host. This mode is not very useful unless the remote host only sends RETURN, but never LINEFEED. The initial value for this toggle is FALSE.

debug Toggle socket level debugging (useful only to the super-user). The initial value for this toggle is FALSE.

localchars

If this is TRUE, then the **flush**, **interrupt**, **quit**, **erase**, and **kill** characters (see **set**) are recognized locally, and transformed into appropriate TELNET control sequences (respectively **ao**, **ip**, **brk**, **ec**, and **el**; see **send**). The initial value for this toggle is TRUE in "line by line" mode, and FALSE in "character at a time" mode.

netdata

Toggle the display of all network data (in hexadecimal format). The initial value for this toggle is FALSE.

options

Toggle the display of some internal **telnet** protocol processing (having to do with TELNET options). The initial value for this toggle is FALSE.

? Display the legal **toggle** commands.

SEE ALSO [rlogin\(1\)](#), [sh\(1\)](#), [stty\(1\)](#), [hosts\(4\)](#), [inet\(7\)](#)

NOTES

There is no adequate way for dealing with flow control.

On some remote systems, echo has to be turned off manually when in “line by line” mode.

There is enough settable state to justify a **file**.

In “line by line” mode, the terminal’s EOF character is only recognized (and sent to the remote system) when it is the first character on a line.

NAME	test – condition evaluation command
SYNOPSIS	<code>/usr/ucb/test</code> <i>expression</i> [<i>expression</i>]
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>test evaluates the expression <i>expression</i> and, if its value is true, sets a zero (true) exit status; otherwise, a non-zero (false) exit status is set; test also sets a non-zero exit status if there are no arguments. When permissions are tested, the effective user ID of the process is used.</p> <p>All operators, flags, and brackets (brackets used as shown in the second SYNOPSIS line) must be separate arguments to the test command; normally these items are separated by spaces.</p>
USAGE	
Primitives	<p>The following primitives are used to construct <i>expression</i>:</p> <ul style="list-style-type: none"> <code>-r filename</code> True if <i>filename</i> exists and is readable. <code>-w filename</code> True if <i>filename</i> exists and is writable. <code>-x filename</code> True if <i>filename</i> exists and is executable. <code>-f filename</code> True if <i>filename</i> exists and is a regular file. Alternatively, if <code>/usr/bin/sh</code> users specify <code>/usr/ucb</code> before <code>/usr/bin</code> in their PATH environment variable, then test will return true if <i>filename</i> exists and is (not-a-directory). This is also the default for <code>/usr/bin/csh</code> users. <code>-d filename</code> True if <i>filename</i> exists and is a directory. <code>-c filename</code> True if <i>filename</i> exists and is a character special file. <code>-b filename</code> True if <i>filename</i> exists and is a block special file. <code>-p filename</code> True if <i>filename</i> exists and is a named pipe (fifo). <code>-u filename</code> True if <i>filename</i> exists and its set-user-ID bit is set. <code>-g filename</code> True if <i>filename</i> exists and its set-group-ID bit is set. <code>-k filename</code> True if <i>filename</i> exists and its sticky bit is set. <code>-s filename</code> True if <i>filename</i> exists and has a size greater than zero. <code>-t [fildes]</code> True if the open file whose file descriptor number is <i>fildes</i> (1 by default) is associated with a terminal device. <code>-z s1</code> True if the length of string <i>s1</i> is zero. <code>-n s1</code> True if the length of the string <i>s1</i> is non-zero. <code>s1 = s2</code> True if strings <i>s1</i> and <i>s2</i> are identical. <code>s1 != s2</code> True if strings <i>s1</i> and <i>s2</i> are <i>not</i> identical. <code>s1</code> True if <i>s1</i> is <i>not</i> the null string.

n1 -eq n2 True if the integers *n1* and *n2* are algebraically equal. Any of the comparisons **-ne**, **-gt**, **-ge**, **-lt**, and **-le** may be used in place of **-eq**.

-L *filename* True if *filename* exists and is a symbolic link. With all other primitives, the symbolic links are followed by default.

Operators

These primaries may be combined with the following operators:

! Unary negation operator.

-a Binary *and* operator.

-o Binary *or* operator (**-a** has higher precedence than **-o**).

(expression) Parentheses for grouping. Notice also that parentheses are meaningful to the shell and, therefore, must be quoted.

SEE ALSO

find(1), sh(1)

NOTES

The **not-a-directory** alternative to the **-f** option is a transition aid for BSD applications and may not be supported in future releases.

The **-L** option is a migration aid for users of other shells which have similar options and may not be supported in future releases.

If you test a file you own (the **-r**, **-w**, or **-x** tests), but the permission tested does not have the *owner* bit set, a non-zero (false) exit status will be returned even though the file may have the *group* or *other* bit set for that permission. The correct exit status will be set if you are super-user.

The **=** and **!=** operators have a higher precedence than the **-r** through **-n** operators, and **=** and **!=** always expect arguments; therefore, **=** and **!=** cannot be used with the **-r** through **-n** operators.

If more than one argument follows the **-r** through **-n** operators, only the first argument is examined; the others are ignored, unless a **-a** or a **-o** is the second argument.

NAME	test – condition evaluation command
SYNOPSIS	test <i>expression</i> [<i>expression</i>]
DESCRIPTION	<p>test evaluates the expression <i>expression</i> and if its value is true, sets a zero (TRUE) exit status; otherwise, a non-zero (FALSE) exit status is set; test also sets a non-zero exit status if there are no arguments. When permissions are tested, the effective user ID of the process is used.</p> <p>All operators, flags, and brackets (brackets used as shown in the second SYNOPSIS line) must be separate arguments to test. Normally these items are separated by spaces.</p>
USAGE Primitives	<p>The following primitives are used to construct <i>expression</i>:</p> <ul style="list-style-type: none"> –r <i>filename</i> True if <i>filename</i> exists and is readable. –w <i>filename</i> True if <i>filename</i> exists and is writable. –x <i>filename</i> True if <i>filename</i> exists and is executable. –f <i>filename</i> True if <i>filename</i> exists and is a regular file. –d <i>filename</i> True if <i>filename</i> exists and is a directory. –c <i>filename</i> True if <i>filename</i> exists and is a character special file. –b <i>filename</i> True if <i>filename</i> exists and is a block special file. –p <i>filename</i> True if <i>filename</i> exists and is a named pipe (fifo). –u <i>filename</i> True if <i>filename</i> exists and its set-user-ID bit is set. –g <i>filename</i> True if <i>filename</i> exists and its set-group-ID bit is set. –k <i>filename</i> True if <i>filename</i> exists and its sticky bit is set. –s <i>filename</i> True if <i>filename</i> exists and has a size greater than zero. –t [<i>fildev</i>] True if the open file whose file descriptor number is <i>fildev</i> (1 by default) is associated with a terminal device. –z <i>s1</i> True if the length of string <i>s1</i> is zero. –n <i>s1</i> True if the length of the string <i>s1</i> is non-zero. <i>s1</i> = <i>s2</i> True if strings <i>s1</i> and <i>s2</i> are identical. <i>s1</i> != <i>s2</i> True if strings <i>s1</i> and <i>s2</i> are <i>not</i> identical. <i>s1</i> True if <i>s1</i> is <i>not</i> the null string.

n1 -eq n2 True if the integers *n1* and *n2* are algebraically equal. Any of the comparisons ***-ne***, ***-gt***, ***-ge***, ***-lt***, and ***-le*** may be used in place of ***-eq***.

Operators

These primaries may be combined with the following operators:

! Unary negation operator.

-a Binary *and* operator.

-o Binary *or* operator (**-a** has higher precedence than **-o**).

`(*expression*)` Parentheses for grouping. Notice also that parentheses are meaningful to the shell and, therefore, must be quoted.

SEE ALSO

find(1), **sh(1)**

NOTES

If you test a file you own (the **-r**, **-w**, or **-x** tests), but the permission tested does not have the *owner* bit set, a non-zero (false) exit status will be returned even though the file may have the *group* or *other* bit set for that permission. The correct exit status will be set if you are super-user.

The **=** and **!=** operators have a higher precedence than the **-r** through **-n** operators, and **=** and **!=** always expect arguments; therefore, **=** and **!=** cannot be used with the **-r** through **-n** operators.

If more than one argument follows the **-r** through **-n** operators, only the first argument is examined; the others are ignored, unless a **-a** or a **-o** is the second argument.

NAME	tftp – trivial file transfer program
SYNOPSIS	tftp [<i>host</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	tftp is the user interface to the Internet TFTP (Trivial File Transfer Protocol), which allows users to transfer files to and from a remote machine. The remote <i>host</i> may be specified on the command line, in which case tftp uses <i>host</i> as the default host for future transfers (see the connect command below).
USAGE	Once tftp is running, it issues the prompt tftp> and recognizes the following commands:
Commands	<p>connect <i>host-name</i> [<i>port</i>]</p> <p>Set the <i>host</i> (and optionally <i>port</i>) for transfers. The TFTP protocol, unlike the FTP protocol, does not maintain connections between transfers; thus, the connect command does not actually create a connection, but merely remembers what host is to be used for transfers. You do not have to use the connect command; the remote host can be specified as part of the get or put commands.</p> <p>mode <i>transfer-mode</i></p> <p>Set the mode for transfers; <i>transfer-mode</i> may be one of ascii or binary. The default is ascii.</p> <p>put <i>filename</i></p> <p>put <i>localfile remotefile</i></p> <p>put <i>filename1 filename2 . . . filenameN remote-directory</i></p> <p>Transfer a file, or a set of files, to the specified remote file or directory. The destination can be in one of two forms: a filename on the remote host if the host has already been specified, or a string of the form:</p> <p style="padding-left: 40px;"><i>host:filename</i></p> <p>to specify both a host and filename at the same time. If the latter form is used, the specified host becomes the default for future transfers. If the remote-directory form is used, the remote host is assumed to be running the UNIX system. Files may be written only if they already exist and are publicly writable (see in.tftpd(1M)).</p> <p>get <i>filename</i></p> <p>get <i>remotename localname</i></p> <p>get <i>filename1 filename2 filename3 . . . filenameN</i></p> <p>Get a file or set of files (three or more) from the specified remote <i>sources</i>. <i>source</i> can be in one of two forms: a filename on the remote host if the host has already been specified, or a string of the form:</p> <p style="padding-left: 40px;"><i>host:filename</i></p> <p>to specify both a host and filename at the same time. If the latter form is used, the last host specified becomes the default for future transfers.</p>

quit Exit **tftp**. An EOF also exits.

verbose Toggle verbose mode.

trace Toggle packet tracing.

status Show current status.

rexmt *retransmission-timeout*
Set the per-packet retransmission timeout, in seconds.

timeout *total-transmission-timeout*
Set the total transmission timeout, in seconds.

ascii Shorthand for **mode ascii**.

binary Shorthand for **mode binary**.

? [*command-name ...*]
Print help information.

NOTES

The default *transfer-mode* is **ascii**. This differs from pre-SunOS 4.0 and pre-4.3BSD systems, so explicit action must be taken when transferring non-ASCII binary files such as executable commands.

Because there is no user-login or validation within the TFTP protocol, many remote sites restrict file access in various ways. Approved methods for file access are specific to each site, and therefore cannot be documented here.

When using the **get** command to transfer multiple files from a remote host, three or more files must be specified. If two files are specified, the second file is used as a local file.

NAME	time – time a command
SYNOPSIS	time <i>command</i>
AVAILABILITY	SUNWaccu
DESCRIPTION	The <i>command</i> is executed; after it is complete, time prints the elapsed time during the command (real), the time spent in the system (sys), and the time spent in execution of the <i>command</i> (user). Times are reported in seconds or in the format <i>hours:minutes:seconds</i> . The times are printed on standard error.
EXAMPLES	<p>The two examples here show the differences between the csh version of time and the version in /usr/bin/time. These examples assume that csh is the shell in use.</p> <pre>example% time find / -name csh.1 -print /usr/share/man/man1/csh.1 95.0u 692.0s 1:17:52 16% 0+0k 0+0io 0pf+0w</pre> <p>For an explanation of the format of time output see csh(1).</p> <pre>example% /usr/bin/time find / -name csh.1 -print /usr/share/man/man1/csh.1 real 1:23:31.5 user 1:33.2 sys 11:28.2</pre>
SEE ALSO	shell_builtins (1), timex (1), times (2)
NOTES	<p>When the time command is run on a multiprocessor machine, the total of the values printed for user and sys can exceed real. This is because on a multiprocessor machine it is possible to divide the task between the various processors.</p> <p>When the command being timed is interrupted, the timing values displayed may not always be accurate.</p>
BUGS	Elapsed time is accurate to the second, while the CPU times are measured to the 100th second. Thus the sum of the CPU times can be up to a second larger than the elapsed time.

NAME	times – shell built-in function to report time usages of the current shell
SYNOPSIS	
sh	times
ksh	† times
DESCRIPTION	
sh	Print the accumulated user and system times for processes run from the shell.
ksh	Print the accumulated user and system times for the shell and for processes run from the shell. On this man page, ksh (1) commands that are preceded by one or two † (daggers) are treated specially in the following ways: <ol style="list-style-type: none">1. Variable assignment lists preceding the command remain in effect when the command completes.2. I/O redirections are processed after variable assignments.3. Errors cause a script that contains them to abort.4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.
SEE ALSO	ksh (1), sh (1), time (1)

NAME	timex – time a command; report process data and system activity
SYNOPSIS	timex [-o] [-p [-fhkmrt]] [-s] <i>command</i>
AVAILABILITY	SUNWaccu
DESCRIPTION	<p>The given <i>command</i> is executed; the elapsed time, user time and system time spent in execution are reported in seconds. Optionally, process accounting data for the <i>command</i> and all its children can be listed or summarized, and total system activity during the execution interval can be reported.</p> <p>The output of timex is written on standard error.</p>
OPTIONS	<p>-o Report the total number of blocks read or written and total characters transferred by <i>command</i> and all its children. This option works only if the process accounting software is installed.</p> <p>-p List process accounting records for <i>command</i> and all its children. This option works only if the process accounting software is installed. Suboptions f, h, k, m, r, and t modify the data items reported. The options are as follows:</p> <ul style="list-style-type: none"> -f Print the fork(2)/exec(2) flag and system exit status columns in the output. -h Instead of mean memory size, show the fraction of total available CPU time consumed by the process during its execution. This “hog factor” is computed as (total CPU time)/(elapsed time). -k Instead of memory size, show total kcore-minutes. -m Show mean core size (the default). -r Show CPU factor (user time/(system-time + user-time)). -t Show separate system and user CPU times. The number of blocks read or written and the number of characters transferred are always reported. <p>-s Report total system activity (not just that due to <i>command</i>) that occurred during the execution interval of <i>command</i>. All the data items listed in sar(1) are reported.</p>
EXAMPLES	<p>A simple example:</p> <pre>example% timex -ops sleep 60</pre> <p>A terminal session of arbitrary complexity can be measured by timing a sub-shell:</p> <pre>example% timex -opskmt sh session commands EOT</pre>

SEE ALSO **sar(1), time(1), times(2)**

NOTES Process records associated with *command* are selected from the accounting file */var/adm/pacct* by inference, since process genealogy is not available. Background processes having the same user ID, terminal ID, and execution time window will be spuriously included.

NAME	tip – connect to remote system
SYNOPSIS	tip [-v] [-speed-entry] <i>hostname</i> <i>phone-number</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>tip establishes a full-duplex terminal connection to a remote host. Once the connection is established, a remote session using tip behaves like an interactive session on a local terminal.</p> <p>The <i>remote</i> file contains entries describing remote systems and line speeds used by tip. Each host has a default baud rate for the connection, or you can specify a speed with the -speed-entry command line argument.</p> <p>When <i>phone-number</i> is specified, tip looks for an entry in the <i>remote</i> file of the form:</p> <p style="padding-left: 40px;">tip -speed-entry</p> <p>When it finds such an entry, it sets the connection speed accordingly. If it finds no such entry, tip interprets -speed-entry as if it were a system name, resulting in an error message.</p> <p>If you omit -speed-entry, tip uses the tip0 entry to set a speed for the connection.</p> <p>When establishing the connection tip sends a connection message to the remote system. The default value for this message can be found in the <i>remote</i> file.</p> <p>When tip attempts to connect to a remote system, it opens the associated device with an exclusive-open ioctl(2) call. Thus only one user at a time may access a device. This is to prevent multiple processes from sampling the terminal line. In addition, tip honors the locking protocol used by uucp(1C).</p> <p>When tip starts up it reads commands from the file .tiprc in your home directory.</p>
OPTIONS	-v Display commands from the .tiprc file as they are executed.
USAGE	<p>Typed characters are normally transmitted directly to the remote machine (which does the echoing as well).</p> <p>At any time that tip prompts for an argument (for example, during setup of a file transfer) the line typed may be edited with the standard erase and kill characters. A null line in response to a prompt, or an interrupt, aborts the dialogue and returns you to the remote machine.</p>
Commands	<p>A tilde (~) appearing as the first character of a line is an escape signal which directs tip to perform some special action. tip recognizes the following escape sequences:</p> <p>^^D ~. Drop the connection and exit (you may still be logged in on the remote machine). ~c [<i>name</i>] Change directory to <i>name</i> (no argument implies change to your home directory).</p>

- ~!** Escape to an interactive shell on the local machine (exiting the shell returns you to **tip**).
- ~>** Copy file from local to remote.
- ~<** Copy file from remote to local.
- ~p** *from* [*to*]
Send a file to a remote host running the UNIX system. When you use the put command, the remote system runs the command string
- ```
cat > to
```
- while **tip** sends it the *from* file. If the *to* file is not specified, the *from* file name is used. This command is actually a UNIX-system-specific version of the “~>” command.
- ~t** *from* [ *to* ]  
Take a file from a remote host running the UNIX system. As in the put command the *to* file defaults to the *from* file name if it is not specified. The remote host executes the command string
- ```
cat from ; echo ^A
```
- to send the file to **tip**.
- ~|** Pipe the output from a remote command to a local process. The command string sent to the local system is processed by the shell.
- ~C** Connect a program to the remote machine. The command string sent to the program is processed by the shell. The program inherits file descriptors 0 as remote line input, 1 as remote line output, and 2 as tty standard error.
- ~\$** Pipe the output from a local process to the remote host. The command string sent to the local system is processed by the shell.
- ~#** Send a BREAK to the remote system.
- ~s** Set a variable (see the discussion below).
- ~~Z** Stop **tip** (only available when run under a shell that supports job control, such as the C shell).
- ~~Y** Stop only the “local side” of **tip** (only available when run under a shell that supports job control, such as the C shell); the “remote side” of **tip**, the side that displays output from the remote host, is left running.
- ~?** Get a summary of the tilde escapes.

Copying files requires some cooperation on the part of the remote host. When a ~> or ~< escape is used to send a file, **tip** prompts for a file name (to be transmitted or received) and a command to be sent to the remote system, in case the file is being transferred from the remote system. While **tip** is transferring a file the number of lines transferred will be continuously displayed on the screen. A file transfer may be aborted with an interrupt.

Auto-call Units

tip may be used to dial up remote systems using a number of auto-call unit's (ACU's). When the remote system description contains the **du** capability, **tip** uses the call-unit (**cu**), ACU type (**at**), and phone numbers (**pn**) supplied. Normally **tip** displays verbose messages as it dials.

Depending on the type of auto-dialer being used to establish a connection the remote host may have garbage characters sent to it upon connection. The user should never assume that the first characters typed to the foreign host are the first ones presented to it. The recommended practice is to immediately type a **kill** character upon establishing a connection (most UNIX systems either support **@** or **CTRL-U** as the initial kill character).

tip currently supports the Ventel MD-212+ modem and DC Hayes-compatible modems.

When **tip** initializes a Hayes-compatible modem for dialing, it sets up the modem to auto-answer. Normally, after the conversation is complete, **tip** drops DTR, which causes the modem to "hang up."

Most modems can be configured such that when DTR drops, they re-initialize themselves to a preprogrammed state. This can be used to reset the modem and disable auto-answer, if desired.

Additionally, it is possible to start the phone number with a Hayes **S** command so that you can configure the modem before dialing. For example, to disable auto-answer, set up all the phone numbers in */etc/remote* using something like **pn=S0=0DT5551212**. The **S0=0** disables auto-answer.

Remote Host Description

Descriptions of remote hosts are normally located in the system-wide file */etc/remote*. However, a user may maintain personal description files (and phone numbers) by defining and exporting the *REMOTE* shell variable. The *remote* file must be readable by **tip**, but a secondary file describing phone numbers may be maintained readable only by the user. This secondary phone number file is */etc/phones*, unless the shell variable *PHONES* is defined and exported. The phone number file contains lines of the form:

system-name phone-number

Each phone number found for a system is tried until either a connection is established, or an end of file is reached. Phone numbers are constructed from **'0123456789--*'**, where the **'=**' and **'*'** are used to indicate a second dial tone should be waited for (ACU dependent).

tip Internal Variables

tip maintains a set of variables which are used in normal operation. Some of these variables are read-only to normal users (root is allowed to change anything of interest). Variables may be displayed and set through the **~s** escape. The syntax for variables is patterned after **vi(1)** and **mail(1)**. Supplying **all** as an argument to the **~s** escape displays all variables that the user can read. Alternatively, the user may request display of a particular variable by attaching a **?** to the end. For example **"~s escape?"** displays the current escape character.

Variables are numeric (num), string (str), character (char), or Boolean (bool) values. Boolean variables are set merely by specifying their name. They may be reset by prepending a **!** to the name. Other variable types are set by appending an **=** and the

value. The entire assignment must not have any blanks in it. A single set command may be used to interrogate as well as set a number of variables.

Variables may be initialized at run time by placing set commands (without the `~s` prefix) in a `.tiprc` file in one's home directory. The `-v` option makes `tip` display the sets as they are made. Comments preceded by a `#` sign can appear in the `.tiprc` file.

Finally, the variable names must either be completely specified or an abbreviation may be given. The following list details those variables known to `tip`.

beautify

(bool) Discard unprintable characters when a session is being scripted; abbreviated `be`. If the `nb` capability is present, `beautify` is initially set to `off`; otherwise, `beautify` is initially set to `on`.

baudrate

(num) The baud rate at which the connection was established; abbreviated `ba`. If a baud rate was specified on the command line, `baudrate` is initially set to the specified value; otherwise, if the `br` capability is present, `baudrate` is initially set to the value of that capability; otherwise, `baudrate` is set to 300 baud. Once `tip` has been started, `baudrate` can only be changed by the super-user.

dialtimeout

(num) When dialing a phone number, the time (in seconds) to wait for a connection to be established; abbreviated `dial`. `dialtimeout` is initially set to 60 seconds, and can only be changed by the super-user.

disconnect

(str) The string to send to the remote host to disconnect from it; abbreviated `di`. If the `di` capability is present, `disconnect` is initially set to the value of that capability; otherwise, `disconnect` is set to a null string (`""`).

echocheck

(bool) Synchronize with the remote host during file transfer by waiting for the echo of the last character transmitted; abbreviated `ec`. If the `ec` capability is present, `echocheck` is initially set to `on`; otherwise, `echocheck` is initially set to `off`.

eofread

(str) The set of characters which signify an end-of-transmission during a `~<` file transfer command; abbreviated `eofr`. If the `ie` capability is present, `eofread` is initially set to the value of that capability; otherwise, `eofread` is set to a null string (`""`).

eofwrite

(str) The string sent to indicate end-of-transmission during a `~>` file transfer command; abbreviated `eofw`. If the `oe` capability is present, `eofread` is initially set to the value of that capability; otherwise, `eofread` is set to a null string (`""`).

- eol** (str) The set of characters which indicate an end-of-line. **tip** will recognize escape characters only after an end-of-line. If the **el** capability is present, **eol** is initially set to the value of that capability; otherwise, **eol** is set to a null string ("").
- escape** (char) The command prefix (escape) character; abbreviated **es**. If the **es** capability is present, **escape** is initially set to the value of that capability; otherwise, **escape** is set to '~'.
- etimeout** (num) The amount of time, in seconds, that **tip** should wait for the echo-check response when **echocheck** is set; abbreviated **et**. If the **et** capability is present, **etimeout** is initially set to the value of that capability; otherwise, **etimeout** is set to 10 seconds.
- exceptions** (str) The set of characters which should not be discarded due to the beautification switch; abbreviated **ex**. If the **ex** capability is present, **exceptions** is initially set to the value of that capability; otherwise, **exceptions** is set to '\t\n\f\b'.
- force** (char) The character used to force literal data transmission; abbreviated **fo**. If the **fo** capability is present, **force** is initially set to the value of that capability; otherwise, **force** is set to \377 (which disables it).
- framesize** (num) The amount of data (in bytes) to buffer between file system writes when receiving files; abbreviated **fr**. If the **fs** capability is present, **framesize** is initially set to the value of that capability; otherwise, **framesize** is set to 1024.
- halfduplex** (bool) Do local echoing because the host is half-duplex; abbreviated **hdx**. If the **hd** capability is present, **halfduplex** is initially set to **on**; otherwise, **halfduplex** is initially set to **off**.
- hardwareflow** (bool) Do hardware flow control; abbreviated **hf**. If the **hf** capability is present, **hardwareflow** is initially set to **on**; otherwise, **hardwareflowcontrol** is initially set to **off**.
- host** (str) The name of the host to which you are connected; abbreviated **ho**. **host** is permanently set to the name given on the command line or in the HOST environment variable.
- localecho** (bool) A synonym for **halfduplex**; abbreviated **le**.
- log** (str) The name of the file to which to log information about outgoing phone calls. **log** is initially set to **/var/adm/aculog**, and can only be inspected or changed by the super-user.

- parity** (str) The parity to be generated and checked when talking to the remote host; abbreviated **par**. The possible values are:
- none**
 - zero** Parity is not checked on input, and the parity bit is set to zero on output.
 - one** Parity is not checked on input, and the parity bit is set to one on output.
 - even** Even parity is checked for on input and generated on output.
 - odd** Odd parity is checked for on input and generated on output.
- If the **pa** capability is present, **parity** is initially set to the value of that capability; otherwise, **parity** is set to **none**.
- phones**
The file in which to find hidden phone numbers. If the environment variable **PHONES** is set, **phones** is set to the value of **PHONES**; otherwise, **phones** is set to **/etc/phones**. The value of **phones** cannot be changed from within **tip**.
- prompt**
(char) The character which indicates an end-of-line on the remote host; abbreviated **pr**. This value is used to synchronize during data transfers. The count of lines transferred during a file transfer command is based on receipt of this character. If the **pr** capability is present, **prompt** is initially set to the value of that capability; otherwise, **prompt** is set to **\n**.
- raise** (bool) Upper case mapping mode; abbreviated **ra**. When this mode is enabled, all lower case letters will be mapped to upper case by **tip** for transmission to the remote machine. If the **ra** capability is present, **raise** is initially set to **on**; otherwise, **raise** is initially set to **off**.
- raisechar**
(char) The input character used to toggle upper case mapping mode; abbreviated **rc**. If the **rc** capability is present, **raisechar** is initially set to the value of that capability; otherwise, **raisechar** is set to **\377** (which disables it).
- rawftp** (bool) Send all characters during file transfers; do not filter non-printable characters, and do not do translations like **\n** to **\r**. Abbreviated **raw**. If the **rw** capability is present, **rawftp** is initially set to **on**; otherwise, **rawftp** is initially set to **off**.
- record** (str) The name of the file in which a session script is recorded; abbreviated **rec**. If the **re** capability is present, **record** is initially set to the value of that capability; otherwise, **record** is set to **tip.record**.
- remote**
The file in which to find descriptions of remote systems. If the environment variable **REMOTE** is set, **remote** is set to the value of **REMOTE**; otherwise, **remote** is set to **/etc/remote**. The value of **remote** cannot be changed from within **tip**.
- script** (bool) Session scripting mode; abbreviated **sc**. When **script** is **on**, **tip** will record everything transmitted by the remote machine in the script record file specified in **record**. If the **beautify** switch is on, only printable ASCII characters will be

included in the script file (those characters between 040 and 0177). The variable **exceptions** is used to indicate characters which are an exception to the normal beautification rules. If the **sc** capability is present, **script** is initially set to **on**; otherwise, **script** is initially set to **off**.

tabexpand

(bool) Expand TAB characters to SPACE characters during file transfers; abbreviated **tab**. When **tabexpand** is **on**, each tab is expanded to 8 SPACE characters. If the **tb** capability is present, **tabexpand** is initially set to **on**; otherwise, **tabexpand** is initially set to **off**.

tandem

(bool) Use XON/XOFF flow control to limit the rate that data is sent by the remote host; abbreviated **ta**. If the **nt** capability is present, **tandem** is initially set to **off**; otherwise, **tandem** is initially set to **on**.

verbose

(bool) Verbose mode; abbreviated **verb**; When verbose mode is enabled, **tip** prints messages while dialing, shows the current number of lines transferred during a file transfer operations, and more. If the **nv** capability is present, **verbose** is initially set to **off**; otherwise, **verbose** is initially set to **on**.

SHELL (str) The name of the shell to use for the **!** command; default value is **/bin/sh**, or taken from the environment.

HOME (str) The home directory to use for the **~c** command; default value is taken from the environment.

EXAMPLES

An example of the dialogue used to transfer files is given below.

```

arpa% tip monet
[connected]
...(assume we are talking to a UNIX system)...
ucbmonet login: sam
Password:
monet% cat > sylvester.c
~> Filename: sylvester.c
32 lines transferred in 1 minute 3 seconds
monet%
monet% ~< Filename: reply.c
List command for remote host: cat reply.c
65 lines transferred in 2 minutes
monet%
...(or, equivalently)...
monet% ~p sylvester.c
...(actually echoes as [put] sylvester.c)...
32 lines transferred in 1 minute 3 seconds
monet%
monet% ~t reply.c
...(actually echoes as [take] reply.c)...

```

65 lines transferred in 2 minutes

monet%

... (to print a file locally)...

monet% ~ | Local command: pr -h sylvester.c | lpr

List command for remote host: cat sylvester.c

monet% ^^D

[EOT]

... (back on the local system)...

ENVIRONMENT

The following environment variables are read by **tip**.

REMOTE The location of the *remote* file.

PHONES The location of the file containing private phone numbers.

HOST A default host to connect to.

HOME One's log-in directory (for chdirs).

SHELL The shell to fork on a '~!' escape.

FILES

/etc/phones

/etc/remote

/var/spool/locks/LCK ..* lock file to avoid conflicts with UUCP

/var/adm/aculog file in which outgoing calls are logged

~/tiprc initialization file

SEE ALSO

cu(1C), mail(1), uucp(1C), vi(1), ioctl(2)

BUGS

There are two additional variables **chardelay** and **linedelay** that are currently not implemented.

NAME	touch – update access time and/or modification time of a file
SYNOPSIS	touch [-amc] [<i>mmddhhmm</i> [<i>yy</i>]] <i>filename</i> ...
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>touch causes the access and modification times of each argument to be updated. The file name is created if it does not exist. If no time is specified (see date(1)) the current time is used.</p> <p>The return code from touch is the number of files for which the times could not be successfully modified (including files that did not exist and were not created).</p>
OPTIONS	<p>-a touch updates only the access time respectively (default is -am).</p> <p>-m touch updates only the modification time.</p> <p>-c Silently prevents touch from creating the file if it did not previously exist.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of touch for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how touch behaves.</p> <p>LC_CTYPE</p> <p>Determines how touch handles characters. When LC_CTYPE is set to a valid value, touch can display and handle text and filenames containing valid characters for that locale. touch can display and handle Extended Unix code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. touch can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p> <p>LC_MESSAGES</p> <p>Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).</p>
SEE ALSO	date (1), utime (2), environ (5)
NOTES	Users familiar with the BSD environment will find that the -f option is accepted, but ignored. The -f option is unnecessary since touch will succeed for all files owned by the user regardless of the permissions on the files.

NAME	tput – initialize a terminal or query terminfo database
SYNOPSIS	<pre>tput [-T<i>type</i>] <i>capname</i> [<i>parms</i> . . .] tput [-T<i>type</i>] init tput [-T<i>type</i>] reset tput [-T<i>type</i>] longname tput -S <<</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>tput uses the terminfo database to make the values of terminal-dependent capabilities and information available to the shell (see sh(1)), to initialize or reset the terminal, or return the long name of the requested terminal type. tput outputs a string if the attribute (<i>capability name</i>) is of type string, or an integer if the attribute is of type integer. If the attribute is of type boolean, tput simply sets the exit code (0 for TRUE if the terminal has the capability, 1 for FALSE if it does not), and produces no output. Before using a value returned on standard output, the user should test the exit code (\$? , see sh(1)) to be sure it is 0. (See the EXIT CODES and DIAGNOSTICS sections.) For a complete list of capabilities and the <i>capname</i> associated with each, see terminfo(4).</p>
OPTIONS	<p>-T<i>type</i> Indicates the <i>type</i> of terminal. Normally this option is unnecessary, because the default is taken from the environment variable TERM. If -T is specified, then the shell variables LINES and COLUMNS and the layer size will not be referenced.</p> <p><i>capname</i> Indicates the attribute from the terminfo database.</p> <p><i>parms</i> If the attribute is a string that takes parameters, the arguments <i>parms</i> will be instantiated into the string. An all numeric argument will be passed to the attribute as a number.</p> <p>-S Allows more than one capability per invocation of tput. The capabilities must be passed to tput from the standard input instead of from the command line (see the example in the EXAMPLES section). Only one <i>capname</i> is allowed per line. The -S option changes the meaning of the 0 and 1 boolean and string exit codes (see the EXIT CODES section).</p> <p>init If the terminfo database is present and an entry for the user's terminal exists (see -T<i>type</i>, above), the following will occur:</p> <ol style="list-style-type: none"> (1) if present, the terminal's initialization strings will be output (is1, is2, is3, if, iprog), (2) any delays (for instance, newline) specified in the entry will be set in the tty driver, (3) tabs expansion will be turned on or off according to the specification in the entry, and

- (4) if tabs are not expanded, standard tabs will be set (every 8 spaces). If an entry does not contain the information needed for any of the four above activities, that activity will silently be skipped.

reset Instead of putting out initialization strings, the terminal's reset strings will be output if present (**rs1**, **rs2**, **rs3**, **rf**). If the reset strings are not present, but initialization strings are, the initialization strings will be output. Otherwise, **reset** acts identically to **init**.

longname If the **terminfo** database is present and an entry for the user's terminal exists (see **-Ttype** above), then the long name of the terminal will be put out. The long name is the last name in the first line of the terminal's description in the **terminfo** database (see **term(5)**).

EXIT CODES

If *capname* is of type boolean, a value of **0** is set for TRUE and **1** for FALSE unless the **-S** option is used.

If *capname* is of type string, a value of **0** is set if the *capname* is defined for this terminal *type* (the value of *capname* is returned on standard output); a value of **1** is set if *capname* is not defined for this terminal *type* (a null value is returned on standard output).

If *capname* is of type boolean or string and the **-S** option is used, a value of **0** is returned to indicate that all lines were successful. No indication of which line failed can be given so exit code **1** will never appear. Exit codes **2**, **3**, and **4** retain their usual interpretation.

If *capname* is of type integer, a value of **0** is always set, whether or not *capname* is defined for this terminal *type*. To determine if *capname* is defined for this terminal *type*, the user must test the value of standard output. A value of **-1** means that *capname* is not defined for this terminal *type*.

Any other exit code indicates an error; see the **DIAGNOSTICS** section.

EXAMPLES

This example initializes the terminal according to the type of terminal in the environment variable **TERM**. This command should be included in everyone's .profile after the environment variable **TERM** has been exported, as illustrated on the **profile(4)** manual page.

```
example% tput init
```

The next example resets an AT&T 5620 terminal, overriding the type of terminal in the environment variable **TERM**.

```
example% tput -T5620 reset
```

The following example sends the sequence to move the cursor to row **0**, column **0** (the upper left corner of the screen, usually known as the "home" cursor position).

```
example% tput cup 0 0
```

The next example echos the clear-screen sequence for the current terminal.

```
example% tput clear
```

The next command prints the number of columns for the current terminal.

```
example% tput cols
```

The following command prints the number of columns for the 450 terminal.

```
example% tput -T450 cols
```

The next example sets the shell variables **bold**, to begin stand-out mode sequence, and **offbold**, to end standout mode sequence, for the current terminal. This might be followed by a prompt:

```
echo "${bold}Please type in your name: ${offbold}\c"
```

```
example% bold='tput smso'
```

```
example% offbold='tput rmso'
```

This example sets the exit code to indicate if the current terminal is a hardcopy terminal.

```
example% tput hc
```

This next example sends the sequence to move the cursor to row **23**, column **4**.

```
example% tput cup 23 4
```

The next command prints the long name from the **terminfo** database for the type of terminal specified in the environment variable **TERM**.

```
example% tput longname
```

This last example shows tput processing several capabilities in one invocation. This example clears the screen, moves the cursor to position 10, 10 and turns on bold (extra bright) mode. The list is terminated by an exclamation mark (!) on a line by itself.

```
example% tput -S <<!
> clear
> cup 10 10
> bold
> !
```

FILES	/usr/include/curses.h	curses(3X) header
	/usr/include/term.h	terminfo header
	/usr/lib/tabset/*	tab settings for some terminals, in a format appropriate to be output to the terminal (escape sequences that set margins and tabs); for more information, see the "Tabs and Initialization" section of terminfo(4)
	/usr/share/lib/terminfo/?/*	compiled terminal description database

SEE ALSO [clear\(1\)](#), [stty\(1\)](#), [tabs\(1\)](#), [profile\(4\)](#), [terminfo\(4\)](#)

DIAGNOSTICS

tput prints the following error messages and sets the corresponding exit codes.

exit code	error message
0	-1 (<i>capname</i> is a numeric variable that is not specified in the terminfo database for this terminal type, for instance, tput -T450 lines and tput -T2621 xmc).
1	no error message is printed, see the EXIT CODES section.
2	usage error.
3	unknown terminal <i>type</i> or no terminfo database.
4	unknown terminfo capability <i>capname</i> .

NAME	tr – translate characters
SYNOPSIS	tr [-c s] [<i>string1</i> [<i>string2</i>]]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>tr copies the standard input to the standard output with substitution or deletion of selected characters. Input characters found in <i>string1</i> are mapped into the corresponding characters of <i>string2</i>.</p> <p>When <i>string2</i> (with any repetitions of characters) contains fewer characters than <i>string1</i>, characters in <i>string1</i> with no corresponding character in <i>string2</i> are not translated.</p>
OPTIONS	<p>Any combination of the options -cs may be used:</p> <ul style="list-style-type: none"> -c Complements the set of characters in <i>string1</i> with respect to the universe of characters whose ASCII codes are 001 through 377 octal. -d Deletes all input characters in <i>string1</i>. -s Squeezes all strings of repeated output characters that are in <i>string2</i> to single characters. <p>The following abbreviation conventions may be used to introduce ranges of characters or repeated characters into the strings:</p> <ul style="list-style-type: none"> [a-z] Stands for the string of characters whose ASCII codes run from character a to character z, inclusive. [a*n] Stands for <i>n</i> repetitions of a. If the first digit of <i>n</i> is 0, <i>n</i> is considered octal; otherwise, <i>n</i> is taken to be decimal. A zero or missing <i>n</i> is taken to be huge; this facility is useful for padding <i>string2</i>. <p>The escape character \ may be used as in the shell to remove special meaning from any character in a string. In addition, \ followed by 1, 2, or 3 octal digits stands for the character whose ASCII code is given by those digits.</p>
EXAMPLES	<p>The following example creates a list of all the words in <i>filename1</i> one per line in <i>filename2</i>, where a word is taken to be a maximal string of alphabetic characters. The strings are quoted to protect the special characters from interpretation by the shell; 012 is the ASCII code for newline.</p> <pre style="margin-left: 40px;">example% tr -cs "[A-Z][a-z]" "[\012*]" <filename1> filename2</pre>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of tr for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how tr behaves.</p>

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

ed(1), **sh(1)**, **ascii(5)**, **environ(5)**

NOTES

Will not handle ASCII NUL in *string1* or *string2*; always deletes NUL from input.

NAME	tr – translate characters
SYNOPSIS	<code>/usr/ucb/tr [-c ds] [string1 [string2]]</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>tr copies the standard input to the standard output with substitution or deletion of selected characters. The arguments <i>string1</i> and <i>string2</i> are considered sets of characters. Any input character found in <i>string1</i> is mapped into the character in the corresponding position within <i>string2</i>. When <i>string2</i> is short, it is padded to the length of <i>string1</i> by duplicating its last character.</p> <p>In either string the notation: $a-b$</p> <p>denotes a range of characters from <i>a</i> to <i>b</i> in increasing ASCII order. The character <code>\</code>, followed by 1, 2 or 3 octal digits stands for the character whose ASCII code is given by those digits. As with the shell, the escape character <code>\</code>, followed by any other character, escapes any special meaning for that character.</p>
OPTIONS	<p>Any combination of the options <code>-c</code>, <code>-d</code>, or <code>-s</code> may be used:</p> <ul style="list-style-type: none"> <code>-c</code> Complement the set of characters in <i>string1</i> with respect to the universe of characters whose ASCII codes are 01 through 0377 octal. <code>-d</code> Delete all input characters in <i>string1</i>. <code>-s</code> Squeeze all strings of repeated output characters that are in <i>string2</i> to single characters.
EXAMPLES	<p>The following example creates a list of all the words in <i>filename1</i> one per line in <i>filename2</i>, where a word is taken to be a maximal string of alphabets. The second string is quoted to protect <code>'\'</code> from the shell. 012 is the ASCII code for NEWLINE.</p> <pre>example% tr -cs A-Za-z '\012' < filename1 > filename2</pre>
SEE ALSO	ed(1), ascii(5)
NOTES	Will not handle ASCII NUL in <i>string1</i> or <i>string2</i> . tr always deletes NUL from input.

NAME	trap, onintr – shell built-in functions to respond to (hardware) signals
SYNOPSIS	
sh	trap [<i>argument</i>] [<i>n</i>] ...
csh	onintr [- <i>label</i>]
ksh	† trap [<i>arg</i>] [<i>sig</i>] ...
DESCRIPTION	
sh	The trap command <i>argument</i> is to be read and executed when the shell receives numeric or symbolic signal(s) (<i>n</i>). (Note: <i>argument</i> is scanned once when the trap is set and once when the trap is taken.) Trap commands are executed in order of signal number or corresponding symbolic names. Any attempt to set a trap on a signal that was ignored on entry to the current shell is ineffective. An attempt to trap on signal 11 (memory fault) produces an error. If <i>argument</i> is absent all trap(s) <i>n</i> are reset to their original values. If <i>argument</i> is the null string this signal is ignored by the shell and by the commands it invokes. If <i>n</i> is 0 the command <i>argument</i> is executed on exit from the shell. The trap command with no arguments prints a list of commands associated with each signal number.
csh	onintr controls the action of the shell on interrupts. With no arguments, onintr restores the default action of the shell on interrupts. (The shell terminates shell scripts and returns to the terminal command input level). With the - argument, the shell ignores all interrupts. With a <i>label</i> argument, the shell executes a goto <i>label</i> when an interrupt is received or a child process terminates because it was interrupted.
ksh	trap uses <i>arg</i> as a command to be read and executed when the shell receives signal(s) <i>sig</i> . (Note that <i>arg</i> is scanned once when the trap is set and once when the trap is taken.) Each <i>sig</i> can be given as a number or as the name of the signal. trap commands are executed in order of signal number. Any attempt to set a trap on a signal that was ignored on entry to the current shell is ineffective. If <i>arg</i> is omitted or is -, then the trap(s) for each <i>sig</i> are reset to their original values. If <i>arg</i> is the NULL (the empty string, e.g., "") string then this signal is ignored by the shell and by the commands it invokes. If <i>sig</i> is ERR then <i>arg</i> will be executed whenever a command has a non-zero exit status. If <i>sig</i> is DEBUG then <i>arg</i> will be executed after each command. If <i>sig</i> is 0 or EXIT and the trap statement is executed inside the body of a function, then the command <i>arg</i> is executed after the function completes. If <i>sig</i> is 0 or EXIT for a trap set outside any function then the command <i>arg</i> is executed on exit from the shell. The trap command with no arguments prints a list of commands associated with each signal number.

On this man page, **ksh**(1) commands that are preceded by one or two † (daggers) are treated specially in the following ways:

1. Variable assignment lists preceding the command remain in effect when the command completes.
2. I/O redirections are processed after variable assignments.
3. Errors cause a script that contains them to abort.
4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.

SEE ALSO **csh**(1), **exit**(1), **ksh**(1), **sh**(1)

NAME	troff – typeset or format documents
SYNOPSIS	troff [-a] [-f] [-Fdir] [-i] [-mname] [-nN] [-olist] [-raN] [-sN] [-Tdest] [-uN] [-z] [<i>filename</i>] ...
AVAILABILITY	SUNWdoc
DESCRIPTION	<p>troff formats text in the <i>filenames</i> for typesetting or laser printing. Input to troff is expected to consist of text interspersed with formatting requests and macros. If no <i>filename</i> argument is present, troff reads standard input. A minus sign (-) as a <i>filename</i> indicates that standard input should be read at that point in the list of input files.</p> <p>The following options may appear in any order, but all must appear before the first <i>filename</i>.</p> <ul style="list-style-type: none"> -a Send an ASCII approximation of formatted output to standard output. -f Do not print a trailer after the final page of output or cause the postprocessor to relinquish control of the device. -Fdir Search directory <i>dir</i> for font width or terminal tables instead of the system default directory. -i Read standard input after all input files are exhausted. -mname Prepend the macro file <code>/usr/share/lib/tmac/name</code> to the input <i>filenames</i>. Note: most references to macro packages include the leading <i>m</i> as part of the name; for example, the man(5) macros reside in <code>/usr/share/lib/tmac/an</code>. The macro directory can be changed by setting the TROFFMACS environment variable to a specific path. Be certain to include the trailing <code>'/'</code> (slash) at the end of the path. -nN Number the first generated page <i>N</i>. -olist Print only pages whose page numbers appear in the comma-separated <i>list</i> of numbers and ranges. A range <i>N-M</i> means pages <i>N</i> through <i>M</i>; an initial -N means from the beginning to page <i>N</i>; and a final N- means from <i>N</i> to the end. -q Quiet mode in nroff; ignored in troff. -raN Set register <i>a</i> (one-character names only) to <i>N</i>. -sN Stop the phototypesetter every <i>N</i> pages. On some devices, troff produces a trailer so you can change cassettes; resume by pressing the typesetter's start button. -Tdest Prepare output for typesetter <i>dest</i>. The following values can be supplied for <i>dest</i>: <ul style="list-style-type: none"> post A PostScript printer; this is the default value. aps Autologic APS-5. -uN Set the emboldening factor for the font mounted in position 3 to <i>N</i>. If <i>N</i> is missing, then set the emboldening factor to 0. -z Suppress formatted output. Only diagnostic messages and messages output using the .tm request are output.

FILES	/tmp/trtmp temporary file
	/usr/share/lib/tmac/* standard macro files
	/usr/lib/font/* font width tables for alternate mounted troff fonts
	/usr/share/lib/nterm/* terminal driving tables for nroff
SEE ALSO	checknr(1), col(1), dpost(1), eqn(1), lp(1), man(1), nroff(1), tbl(1), man(5), me(5), ms(5)
NOTES	troff is not 8-bit clean because it is by design based on 7-bit ASCII.

NAME	true, false – provide truth values
SYNOPSIS	true false
AVAILABILITY	SUNWcsu
DESCRIPTION	true does nothing, successfully. false does nothing, unsuccessfully. They are typically used in a shell script sh as: <pre> while true do <i>command</i> done</pre>
SEE ALSO	sh (1)
DIAGNOSTICS	true has exit status zero. false has non-zero exit status.

NAME	truss – trace system calls and signals
SYNOPSIS	truss [-fcaeil] [-tvx] [!]syscall... [-s [!]signal...] [-m [!]fault...] [-rw] [!]fd... [-o outfile] <i>command</i> -p <i>pid</i>
DESCRIPTION	truss executes the specified command and produces a trace of the system calls it performs, the signals it receives, and the machine faults it incurs. Each line of the trace output reports either the fault or signal name or the system call name with its arguments and return value(s). System call arguments are displayed symbolically when possible using defines from relevant system headers; for any pathname pointer argument, the pointed-to string is displayed. Error returns are reported using the error code names described in intro(2) .
OPTIONS	<p>The following options are recognized. For those options that take a list argument, the name all can be used as a shorthand to specify all possible members of the list. If the list begins with a !, the meaning of the option is negated (for example, exclude rather than trace). Multiple occurrences of the same option may be specified. For the same name in a list, subsequent options (those to the right) override previous ones (those to the left).</p> <p>-p Interpret the <i>command</i> arguments to truss as a list of process-ids for existing processes (see ps(1)) rather than as a command to be executed. truss takes control of each process and begins tracing it provided that the userid and groupid of the process match those of the user or that the user is a privileged user. Processes may also be specified by their names in the /proc directory, for example, /proc/12345.</p> <p>-f Follow all children created by fork() or vfork() and include their signals, faults, and system calls in the trace output. Normally, only the first-level command or process is traced. When -f is specified, the process-id is included with each line of trace output to indicate which process executed the system call or received the signal.</p> <p>-c Count traced system calls, faults, and signals rather than displaying the trace line-by-line. A summary report is produced after the traced command terminates or when truss is interrupted. If -f is also specified, the counts include all traced system calls, faults, and signals for child processes.</p> <p>-a Show the argument strings that are passed in each exec() system call.</p> <p>-e Show the environment strings that are passed in each exec() system call.</p> <p>-i Do not display interruptible sleeping system calls. Certain system calls, such as open() and read() on terminal devices or pipes can sleep for indefinite periods and are interruptible. Normally, truss reports such sleeping system calls if they remain asleep for more than one second. The system call is reported again a second time when it completes. The -i option causes such system calls to be reported only once, when they complete.</p>

- l** Include the id of the responsible lightweight process with each line of trace output. If **-f** is also specified, both the process-id and the lightweight process id are included.
- t [!]syscall,...** System calls to trace or exclude. Those system calls specified in the comma-separated list are traced. If the list begins with a **!**, the specified system calls are excluded from the trace output. Default is **-tall**.
- v [!]syscall,...** Verbose. Display the contents of any structures passed by address to the specified system calls (if traced). Input values as well as values returned by the operating system are shown. For any field used as both input and output, only the output value is shown. Default is **-v!all**.
- x [!]syscall,...** Display the arguments to the specified system calls (if traced) in raw form, usually hexadecimal, rather than symbolically. This is for unredeemed hackers who must see the raw bits to be happy. Default is **-x!all**.
- s [!]signal,...** Signals to trace or exclude. Those signals specified in the comma-separated list are traced. The trace output reports the receipt of each specified signal, even if the signal is being ignored (not blocked). (Blocked signals are not received until they are unblocked.) Signals may be specified by name or number (see **<sys/signal.h>**). If the list begins with a **!**, the specified signals are excluded from the trace output. Default is **-sall**.
- m [!]fault,...** Machine faults to trace or exclude. Those machine faults specified in the comma-separated list are traced. Faults may be specified by name or number (see **<sys/fault.h>**). If the list begins with a **!**, the specified faults are excluded from the trace output. Default is **-mall -m!ftpage**.
- r [!]fd,...** Show the full contents of the I/O buffer for each **read()** on any of the specified file descriptors. The output is formatted 32 bytes per line and shows each byte as an ascii character (preceded by one blank) or as a 2-character C language escape sequence for control characters such as horizontal tab (**\t**) and newline (**\n**). If ascii interpretation is not possible, the byte is shown in 2-character hexadecimal representation. (The first 12 bytes of the I/O buffer for each traced **read()** are shown even in the absence of **-r**.) Default is **-r!all**.
- w [!]fd,...** Show the contents of the I/O buffer for each **write()** on any of the specified file descriptors (see **-r**). Default is **-w!all**.
- o outfile** File to be used for the trace output. By default, the output goes to standard error.

See Section 2 of the *man Pages(2): System Calls* for system call names accepted by the **-t**, **-v**, and **-x** options. System call numbers are also accepted.

If **truss** is used to initiate and trace a specified command and if the **-o** option is used or if standard error is redirected to a non-terminal file, then **truss** runs with hangup, interrupt, and quit signals ignored. This facilitates tracing of interactive programs that catch interrupt and quit signals from the terminal.

If the trace output remains directed to the terminal, or if existing processes are traced (the **-p** option), then **truss** responds to hangup, interrupt, and quit signals by releasing all traced processes and exiting. This enables the user to terminate excessive trace output and to release previously-existing processes. Released processes continue normally, as though they had never been touched.

EXAMPLES

This example produces a trace of the **find(1)** command on the terminal:

```
example% truss find . -print >find.out
```

Or, to see only a trace of the open, close, read, and write system calls:

```
example% truss -t open,close,read,write find . -print >find.out
```

This produces a trace of the **spell(1)** command on the file **truss.out**:

```
example% truss -f -o truss.out spell document
```

spell is a shell script, so the **-f** flag is needed to trace not only the shell but also the processes created by the shell. (The **spell** script runs a pipeline of eight concurrent processes.)

A particularly boring example is:

```
example% truss nroff -mm document >nroff.out
```

because 97% of the output reports **lseek()**, **read()**, and **write()** system calls. To abbreviate it:

```
example% truss -t !lseek,read,write nroff -mm document >nroff.out
```

This example verbosely traces the activity of process #1, **init(1M)** (if you are a privileged user):

```
example% truss -p -v all 1
```

Interrupting **truss** returns **init** to normal operation.

FILES /**proc/nmmnn** process files
 /**proc/process-id**

SEE ALSO **intro(2)**, **proc(4)**

NOTES Some of the system calls described in Section 2 differ from the actual operating system interfaces. Do not be surprised by minor deviations of the trace output from the descriptions in Section 2.

Every machine fault (except a page fault) results in the posting of a signal to the light-weight process that incurred the fault. A report of a received signal will immediately follow each report of a machine fault (except a page fault) unless that signal is being blocked.

The operating system enforces certain security restrictions on the tracing of processes. In particular, any command whose object file (**a.out**) cannot be read by a user cannot be traced by that user; set-uid and set-gid commands can be traced only by a privileged user. Unless it is run by a privileged user, **truss** loses control of any process that performs an **exec()** of a set-id or unreadable object file; such processes continue normally, though independently of **truss**, from the point of the **exec()**.

To avoid collisions with other controlling processes, **truss** will not trace a process that it detects is being controlled by another process via the **/proc** interface. This allows **truss** to be applied to **proc(4)**-BASED debuggers as well as to another instance of itself.

The trace output contains tab characters under the assumption that standard tab stops are set (every eight positions).

The trace output for multiple processes or for a multithreaded process (one that contains more than one lightweight process) is not produced in strict time order. For example, a **read()** on a pipe may be reported before the corresponding **write()**. For any one lightweight process (a traditional process contains only one), the output is strictly time-ordered.

The system may run out of per-user process slots when tracing of children is requested. When tracing more than one process, **truss** runs as one controlling process for each process being traced. For the example of the **spell** command shown above, **spell** itself uses 9 process slots, one for the shell and 8 for the 8-member pipeline, while **truss** adds another 9 processes, for a total of 18. This is perilously close to the usual system-imposed limit of 25 processes per user.

Not all possible structures passed in all possible system calls are displayed under the **-v** option.

NAME	tset, reset – establish or restore terminal characteristics
SYNOPSIS	<pre>tset [-InQrs] [-ec] [-kc] [-m [port-ID [baudrate]: type] ...] [type] reset [-] [-ec] [-I] [-kc] [-n] [-Q] [-r] [-s] [-m [indent] [test baudrate]: type] ... [type]</pre>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>tset sets up your terminal, typically when you first log in. It does terminal dependent processing such as setting erase and kill characters, setting or resetting delays, sending any sequences needed to properly initialize the terminal, and the like. tset first determines the <i>type</i> of terminal involved, and then does necessary initializations and mode settings. If a port is not wired permanently to a specific terminal (not hardwired) it is given an appropriate generic identifier such as dialup.</p> <p>reset clears the terminal settings by turning off CBREAK and RAW modes, output delays and parity checking, turns on NEWLINE translation, echo and TAB expansion, and restores undefined special characters to their default state. It then sets the modes as usual, based on the terminal type (which will probably override some of the above). See stty(1) for more information. All arguments to tset may be used with reset. reset also uses rs= and rf= to reset the initialization string and file. This is useful after a program dies and leaves the terminal in a funny state. Often in this situation, characters will not echo as you type them. You may have to type '<LINEFEED>reset<LINEFEED>' since '<RETURN>' may not work.</p> <p>When no arguments are specified, tset reads the terminal type from the TERM environment variable and re-initializes the terminal, and performs initialization of mode, environment and other options at login time to determine the terminal type and set up terminal modes.</p> <p>When used in a startup script (.profile for sh(1) users or .login for csh(1) users) it is desirable to give information about the type of terminal you will usually use on ports that are not hardwired. Any of the alternate generic names given in /usr/share/lib/terminfo/* may be used for the identifier. Refer to the -m option below for more information. If no mapping applies and a final <i>type</i> option, not preceded by a -m, is given on the command line then that type is used.</p> <p>It is usually desirable to return the terminal type, as finally determined by tset, and information about the terminal's capabilities, to a shell's environment. This can be done using the -, -s, or -S options.</p>

For the Bourne shell, put this command in your **.profile** file:

```
eval `tset -s options...`
```

or using the C shell, put these commands in your **.login** file:

```
set noglob
eval `tset -s options...`
unset noglob
```

With the C shell, it is also convenient to make an alias in your **.cshrc** file:

```
alias ts `eval `tset -s \!*`
```

This also allows the command:

```
ts 2621
```

to be invoked at any time to set the terminal and environment. It is not possible to get this aliasing effect with a Bourne shell script, because shell scripts cannot set the environment of their parent. If a process could set its parent's environment, none of this nonsense would be necessary in the first place.

Once the terminal type is known, **tset** sets the terminal driver mode. This normally involves sending an initialization sequence to the terminal, setting the single character erase (and optionally the line-kill (full line erase)) characters, and setting special character delays. TAB and NEWLINE expansion are turned off during transmission of the terminal initialization sequence.

On terminals that can backspace but not overstrike (such as a CRT), and when the erase character is '#', the erase character is changed as if **-e** had been used.

OPTIONS

- The name of the terminal finally decided upon is output on the standard output. This is intended to be captured by the shell and placed in the **TERM** environment variable.
- ec** Set the erase character to be the named character *c* on all terminals. Default is the BACKSPACE key on the keyboard, usually **^H** (CTRL-H). The character *c* can either be typed directly, or entered using the circumflex-character notation used here.
- ic** Set the interrupt character to be the named character *c* on all terminals. Default is **^C** (CTRL-C). The character *c* can either be typed directly, or entered using the circumflex-character notation used here.
- I** Suppress transmitting terminal-initialization strings.
- kc** Set the line kill character to be the named character *c* on all terminals. Default is **^U** (CTRL-U). The kill character is left alone if **-k** is not specified. Control characters can be specified by prefixing the alphabetical character with a circumflex (as in CTRL-U) instead of entering the actual control key itself. This allows you to specify control keys that are currently assigned.
- n** Specify that the new tty driver modes should be initialized for this terminal. Probably useless since **stty new** is the default.

- Q** Suppress printing the **'Erase set to'** and **'Kill set to'** messages.
- r** In addition to other actions, reports the terminal type.
- s** Output commands to set and export **TERM**. This can be used with


```
set noglob
eval `tset -s ...`
unset noglob
```

 to bring the terminal information into the environment. Doing so makes programs such as **vi(1)** start up faster. If the **SHELL** environment variable ends with **csH**, C shell commands are output, otherwise Bourne shell commands are output.

-m [*port-ID* [*baudrate*] : *type*] ...

Specify (map) a terminal type when connected to a generic port (such as *dialup* or *plugboard*) identified by *port-ID*. The *baudrate* argument can be used to check the baudrate of the port and set the terminal type accordingly. The target rate is prefixed by any combination of the following operators to specify the conditions under which the mapping is made:

- > Greater than
- @ Equals or "at"
- < Less than
- ! It is not the case that (negates the above operators)
- ? Prompt for the terminal type. If no response is given, then *type* is selected by default.

In the following example, the terminal type is set to **adm3a** if the port is a dialup with a speed of greater than 300 or to **dw2** if the port is a dialup at 300 baud or less. In the third case, the question mark preceding the terminal type indicates that the user is to verify the type desired. A **NULL** response indicates that the named type is correct. Otherwise, the user's response is taken to be the type desired.

```
tset -m 'dialup>300:adm3a' -m 'dialup:dw2' -m \
'plugboard:?adm3a'
```

To prevent interpretation as metacharacters, the entire argument to **-m** should be enclosed in single quotes. When using the C shell, exclamation points should be preceded by a backslash (****).

EXAMPLES

These examples all use the **'-'** option. A typical use of **tset** in a **.profile** or **.login** will also use the **-e** and **-k** options, and often the **-n** or **-Q** options as well. These options have been omitted here to keep the examples short.

To select a 2621, you might put the following sequence of commands in your **.login** file (or **.profile** for Bourne shell users).

```
set noglob
eval `tset -s 2621`
unset noglob
```

If you have a switch which connects to various ports (making it impractical to identify which port you may be connected to), and use various terminals from time to time, you can select from among those terminals according to the *speed* or baud rate. In the example below, **tset** will prompt you for a terminal type if the baud rate is greater than 1200 (say, 9600 for a terminal connected by an RS-232 line), and use a Wyse® 50 by default. If the baud rate is less than or equal to 1200, it will select a 2621. Note the placement of the question mark, and the quotes to protect the > and ? from interpretation by the shell.

```
set noglob
eval `tset -s -m 'switch>1200:?wy' -m 'switch<=1200:2621'`
unset noglob
```

The following entry is appropriate if you always dial up, always at the same baud rate, on many different kinds of terminals, and the terminal you use most often is an **adm3a**.

```
set noglob
eval `tset -s ?adm3a`
unset noglob
```

If you want to make the selection based only on the baud rate, you might use the following:

```
set noglob
eval `tset -s -m '>1200:wy' 2621`
unset noglob
```

The following example quietly sets the erase character to BACKSPACE, and kill to CTRL-U. If the port is switched, it selects a Concept™ 100 for speeds less than or equal to 1200, and asks for the terminal type otherwise (the default in this case is a Wyse 50). If the port is a direct dialup, it selects Concept 100 as the terminal type. If logging in over the ARPANET, the terminal type selected is a Datamedia® 2500 terminal or emulator. Note the backslash escaping the NEWLINE at the end of the first line in the example.

```
set noglob
eval `tset -e -k^U -Q -s -m 'switch<=1200:concept100' -m \
'switch:?wy' -m dialup:concept100 -m arpanet:dm2500`
unset noglob
```

FILES

```
.login
.profile
/usr/share/lib/terminfo/*
```

SEE ALSO

```
csh(1), sh(1), stty(1), vi(1), environ(5)
```

NOTES

The **tset** command is one of the first commands a user must master when getting started on a UNIX system. Unfortunately, it is one of the most complex, largely because of the extra effort the user must go through to get the environment of the login shell set. Something needs to be done to make all this simpler, either the **login** program should do this stuff, or a default shell alias should be made, or a way to set the environment of the parent should exist.

This program cannot intuit personal choices for erase, interrupt and line kill characters, so it leaves these set to the local system standards.

It could well be argued that the shell should be responsible for ensuring that the terminal remains in a sane state; this would eliminate the need for the **reset** program.

NAME	tsort – topological sort
SYNOPSIS	<code>/usr/ccs/bin/tsort</code> [<i>filename</i>]
DESCRIPTION	<p>The tsort command produces on the standard output a totally ordered list of items consistent with a partial ordering of items mentioned in the input <i>filename</i>. If no <i>filename</i> is specified, the standard input is understood.</p> <p>The input consists of pairs of items (nonempty strings) separated by blanks. Pairs of different items indicate ordering. Pairs of identical items indicate presence, but not ordering.</p>
SEE ALSO	lorder (1)
DIAGNOSTICS	<p>Odd data: there is an odd number of fields in the input file.</p> <p>tsort returns exit code 1 if the file specified is not found.</p>

NAME	tty – get the name of the terminal
SYNOPSIS	tty [-l] [-s]
AVAILABILITY	SUNWcsu
DESCRIPTION	tty prints the path name of the user's terminal.
OPTIONS	<p>-l Prints the synchronous line number to which the user's terminal is connected, if it is on an active synchronous line.</p> <p>-s Inhibits printing of the terminal path name, allowing one to test just the exit code.</p>
EXIT CODES	<p>2 if invalid options were specified,</p> <p>0 if standard input is a terminal,</p> <p>1 otherwise.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of tty for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how tty behaves.</p> <p>LC_CTYPE Determines how tty handles characters. When LC_CTYPE is set to a valid value, tty can display and handle text and filenames containing valid characters for that locale. tty can display and handle Extended Unix Code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. tty can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p> <p>LC_MESSAGES Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).</p>
SEE ALSO	environ(5)
DIAGNOSTICS	<p>“not on an active synchronous line“ The standard input is not a synchronous terminal and -l is specified.</p> <p>“not a tty” The standard input is not a terminal and -s is not specified.</p>

NAME	typeset, type, whence – shell built-in functions to set/get attributes and values for shell variables and functions
SYNOPSIS	
sh	type [<i>name</i> ...]
ksh	†† typeset [±HLRZfilrtux[<i>n</i>]] [<i>name</i> [= <i>value</i>]] ... whence [-pv] <i>name</i> ...
DESCRIPTION	
sh	For each <i>name</i> , the type command will indicate how it would be interpreted if used as a command name.
ksh	typeset sets attributes and values for shell variables and functions. When typeset is invoked inside a function, a new instance of the variables <i>name</i> is created. The variables <i>value</i> and <i>type</i> are restored when the function completes. The following list of attributes may be specified: <ul style="list-style-type: none"> -H This flag provides UNIX to host-name file mapping on non-UNIX machines. -L Left justify and remove leading blanks from <i>value</i>. If <i>n</i> is non-zero it defines the width of the field; otherwise, it is determined by the width of the value of first assignment. When the variable is assigned to, it is filled on the right with blanks or truncated, if necessary, to fit into the field. Leading zeros are removed if the -Z flag is also set. The -R flag is turned off. -R Right justify and fill with leading blanks. If <i>n</i> is non-zero it defines the width of the field, otherwise it is determined by the width of the value of first assignment. The field is left filled with blanks or truncated from the end if the variable is re-assigned. The -L flag is turned off. -Z Right justify and fill with leading zeros if the first non-blank character is a digit and the -L flag has not been set. If <i>n</i> is non-zero it defines the width of the field; otherwise, it is determined by the width of the value of first assignment. -f The names refer to function names rather than variable names. No assignments can be made and the only other valid flags are -t, -u and -x. The flag -t turns on execution tracing for this function. The flag -u causes this function to be marked undefined. The FPATH variable will be searched to find the function definition when the function is referenced. The flag -x allows the function definition to remain in effect across shell procedures invoked by name. -i Parameter is an integer. This makes arithmetic faster. If <i>n</i> is non-zero it defines the output arithmetic base; otherwise, the first assignment determines the output base. -l All upper-case characters are converted to lower-case. The upper-case flag, -u is turned off. -r The given <i>names</i> are marked readonly and these names cannot be changed by subsequent assignment.

- t Tags the variables. Tags are user definable and have no special meaning to the shell.
- u All lower-case characters are converted to upper-case characters. The lower-case flag, -l is turned off.
- x The given *names* are marked for automatic export to the **environment** of subsequently-executed commands.

The -i attribute can not be specified along with -R, -L, -Z, or -f.

Using + rather than - causes these flags to be turned off. If no *name* arguments are given but flags are specified, a list of *names* (and optionally the *values*) of the *variables* which have these flags set is printed. (Using + rather than - keeps the values from being printed.) If no *names* and flags are given, the *names* and *attributes* of all *variables* are printed.

For each *name*, **whence** indicates how it would be interpreted if used as a command name.

The -v flag produces a more verbose report.

The -p flag does a path search for *name* even if name is an alias, a function, or a reserved word.

On this man page, **ksh(1)** commands that are preceded by one or two † (daggers) are treated specially in the following ways:

1. Variable assignment lists preceding the command remain in effect when the command completes.
2. I/O redirections are processed after variable assignments.
3. Errors cause a script that contains them to abort.
4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.

SEE ALSO **ksh(1), set(1), sh(1)**

NAME	ucblinks – adds /dev entries to give SunOS 4.1 compatible names to SunOS 5.x devices
SYNOPSIS	/usr/ucb/ucblinks [-e rulebase] [-r rootdir]
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>ucblinks creates symbolic links under the /dev directory for devices whose SunOS 5.x names differ from their SunOS 4.1 names. Where possible, these symbolic links point to the device's SunOS 5.x name rather than to the actual /devices entry.</p> <p>ucblinks does not remove unneeded compatibility links; these must be removed by hand.</p> <p>ucblinks should be called each time the system is reconfiguration-booted, after any new SunOS 5.x links that are needed have been created, since the reconfiguration may have resulted in more compatibility names being needed.</p> <p>In releases prior to SunOS 5.4, ucblinks used a nawk rule-base to construct the SunOS 4.1 compatible names. ucblinks no longer uses nawk for the default operation, although nawk rule-bases can still be specified with the -e option. The nawk rule-base equivalent to the SunOS 5.4 default operation can be found in /usr/ucblib/ucblinks.awk.</p>
OPTIONS	<p>-e rulebase Specify <i>rulebase</i> as the file containing nawk(1) pattern-action statements.</p> <p>-r rootdir Specify <i>rootdir</i> as the directory under which dev and devices will be found, rather than the standard root directory /.</p>
FILES	/usr/ucblib/ucblinks.awk sample rule-base for compatibility links
SEE ALSO	devlinks (1M), disks (1M), ports (1M), tapes (1M)

NAME	ul – do underlining
SYNOPSIS	ul [-i] [-t terminal] [<i>filename...</i>]
AVAILABILITY	SUNWdoc
DESCRIPTION	ul reads the named <i>filenames</i> (or the standard input if none are given) and translates occurrences of underscores to the sequence which indicates underlining for the terminal in use, as specified by the environment variable TERM . ul uses the /usr/share/lib/terminfo entry to determine the appropriate sequences for underlining. If the terminal is incapable of underlining, but is capable of a standout mode then that is used instead. If the terminal can overstrike, or handles underlining automatically, ul degenerates to cat(1) . If the terminal cannot underline, underlining is ignored.
OPTIONS	-t terminal Override the terminal kind specified in the environment. If the terminal cannot underline, underlining is ignored. If the terminal name is not found, no underlining is attempted. -i Indicate underlining by a separate line containing appropriate dashes ‘-’; this is useful when you want to look at the underlining which is present in an nroff(1) output stream on a CRT-terminal.
RETURN VALUES	ul returns exit code 1 if the file specified is not found.
FILES	/usr/share/lib/terminfo/*
SEE ALSO	cat(1) , man(1) , nroff(1)
BUGS	nroff usually generates a series of backspaces and underlines intermixed with the text to indicate underlining. ul makes attempt to optimize the backward motion.

NAME	umask – shell built-in function to restrict read/write/execute permissions
SYNOPSIS	
sh	umask [<i>ooo</i>]
cs h	umask [<i>ooo</i>]
ksh	umask [<i>mask</i>]
DESCRIPTION	
sh	<p>The user file-creation mode mask is set to <i>ooo</i>. The three octal digits refer to read/write/execute permissions for owner, group, and other, respectively (see chmod(1), chmod(2), and umask(2)). The value of each specified digit is subtracted from the corresponding “digit” specified by the system for the creation of a file (see creat(2)). For example, umask 022 removes write permission for group and other (files normally created with mode 777 become mode 755; files created with mode 666 become mode 644).</p> <p>If <i>ooo</i> is omitted, the current value of the mask is printed.</p> <p>umask is recognized and executed by the shell.</p> <p>umask can be included in the user’s .profile (see profile(4)) and invoked at login to automatically set the user’s permissions on files or directories created.</p>
cs h	See the description above for the Bourne shell (sh) umask built-in.
ksh	The user file-creation mask is set to <i>mask</i> . <i>mask</i> can either be an octal number or a symbolic value as described in chmod (1). If a symbolic value is given, the new umask value is the complement of the result of applying <i>mask</i> to the complement of the previous umask value. If <i>mask</i> is omitted, the current value of the mask is printed.
SEE ALSO	chmod (1), cs h (1), ksh (1), sh (1), chmod (2), creat (2), profile (4)

NAME	uname – print name of current system
SYNOPSIS	uname [-amnprsv] uname [-S <i>system_name</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	uname prints information about the current system on the standard output. If no options are specified, uname prints the current operating system's name. The options print selected information returned by uname (2) and/or sysinfo (2).
OPTIONS	<p>-a Print all information.</p> <p>-m Print the machine hardware name.</p> <p>-n Print the nodename (the nodename is the name by which the system is known to a communications network).</p> <p>-p Print the current host's processor type.</p> <p>-r Print the operating system release.</p> <p>-s Print the name of the operating system. This is the default.</p> <p>-v Print the operating system version.</p> <p>-S <i>systemname</i> The nodename may be changed by specifying a system name argument. The system name argument is restricted to SYS_NMLN characters. SYS_NMLN is an implementation specific value defined in <sys/utsname.h>. Only the super-user is allowed this capability.</p>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of uname for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how uname behaves.</p> <p>LC_CTYPE Determines how uname handles characters. When LC_CTYPE is set to a valid value, uname can display and handle text and filenames containing valid characters for that locale. uname can display and handle Extended Unix code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. uname can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.</p>

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO [sysinfo\(2\)](#), [uname\(2\)](#), [environ\(5\)](#)

NAME	unifdef – resolve and remove <code>ifdef</code> ed lines from C program source
SYNOPSIS	unifdef [<code>-c</code>] [<code>-Dname</code>] [<code>-Uname</code>] [<code>-iDname</code>] [<code>-iUname</code>] ... [<i>filename</i>]
DESCRIPTION	<p>unifdef removes ifdefed lines from a file while otherwise leaving the file alone. It is smart enough to deal with the nested ifdefs, comments, single and double quotes of C syntax, but it does not do any including or interpretation of macros. Neither does it strip out comments, though it recognizes and ignores them. You specify which symbols you want defined with <code>-D</code> options, and which you want undefined with <code>-U</code> options. Lines within those ifdefs will be copied to the output, or removed, as appropriate. Any ifdef, ifndef, else, and endif lines associated with <i>filename</i> will also be removed.</p> <p>ifdefs involving symbols you do not specify are untouched and copied out along with their associated ifdef, else, and endif lines.</p> <p>If an ifdefX occurs nested inside another ifdefX, then the inside ifdef is treated as if it were an unrecognized symbol. If the same symbol appears in more than one argument, only the first occurrence is significant.</p> <p>unifdef copies its output to the standard output and will take its input from the standard input if no <i>filename</i> argument is given.</p>
OPTIONS	<p><code>-c</code> Complement the normal operation. Lines that would have been removed or blanked are retained, and vice versa.</p> <p><code>-l</code> Replace “lines removed” lines with blank lines.</p> <p><code>-t</code> Plain text option. unifdef refrains from attempting to recognize comments and single and double quotes.</p> <p><code>-Dname</code> Lines associated with the defined symbol <i>name</i>.</p> <p><code>-Uname</code> Lines associated with the undefined symbol <i>name</i>.</p> <p><code>-iDname</code> Ignore, but print out, lines associated with the defined symbol <i>name</i>. If you use ifdefs to delimit non-C lines, such as comments or code which is under construction, then you must tell unifdef which symbols are used for that purpose so that it will not try to parse for quotes and comments within them.</p> <p><code>-iUname</code> Ignore, but print out, lines associated with the undefined symbol <i>name</i>.</p>
SEE ALSO	diff(1)
DIAGNOSTICS	<p>Premature EOF Inappropriate else or endif.</p> <p>Exit status is 1 if unifdef encounters problems, and 0 otherwise.</p>

NAME	uniq – report repeated lines in a file
SYNOPSIS	uniq [[-u] [-d] [-c] [+n] [-n]] [<i>input</i> [<i>output</i>]]
AVAILABILITY	SUNWesu
DESCRIPTION	uniq reads the input file comparing adjacent lines. In the normal case, the second and succeeding copies of repeated lines are removed; the remainder is written on the output file. <i>Input</i> and <i>output</i> should always be different. Note that repeated lines must be adjacent in order to be found; see sort (1).
OPTIONS	<p>If the -d and the -u options are used, the last option used on the command line will supersede the other option.</p> <p>-u If the -u flag is used, just the lines that are not repeated in the original file are output.</p> <p>-d The -d option specifies that one copy of just the repeated lines is to be written. The normal mode output is the union of the -u and -d mode outputs.</p> <p>-c The -c option supersedes -u and -d and generates an output report in default style but with each line preceded by a count of the number of times it occurred.</p> <p>The <i>n</i> arguments specify skipping an initial portion of each line in the comparison:</p> <p>-n The first <i>n</i> fields together with any blanks before each are ignored. A field is defined as a string of non-space, non-tab characters separated by tabs and spaces from its neighbors.</p> <p>+n The first <i>n</i> characters are ignored. Fields are skipped before characters.</p>
EXAMPLES	<p>The following example lists the contents of the uniq.test file and outputs a copy of the repeated lines.</p> <pre>example% cat uniq.test This is a test. This is a test. TEST. Computer. TEST. TEST. Software.</pre> <pre>example% uniq -d uniq.test This is a test. TEST. example%</pre>

The next example outputs just those lines that are not repeated in the **uniq.test** file.

```
example% uniq -u uniq.test
TEST.
Computer.
Software.
example%
```

The last example outputs a report with each line preceded by a count of the number of times each line occurred in the file.

```
example% uniq -c uniq.test
2 This is a test.
1 TEST.
1 Computer.
2 TEST.
1 Software.
example%
```

ENVIRONMENT

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

LC_CTYPE

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

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

SEE ALSO

comm(1), **sort(1)**, **environ(5)**

NAME	units – converts quantities expressed in standard scales to other scales
SYNOPSIS	units
AVAILABILITY	SUNWesu
DESCRIPTION	<p>units converts quantities expressed in various standard scales to their equivalents in other scales. It works interactively in this fashion:</p> <p style="padding-left: 40px;">You have: inch You want: cm * 2.540000e+00 / 3.937008e-01</p> <p>A quantity is specified as a multiplicative combination of units optionally preceded by a numeric multiplier. Powers are indicated by suffixed positive integers, division by the usual sign:</p> <p style="padding-left: 40px;">You have: 15 lbs force/in2 You want: atm * 1.020689e+00 / 9.797299e-01</p> <p>units only does multiplicative scale changes; thus it can convert Kelvin to Rankine, but not Celsius to Fahrenheit. Most familiar units, abbreviations, and metric prefixes are recognized, together with a generous leavening of exotica and a few constants of nature including:</p> <p style="padding-left: 40px;">pi ratio of circumference to diameter, c speed of light, e charge on an electron, g acceleration of gravity, force same as g, mole Avogadro's number, water pressure head per unit height of water, au astronomical unit.</p> <p>Pound is not recognized as a unit of mass; lb is. Compound names are run together, (for example, lightyear). British units that differ from their U.S. counterparts are prefixed thus: brgallon. For a complete list of units, type:</p> <p style="padding-left: 40px;">cat /usr/share/lib/unittab</p>
FILES	/usr/share/lib/unittab

NAME	unix2dos – convert text file from ISO format to DOS format
SYNOPSIS	unix2dos [-ascii] [-iso] [-7] <i>originalfile convertedfile</i>
AVAILABILITY	SUNWesu
DESCRIPTION	<p>unix2dos converts ISO standard characters to the corresponding characters in the DOS extended character set.</p> <p>This command may be invoked from either DOS or SunOS. However, the filenames must conform to the conventions of the environment in which the command is invoked.</p> <p>If the original file and the converted file are the same, unix2dos will rewrite the original file after converting it.</p>
OPTIONS	<p>-ascii Adds carriage returns and converts end of file characters in SunOS format text files to conform to DOS requirements.</p> <p>-iso This is the default. Converts ISO standard characters to the corresponding character in the DOS extended character set.</p> <p>-7 Convert 8 bit SunOS characters to 7 bit DOS characters.</p>
DIAGNOSTICS	<p>File <i>filename</i> not found, or no read permission The input file you specified does not exist, or you do not have read permission (check with the SunOS command ls -l).</p> <p>Bad output filename <i>filename</i>, or no write permission The output file you specified is either invalid, or you do not have write permission for that file or the directory that contains it. Check also that the drive or diskette is not write-protected.</p> <p>Error while writing to temporary file An error occurred while converting your file, possibly because there is not enough space on the current drive. Check the amount of space on the current drive using the DIR command. Also be certain that the default diskette or drive is write-enabled (not write-protected). Note that when this error occurs, the original file remains intact.</p> <p>Could not rename tmpfile to <i>filename</i>. Translated tmpfile name = <i>filename</i>. The program could not perform the final step in converting your file. Your converted file is stored under the name indicated on the second line of this message.</p>
SEE ALSO	dos2unix(1)

NAME	uptime – show how long the system has been up
SYNOPSIS	uptime
AVAILABILITY	SUNWcsu
DESCRIPTION	The uptime command prints the current time, the length of time the system has been up, and the average number of jobs in the run queue over the last 1, 5 and 15 minutes. It is, essentially, the first line of a w(1) command.
EXAMPLE	Below is an example of the output uptime provides: example% uptime 10:47am up 27 day(s), 50 mins, 1 user, load average: 0.18, 0.26, 0.20
SEE ALSO	w(1) , who(1) , whodo(1M)
NOTES	who -b gives the time the system was last booted.

NAME	users – display a compact list of users logged in
SYNOPSIS	/usr/ucb/users [<i>filename</i>]
AVAILABILITY	SUNWscpu
DESCRIPTION	users lists the login names of the users currently on the system in a compact, one-line format. Specifying <i>filename</i> , tells users where to find its information; by default it checks /var/adm/utmp . Typing <i>users</i> is equivalent to typing who -q .
EXAMPLES	example% users paul george ringo example%
FILES	/var/adm/utmp
SEE ALSO	who(1)

NAME	uucp , uulog , uuname – UNIX-to-UNIX system copy
SYNOPSIS	<p>uucp [-c -C] [-d -f] [-ggrade] [-j] [-m] [-nuser] [-r] [-sfile] [-xdebug_level] <i>source-file destination-file</i></p> <p>uulog [-ssys] [-f system] [-x] [-number] <i>system</i></p> <p>uuname [-c] [-l]</p>
AVAILABILITY	SUNWbnuu
DESCRIPTION	<p>uucp copies files named by the <i>source-file</i> arguments to the <i>destination-file</i> argument. A source file name may be a path name on your machine, or may have the form:</p> <p style="padding-left: 40px;"><i>system-name!pathname</i></p> <p>where <i>system-name</i> is taken from a list of system names that uucp knows about. <i>source_file</i> is restricted to no more than one <i>system-name</i>. The destination <i>system-name</i> may also include a list of system names such as</p> <p style="padding-left: 40px;"><i>system-name!system-name!...!system-name!pathname</i></p> <p>In this case, an attempt is made to send the file, using the specified route, to the destination. Care should be taken to ensure that intermediate nodes in the route are willing to forward information (see NOTES below for restrictions).</p> <p>For C-Shell users, the “!” character must be surrounded by single quotes (’), or preceded by a backslash (\).</p> <p>The shell metacharacters ?, * and [...] appearing in <i>path name</i> will be expanded on the appropriate system.</p> <p>Path names may be one of:</p> <ol style="list-style-type: none"> (1) a full path name; (2) a path name preceded by <i>~user</i> where <i>user</i> is a login name on the specified system and is replaced by that user’s login directory; (3) a path name preceded by <i>~/destination</i> where <i>destination</i> is appended to /var/spool/uucppublic; (Note: This destination will be treated as a file name unless more than one file is being transferred by this request or the destination is already a directory. To ensure that it is a directory, follow the destination with a ‘/’. For example <i>~/dan/</i> as the destination will make the directory /var/spool/uucppublic/dan if it does not exist and put the requested file(s) in that directory). (4) anything else is prefixed by the current directory. <p>If the result is an erroneous path name for the remote system, the copy will fail. If the <i>destination-file</i> is a directory, the last part of the <i>source-file</i> name is used.</p> <p>Invoking uucp with shell wildcard characters as the remote <i>source-file</i> invokes the uux(1C) command to execute the uucp command on the remote machine. The remote uucp command spools the files on the remote machine. After the first session terminates,</p>

if the remote machine is configured to transfer the spooled files to the local machine, the remote machine will initiate a call and send the files; otherwise, the user must "call" the remote machine to transfer the files from the spool directory to the local machine. This call can be done manually using **Uutry(1M)**, or as a side effect of another **uux(1C)** or **uucp** call.

Note that the local machine has to have permission to execute the **uucp** command on the remote machine in order for the remote machine to send the spooled files.

uucp removes execute permissions across the transmission and gives **0666** read and write permissions (see **chmod(2)**).

uulog queries a log file of **uucp** or **uuxqt** transactions in file **/var/uucp/.Log/uucico/system** or **/var/uucp/.Log/uuxqt/system**.

uname lists the names of systems known to **uucp**.

OPTIONS

The following options are interpreted by **uucp**:

- c** Do not copy local file to the spool directory for transfer to the remote machine (default).
- C** Force the copy of local files to the spool directory for transfer.
- d** Make all necessary directories for the file copy (default).
- f** Do not make intermediate directories for the file copy.
- ggrade** *grade* can be either a single letter, number, or a string of alphanumeric characters defining a service grade. The **uuglist** command can determine whether it is appropriate to use the single letter, number, or a string of alphanumeric characters as a service grade. The output from the **uuglist** command will be a list of service grades that are available, or a message that says to use a single letter or number as a grade of service.
- j** Print the **uucp** job identification string on standard output. This job identification can be used by **uustat** to obtain the status of a **uucp** job or to terminate a **uucp** job. The **uucp** job is valid as long as the job remains queued on the local system.
- m** Send mail to the requester when the copy is complete.
- nuser** Notify *user* on the remote system that a file was sent.
- r** Do not start the file transfer, just queue the job.
- sfile** Report status of the transfer to *file*. This option is accepted for compatibility, but it is ignored because it is insecure.
- xdebug_level** Produce debugging output on standard output. *debug_level* is a number between 0 and 9; as it increases to 9, more detailed debugging information is given. This option may not be available on all systems.

The following options cause **uulog** to print logging information:

- ssys** Print information about file transfer work involving system *sys*.
- fsystem** Do a "**tail -f**" of the file transfer log for **system**. (You must hit BREAK to

exit this function.)

Other options used in conjunction with the above options are:

- x Look in the **uuxqt** log file for the given system.
- number* Execute a **tail** command of *number* lines.

uname options are:

- c Display the names of systems known to **cu**. The two lists are the same, unless your machine is using different **Systems** files for **cu** and **uucp**. See the **Sysfiles** file.
- l Display the local system name.

FILES	/etc/uucp/*	other data files
	/var/spool/uucp	spool directories
	/usr/lib/uucp/*	other program files
	/var/spool/uucppublic/*	public directory for receiving and sending

SEE ALSO **mail(1)**, **uuglist(1C)**, **uustat(1C)**, **uux(1C)**, **Uutry(1M)**, **uuxqt(1M)**, **chmod(2)**

NOTES

For security reasons, the domain of remotely accessible files may be severely restricted. You will probably not be able to access files by path name; ask a responsible person on the remote system to send them to you. For the same reasons you will probably not be able to send files to arbitrary path names. As distributed, the remotely accessible files are those whose names begin **/var/spool/uucppublic** (equivalent to **~/**).

All files received by **uucp** will be owned by **uucp**.

The **-m** option will only work when sending files or receiving a single file. Receiving multiple files specified by special shell characters **?**, **&**, and **[...]** will not activate the **-m** option.

The forwarding of files through other systems may not be compatible with the previous version of **uucp**. If forwarding is used, all systems in the route must have compatible versions of **uucp**.

Protected files and files that are in protected directories that are owned by the requester can be sent by **uucp**. However, if the requester is root, and the directory is not searchable by "other" or the file is not readable by "other", the request will fail.

NAME	uuencode, uudecode – encode a binary file, or decode its ASCII representation
SYNOPSIS	uuencode [<i>source-file</i>] <i>file-label</i> uudecode [-p] [<i>encoded-file</i>]
AVAILABILITY	SUNWesu
DESCRIPTION	<p>uuencode converts a binary file into an ASCII-encoded representation that can be sent using mail(1). It encodes the contents of <i>source-file</i>, or the standard input if no <i>source-file</i> argument is given. The <i>file-label</i> argument is required. The <i>file-label</i> is included in the encoded file's header as the name of the file into which uudecode is to place the binary (decoded) data. uuencode also includes the permission modes of <i>source-file</i>, (except setuid, setgid, and sticky-bits), so that <i>file-label</i> is recreated with those same permission modes.</p> <p>uudecode reads an <i>encoded-file</i>, strips off any leading and trailing lines added by mailer programs, and recreates the original binary data with the filename and the mode specified in the header.</p> <p>The encoded file is an ordinary ASCII text file; it can be edited by any text editor. But it is best only to change the mode or file-label in the header to avoid corrupting the decoded binary.</p>
OPTIONS	-p decode <i>encoded-file</i> and send it to standard output. This allows uudecode to be used in a pipeline.
SEE ALSO	mail (1), uucp (1C), uux (1C)
NOTES	<p>The encoded file's size is expanded by 35% (3 bytes become 4, plus control information), causing it to take longer to transmit than the equivalent binary.</p> <p>The user on the remote system who is invoking uudecode (typically uucp) must have write permission on the file specified in the <i>file-label</i>.</p> <p>If you uuencode then uudecode a file in the same directory, you will overwrite the original file.</p>

NAME	uuglist – print the list of service grades that are available on this UNIX system
SYNOPSIS	uuglist [-u]
AVAILABILITY	SUNWbnuu
DESCRIPTION	uuglist prints the list of service grades that are available on the system to use with the -g option of uucp (1C) and uux (1C).
OPTIONS	-u List the names of the service grades that the user is allowed to use with the -g option of the uucp and uux commands.
FILES	/etc/uucp/Grades contains the list of service grades
SEE ALSO	uucp (1C), uux (1C)

NAME	uustat – uucp status inquiry and job control
SYNOPSIS	uustat [-m] [-p] [-q] [-kjobid [-n]] [-rjobid [-n]] uustat [-a] [-ssystem [-j]] [-uuser] [-Sgric] uustat -tssystem [-c] [-dnumber]
AVAILABILITY	SUNWbnuu
DESCRIPTION	uustat functions in the following three areas: displays the general status of, or cancels, previously specified uucp commands; provides remote system performance information, in terms of average transfer rates or average queue times; and provides general remote system-specific and user-specific status of uucp connections to other systems.
OPTIONS	Here are the options that obtain general status of, or cancel, previously specified uucp commands: <ul style="list-style-type: none"> -a List all jobs in queue. -j List the total number of jobs displayed. The -j option can be used in conjunction with the -a or the -s option. -kjobid Kill the uucp request whose job identification is <i>jobid</i>. The killed uucp request must belong to the user issuing the uustat command unless the user is the super-user or uucp administrator. If the job is killed by the super-user or uucp administrator, electronic mail is sent to the user. -m Report the status of accessibility of all machines. -n Suppress all standard output, but not standard error. The -n option is used in conjunction with the -k and -r options. -p Execute the command ps -flp for all the process-ids that are in the lock files. -q List the jobs queued for each machine. If a status file exists for the machine, its date, time and status information are reported. In addition, if a number appears in parentheses next to the number of C or X files, it is the age in days of the oldest C./X. file for that system. The Retry field represents the number of hours until the next possible call. The Count is the number of failure attempts. Note: For systems with a moderate number of outstanding jobs, this could take 30 seconds or more of real-time to execute. An example of the output produced by the -q option is: <pre style="margin-left: 40px;"> eagle 3C 04/07-11:07 NO DEVICES AVAILABLE mh3bs3 2C 07/07-10:42 SUCCESSFUL </pre> <p>This indicates the number of command files that are waiting for each system. Each command file may have zero or more files to be sent (zero means to call the system and see if work is to be done). The date and time refer to the previous interaction with the system followed by the status of the interaction.</p>

-rjobid Rejuvenate *jobid*. The files associated with *jobid* are touched so that their modification time is set to the current time. This prevents the cleanup daemon from deleting the job until the jobs' modification time reaches the limit imposed by the daemon.

Here are the options that provide remote system performance information, in terms of average transfer rates or average queue times; the **-c** and **-d** options can only be used in conjunction with the **-t** option:

-tsystem Report the average transfer rate or average queue time for the past 60 minutes for the remote *system*. The following parameters can only be used with this option:

-c Average queue time is calculated when the **-c** parameter is specified and average transfer rate when **-c** is not specified. For example, the command:

```
example% uustat -teagle -d50 -c
```

produces output in the following format:

```
average queue time to eagle for last 50 minutes: 5 seconds
```

The same command without the **-c** parameter produces output in the following format:

```
average transfer rate with eagle for last 50 minutes: 2000.88 bytes/sec
```

-dnumber *number* is specified in minutes. Used to override the 60 minute default used for calculations. These calculations are based on information contained in the optional performance log and therefore may not be available. Calculations can only be made from the time that the performance log was last cleaned up.

Here are the options that provide general remote system-specific and user-specific status of **uucp** connections to other systems. Either or both of the following options can be specified with **uustat**. The **-j** option can be used in conjunction with the **-s** option to list the total number of jobs displayed:

-ssystem Report the status of all **uucp** requests for remote system *system*.

-uuser Report the status of all **uucp** requests issued by *user*.

Output for both the **-s** and **-u** options has the following format:

```
eagleN1bd7 4/07-11:07 S eagle dan 522 /home/dan/A  
eagleC1bd8 4/07-11:07 S eagle dan 59 D.3b2al2ce4924  
4/07-11:07 S eagle dan rmail mike
```

With the above two options, the first field is the *jobid* of the job. This is followed by the date/time. The next field is an **S** if the job is sending a file or an **R** if the job is requesting a file. The next field is the machine where the file is to be transferred. This is followed by the user-id of the user who queued the job. The next field contains the size of the file, or in the case of a remote execution (**rmail** is the command used for remote mail), the name of the command. When the size appears in this field, the file name is also given. This can either be the name given by the user or an internal name (for example, **D.3b2alce4924**) that is created for data files associated with remote executions (**rmail** in this example).

-Sqrlic Report the job state: **q** for queued jobs, **r** for running jobs, **i** for interrupted jobs, and **c** for completed jobs.

A job is queued if the transfer has not started. A job is running when the transfer has begun. A job is interrupted if the transfer began but was terminated before the file was completely transferred. A completed job is a job that successfully transferred. The completed state information is maintained in the accounting log, which is optional and therefore may be unavailable. The parameters can be used in any combination, but at least one parameter must be specified. The **-S** option can also be used with **-s** and **-u** options. The output for this option is exactly like the output for **-s** and **-u** except that the job states are appended as the last output word. Output for a completed job has the following format:

eagleC1bd3 completed

When no options are given, **uustat** outputs the status of all **uucp** requests issued by the current user.

FILES

/var/spool/uucp/*	spool directories
/var/uucp/.Admin/account	accounting log
/var/uucp/.Admin/perflog	performance log

SEE ALSO **uucp(1C)**

DIAGNOSTICS The **-t** option produces no message when the data needed for the calculations is not being recorded.

NOTES After the user has issued the **uucp** request, if the file to be transferred is moved, deleted or was not copied to the spool directory (**-C** option) when the **uucp** request was made, **uustat** reports a file size of **-99999**. This job will eventually fail because the file(s) to be transferred can not be found.

NAME	uuto , uupick – public UNIX-to-UNIX system file copy
SYNOPSIS	uuto [<i>options</i>] <i>source-files destination</i> uupick [<i>-s system</i>]
AVAILABILITY	SUNWbnuu
DESCRIPTION	<p>uuto sends <i>source-files</i> to <i>destination</i>. uuto uses the uucp(1C) facility to send files, while it allows the local system to control the file access. A source-file name is a path name on your machine. Destination has the form:</p> <p style="text-align: center;">system[<i>!system</i>] ... <i>!user</i></p> <p>where system is taken from a list of system names that uucp knows about. <i>User</i> is the login name of someone on the specified system.</p> <p>The files (or sub-trees if directories are specified) are sent to <i>PUBDIR</i> on system, where <i>PUBDIR</i> is a public directory defined in the uucp source. By default, this directory is <i>/var/spool/uucppublic</i>. Specifically the files are sent to</p> <p style="text-align: center;"><i>PUBDIR/receive/user/mysystem/files</i>.</p> <p>The recipient is notified by mail(1) of the arrival of files.</p>
OPTIONS	<p>The following options apply to uuto:</p> <p>-p Copy the source file into the spool directory before transmission.</p> <p>-m Send mail to the sender when the copy is complete.</p> <p>The following option applies to uupick:</p> <p>-s system Search only the <i>PUBDIR</i> for files sent from system.</p>
USAGE	<p>uupick accepts or rejects the files transmitted to the user. Specifically, uupick searches <i>PUBDIR</i> for files destined for the user. For each entry (file or directory) found, the following message is printed on standard output:</p> <p style="text-align: center;">from system sysname: [file <i>file-name</i>] [dir <i>dirname</i>] ?</p> <p>uupick then reads a line from standard input to determine the disposition of the file:</p> <p><new-line> Go to next entry.</p> <p>d Delete the entry.</p> <p>m [<i>dir</i>] Move the entry to named directory <i>dir</i>. If <i>dir</i> is not specified as a complete path name (in which <i>\$HOME</i> is legitimate), a destination relative to the current directory is assumed. If no destination is given, the default is the current directory.</p> <p>a [<i>dir</i>] Same as m above, except it moves all the files sent from system.</p>

p Print the content of the file.
q Stop.
 EOT (control-d) Same as **q**.
!command Escape to the shell to do *command*.
***** Print a command summary.

FILES *PUBDIR* `/var/spool/uucppublic` public directory

SEE ALSO `mail(1)`, `uucp(1C)`, `uustat(1C)`, `uux(1C)`, `uucleanup(1M)`

NOTES In order to send files that begin with a dot (for instance, `.profile`), the files must be qualified with a dot. For example, the following files are correct:

`.profile` `.prof*` `.profil?`

The following files are incorrect:

`*prof*` `?profile`

NAME	uux – UNIX-to-UNIX system command execution
SYNOPSIS	uux [-] [-aname] [-b] [-c] [-C] [-ggrade] [-j] [-n] [-p] [-r] [-sfilename] [-xdebug_level] [-z] <i>command-string</i>
AVAILABILITY	SUNWbnuu
DESCRIPTION	<p>uux will gather zero or more files from various systems, execute a command on a specified system and then send standard output to a file on a specified system.</p> <p>Note: For security reasons, most installations limit the list of commands executable on behalf of an incoming request from uux, permitting only the receipt of mail (see mail(1)). (Remote execution permissions are defined in /etc/uucp/Permissions.)</p> <p>The <i>command-string</i> is made up of one or more arguments that look like a shell command line, except that the command and file names may be prefixed by <i>system-name!</i>. A null <i>system-name</i> is interpreted as the local system.</p> <p>File names may be one of:</p> <ul style="list-style-type: none"> • a full pathname; • a pathname preceded by <i>~xxx</i>, where <i>xxx</i> is a login name on the specified system and is replaced by that user's login directory; • anything else is prefixed by the current directory. <p>As an example, the command:</p> <pre>example% uux "!diff sys1!/home/dan/filename1 sys2!/a4/dan/filename2 > !~/dan/filename.diff"</pre> <p>will get the <i>filename1</i> and <i>filename2</i> files from the "sys1" and "sys2" machines, execute a diff(1) command and put the results in <i>filename.diff</i> in the local <i>PUBDIR/dan/</i> directory. <i>PUBDIR</i> is a public directory defined in the uucp source. By default, this directory is /var/spool/uucppublic.</p> <p>Any special shell characters such as <i><</i>, <i>></i>, <i>;</i>, <i> </i> should be quoted either by quoting the entire <i>command-string</i>, or quoting the special characters as individual arguments.</p> <p>uux will attempt to get all appropriate files to the specified system where they will be processed. For files that are output files, the file name must be escaped using parentheses. For example, the command:</p> <pre>example% uux "a!cut -f1 b!/usr/filename > c!/usr/filename"</pre> <p>gets /usr/filename from system "b" and sends it to system "a", performs a cut command on that file and sends the result of the cut command to system "c".</p> <p>uux will notify you if the requested command on the remote system was disallowed. This notification can be turned off by the -n option. The response comes by remote mail from the remote machine.</p>

OPTIONS	<ul style="list-style-type: none"> - The standard input to uux is made the standard input to the <i>command-string</i>. -a<i>name</i> Use <i>name</i> as the user job identification replacing the initiator user-id. (Notification will be returned to user-id <i>name</i>.) -b Return whatever standard input was provided to the uux command if the exit status is non-zero. -c Do not copy local file to the spool directory for transfer to the remote machine (default). -C Force the copy of local files to the spool directory for transfer. -g<i>grade</i> <i>grade</i> can be either a single letter, number, or a string of alphanumeric characters defining a service grade. The uuglist(1C) command determines whether it is appropriate to use the single letter, number, or a string of alphanumeric characters as a service grade. The output from the <i>uuglist</i> command will be a list of service grades that are available or a message that says to use a single letter or number as a grade of service. -j Output the jobid string on the standard output which is the job identification. This job identification can be used by uustat(1C) to obtain the status or terminate a job. -n Do not notify the user if the command fails. -p Same as -: The standard input to uux is made the standard input to the <i>command-string</i>. -r Do not start the file transfer, just queue the job. -s <i>filename</i> Report status of the transfer in <i>filename</i>. This option is accepted for compatibility, but it is ignored because it is insecure. -x<i>debug_level</i> Produce debugging output on the standard output. <i>debug_level</i> is a number between 0 and 9; as it increases to 9, more detailed debugging information is given. -z Send success notification to the user. 								
FILES	<table border="0"> <tr> <td style="padding-right: 20px;">/etc/uucp/*</td> <td>other data and programs</td> </tr> <tr> <td>/etc/uucp/Permissions</td> <td>remote execution permissions</td> </tr> <tr> <td>/usr/lib/uucp/*</td> <td>other programs</td> </tr> <tr> <td>/var/spool/uucp</td> <td>spool directories</td> </tr> </table>	/etc/uucp/*	other data and programs	/etc/uucp/Permissions	remote execution permissions	/usr/lib/uucp/*	other programs	/var/spool/uucp	spool directories
/etc/uucp/*	other data and programs								
/etc/uucp/Permissions	remote execution permissions								
/usr/lib/uucp/*	other programs								
/var/spool/uucp	spool directories								
SEE ALSO	cut(1), mail(1), uucp(1C), uuglist(1C), uustat(1C)								
NOTES	<p>Only the first command of a shell pipeline may have a <i>system-name</i>!. All other commands are executed on the system of the first command.</p> <p>The use of the shell metacharacter * will probably not do what you want it to do. The shell tokens << and >> are not implemented.</p>								

The execution of commands on remote systems takes place in an execution directory known to the **uucp** system. All files required for the execution will be put into this directory unless they already reside on that machine. Therefore, the simple file name (without path or machine reference) must be unique within the **uux** request. The following command will NOT work:

```
example% uux "a!diff b!/home/dan/xyz c!/home/dan/xyz > !xyz.diff"
```

but the command:

```
example% uux "a!diff a!/home/dan/xyz c!/home/dan/xyz > !xyz.diff"
```

will work. (If **diff** is a permitted command.)

Protected files and files that are in protected directories that are owned by the requester can be sent in commands using **uux**. However, if the requester is root, and the directory is not searchable by "other", the request will fail.

NAME	vacation – reply to mail automatically
SYNOPSIS	vacation [-I] vacation [-j] [-a alias] [-tN] <i>username</i>
AVAILABILITY	SUNWcsu
DESCRIPTION	vacation automatically replies to incoming mail.
Installation	<p>The installation consists of an interactive program which sets up vacation's basic configuration.</p> <p>To install vacation, type it with no arguments on the command line. The program creates a .vacation.msg file, which contains the message that is automatically sent to all senders when vacation is enabled, and starts an editor for you to modify the message. (See USAGE section.) Which editor is invoked is determined by the VISUAL or EDITOR environment variable, or vi(1) if neither of those environment variables are set.</p> <p>A .forward file is also created if one does not exist in your home directory. Once created, the .forward file will contain a line of the form:</p> <pre style="margin-left: 40px;">\username, " /usr/bin/vacation username"</pre> <p>One copy of an incoming message is sent to the <i>username</i> and another copy is piped into vacation.</p> <p>If a .forward file is present in your home directory, it will ask whether you want to remove it, which disables vacation and ends the installation.</p> <p>The program automatically creates .vacation.pag and .vacation.dir, which contain a list of senders when vacation is enabled.</p>
Activation and Deactivation	The presence of the .forward file determines whether or not vacation is disabled or enabled. To disable vacation remove the .forward file, or move it to a new name.
Initialization	vacation -I clears the vacation log files, .vacation.pag and .vacation.dir , erasing the list of senders from a previous vacation session. (See OPTIONS section).
Additional Configuration	vacation provides configuration options that are not part of the installation, these being -j , -a , -t . (See OPTIONS section).
OPTIONS	-I Initialize the .vacation.pag and .vacation.dir files and enables vacation . If the -I flag is not specified, and a <i>user</i> argument is given, vacation reads the first line from the standard input (for a From: line, no colon). If absent, it produces an error message.

Options **-j**, **-a**, **-t** are configuration options to be used in conjunction with **vacation** in the **.forward** file, not on the command line. For example,

```
\username, " |/usr/bin/vacation -t1m username"
```

repeats replies to the sender every minute.

-j Do not check whether the recipient appears in the **To:** or the **Cc:** line.

-a alias Indicate that *alias* is one of the valid aliases for the user running **vacation**, so that mail addressed to that alias generates a reply.

-tN Change the interval between repeat replies to the same sender. The default is 1 week. A trailing **s**, **m**, **h**, **d**, or **w** scales *N* to seconds, minutes, hours, days, or weeks respectively.

USAGE

Files

.vacation.msg should include a header with at least a **Subject:** line (it should not include a **From:** or a **To:** line). For example:

```
Subject: I am on vacation
I am on vacation until July 22. If you have something urgent,
please contact Joe Jones (jones@fB0).
--John
```

If the string **SSUBJECT** appears in the **.vacation.msg** file, it is replaced with the subject of the original message when the reply is sent; thus, a **.vacation.msg** file such as

```
Subject: I am on vacation
I am on vacation until July 22.
Your mail regarding "$SUBJECT" will be read when I return.
If you have something urgent, please contact
Joe Jones (jones@fB0).
--John
```

will include the subject of the message in the reply.

No message is sent if the **To:** or the **Cc:** line does not list the user to whom the original message was sent or one of a number of aliases for them, if the initial **From** line includes the string **-REQUEST@**, or if a **Precedence: bulk** or **Precedence: junk** line is included in the header.

vacation will also not respond to mail from either **postmaster** or **Mailer-Daemon**.

FILES

```
~/forward
~/vacation.msg
```

A list of senders is kept in the **dbm** format files **.vacation.pag** and **.vacation.dir** in your home directory. These files are **dbm** files and cannot be viewed directly with text editors.

SEE ALSO

vi(1), **sendmail(1M)**, **dbm(3B)**, **aliases(4)**

NAME	vc – version control
SYNOPSIS	vc [-a] [-t] [-cchar] [-s] [<i>keyword=value ... keyword=value</i>]
DESCRIPTION	<p>This command is obsolete and will be removed in the next release.</p> <p>The vc command copies lines from the standard input to the standard output under control of its arguments and of “control statements” encountered in the standard input. In the process of performing the copy operation, user-declared <i>keywords</i> may be replaced by their string <i>value</i> when they appear in plain text and/or control statements.</p> <p>The copying of lines from the standard input to the standard output is conditional, based on tests (in control statements) of keyword values specified in control statements or as vc command arguments.</p> <p>A control statement is a single line beginning with a control character, except as modified by the -t keyletter (see below). The default control character is colon (:), except as modified by the -c keyletter (see below). Input lines beginning with a backslash (\) followed by a control character are not control lines and are copied to the standard output with the backslash removed. Lines beginning with a backslash followed by a non-control character are copied in their entirety.</p> <p>A keyword is composed of 9 or less alphanumeric characters; the first must be alphabetic. A value is any ASCII string that can be created with ed; a numeric value is an unsigned string of digits. Keyword values may not contain blanks or tabs.</p> <p>Replacement of keywords by values is done whenever a keyword surrounded by control characters is encountered on a version control statement. The -a keyletter (see below) forces replacement of keywords in all lines of text. An uninterpreted control character may be included in a value by preceding it with \. If a literal \ is desired, then it too must be preceded by \.</p>
OPTIONS	<p>-a Forces replacement of keywords surrounded by control characters with their assigned value in all text lines and not just in vc statements.</p> <p>-t All characters from the beginning of a line up to and including the first tab character are ignored for the purpose of detecting a control statement. If a control statement is found, all characters up to and including the tab are discarded.</p> <p>-cchar Specifies a control character to be used in place of the “:” default.</p> <p>-s Silences warning messages (not error) that are normally printed on the diagnostic output.</p> <p>vc recognizes the following version control statements:</p> <p>:dcl <i>keyword</i> [, ..., <i>keyword</i>] Declare keywords. All keywords must be declared.</p>

:asg *keyword=value*

Assign values to keywords. An **asg** statement overrides the assignment for the corresponding keyword on the **vc** command line and all previous **asg** statements for that keyword. Keywords that are declared but are not assigned values have null values.

:if *condition*

...

:end

Skip lines of the standard input. If the condition is true, all lines between the **if** statement and the matching **end** statement are copied to the standard output. If the condition is false, all intervening lines are discarded, including control statements. Note: Intervening **if** statements and matching **end** statements are recognized solely for the purpose of maintaining the proper **if-end** matching.

The syntax of a condition is:

```

<cond> ::= [ "not" ] <or>
<or>   ::= <and> / <and> " | " <or>
<and>  ::= <exp> / <exp> "&" <and>
<exp>  ::= "(" <or> ")" / <value> <op> <value>
<op>   ::= "=" / "!=" / "<" / ">"
<value> ::= <arbitrary ASCII string> / <numeric string>

```

The available operators and their meanings are:

=	equal
!=	not equal
&	and
	or
>	greater than
<	less than
()	used for logical groupings
not	may only occur immediately after the if , and when present, inverts the value of the entire condition

The > and < operate only on unsigned integer values (for example, : **012** > **12** is false). All other operators take strings as arguments (for example, : **012** != **12** is true).

The precedence of the operators (from highest to lowest) is:

```

= != > <   all of equal precedence
&
|

```

Parentheses may be used to alter the order of precedence.

Values must be separated from operators or parentheses by at least one blank or tab.

::text Replace keywords on lines that are copied to the standard output. The two leading control characters are removed, and keywords surrounded by control characters in text are replaced by their value before the line is copied to the output file. This action is independent of the **-a** keyletter.

:on

:off Turn on or off keyword replacement on all lines.

:ctl *char*

Change the control character to *char*.

:msg *message*

Print *message* on the diagnostic output.

:err *message*

Print *message* followed by:

ERROR: err statement on line ... (915)

on the diagnostic output. **vc** halts execution, and returns an exit code of 1.

SEE ALSO

ed(1)

NAME	vgrind – grind nice program listings
SYNOPSIS	vgrind [-2fntwWx] [-d <i>defs-file</i>] [-h <i>header</i>] [-l <i>language</i>] [-sn] [-opagelist] [-P <i>printer</i>] [-T <i>output-device</i>] <i>filename...</i>
AVAILABILITY	SUNWdoc
DESCRIPTION	<p>vgrind formats the program sources named by the <i>filename</i> arguments in a nice style using troff(1). Comments are placed in italics, keywords in bold face, and as each function is encountered its name is listed on the page margin.</p> <p>vgrind runs in two basic modes, filter mode or regular mode. In filter mode vgrind acts as a filter in a manner similar to tbl(1). The standard input is passed directly to the standard output except for lines bracketed by the troff-like macros:</p> <pre style="margin-left: 40px;">.vS starts processing .vE ends processing</pre> <p>These lines are formatted as described above. The output from this filter can be passed to troff for output. There need be no particular ordering with eqn(1) or tbl.</p> <p>In regular mode vgrind accepts input <i>filenames</i>, processes them, and passes them to troff for output. If no <i>filename</i> is given, or if the ‘-’ argument is given, vgrind reads from the standard input (default if -f is specified).</p> <p>In both modes vgrind passes any lines beginning with a decimal point without conversion.</p>
OPTIONS	<p>Note: The syntax of options with arguments is important. Some require a SPACE between the option name and the argument, while those that do not have a SPACE below will not tolerate one.</p> <pre style="margin-left: 20px;">-2 Produce two column output. Specifying this option changes the default point size to 8 (as if the -s8 option were supplied). It also arranges for output to appear in landscape mode, by supplying the -L flag to the for- matter and changing the page height and width accordingly. -f Force filter mode. -n Do not make keywords boldface. -w Consider TAB characters to be spaced four columns apart instead of the usual eight. -x Output the index file in a “pretty” format. The index file itself is pro- duced whenever vgrind is run with a file called index present in the current directory. The index of function definitions can then be run off by giving vgrind the -x option and the file index as argument.</pre>

- d** *defs-file* Specify an alternate language definitions file (default is `/usr/lib/vgrindefs`).
- h** *header* Specify a header to appear in the center of every output page.
- l** *language* Specify the language to use. Among the languages currently known are: Bourne shell (**-lsh**), C (**-lc**, the default), C++ (**-lc++**), C shell (**-lcs**), emacs MLisp, (**-lml**), FORTRAN (**-lf**), Icon (**-II**), ISP (**-i**), LDL (**-ILD**), Model (**-lm**), Pascal (**-lp**), and RATFOR (**-lr**).
- sn** Specify a point size to use on output (exactly the same as the argument of a **troff** **.ps** point size request).

vgrind passes the following options to the formatter specified by the **TROFF** environment variable, see **ENVIRONMENT** below.

- t** Similar to the same option in **troff**; that is, formatted text goes to the standard output.
- W** Force output to the (wide) Versatec printer rather than the (narrow) Varian.
- opagelist**
Print only those pages whose page numbers appear in the comma-separated *pagelist* of numbers and ranges. A range *N-M* means pages *N* through *M*; an initial *-N* means from the beginning to page *N*; and a final *N-* means from *N* to the end.
- Pprinter**
Send output to the named *printer*.
- Toutput-device**
Format output for the specified *output-device*.

ENVIRONMENT

In regular mode **vgrind** feeds its intermediate output to the text formatter given by the value of the **TROFF** environment variable, or to **troff** if this variable is not defined in the environment. This mechanism allows for local variations in **troff**'s name.

FILES

index	file where source for index is created
<code>/usr/lib/vgrindefs</code>	language descriptions
<code>/usr/lib/vfontedpr</code>	preprocessor
<code>/usr/share/lib/tmac/tmac.vgrind</code>	macro package

SEE ALSO

troff(1)

BUGS

vgrind assumes that a certain programming style is followed:

- C** Function names can be preceded on a line only by SPACE, TAB, or an asterisk. The parenthesized arguments must also be on the same line.
- FORTRAN** Function names need to appear on the same line as the keywords *function* or *subroutine*.
- MLisp** Function names should not appear on the same line as the preceding *defun*.
- Model** Function names need to appear on the same line as the keywords *is beginproc*.

Pascal Function names need to appear on the same line as the keywords *function* or *procedure*.

If these conventions are not followed, the indexing and marginal function name comment mechanisms will fail.

More generally, arbitrary formatting styles for programs mostly look bad. The use of SPACE characters to align source code fails miserably; if you plan to **vgrind** your program you should use TAB characters. This is somewhat inevitable since the fonts **vgrind** uses are variable width.

The mechanism of **ctags**(1) in recognizing functions should be used here.

The **-w** option is a crock, but there is no other way to achieve the desired effect.

The macros defined in **tmac.vgrind** do not coexist gracefully with those of other macro packages, making filter mode difficult to use effectively.

vgrind does not process certain special characters in **cs**h(1) scripts correctly.

The **tmac.vgrind** formatting macros wire in the page height and width used in two column mode, effectively making two column output useless for paper sizes other than the standard American size of 8.5 by 11 inches. For other paper sizes, it is necessary to edit the size values given in **tmac.vgrind**. A better solution would be to create a **troff** output device specification intended specifically for landscape output and record size information there.

NAME	vi, view, vedit – screen-oriented (visual) display editor based on ex
SYNOPSIS	<pre>vi [- -s] [-I] [-L] [-R] [-r [filename]] [-t tag] [-v] [-V] [-x] [-wn] [-C] [+command -c command] filename. . . view [- -s] [-I] [-L] [-R] [-r [filename]] [-t tag] [-v] [-V] [-x] [-wn] [-C] [+command -c command] filename. . . vedit [- -s] [-I] [-L] [-R] [-r [filename]] [-t tag] [-v] [-V] [-x] [-wn] [-C] [+command -c command] filename. . .</pre>
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>vi (visual) is a display-oriented text editor based on an underlying line editor ex. It is possible to use the command mode of ex from within vi and vice-versa. The visual commands are described on this manual page; how to set options (like automatically numbering lines and automatically starting a new output line when you type carriage return) and all ex line editor commands are described on the ex(1) manual page.</p> <p>When using vi, changes you make to the file are reflected in what you see on your terminal screen. The position of the cursor on the screen indicates the position within the file.</p>
OPTIONS Invocation Options	<p>The following invocation options are interpreted by vi (previously documented options are discussed in the NOTES section of this manual page):</p> <ul style="list-style-type: none"> - -s Suppress all interactive user feedback. This is useful when processing editor scripts. -I Set up for editing LISP programs. -L List the name of all files saved as the result of an editor or system crash. -R Readonly mode; the readonly flag is set, preventing accidental overwriting of the file. -r filename Edit <i>filename</i> after an editor or system crash. (Recovers the version of <i>filename</i> that was in the buffer when the crash occurred.) -t tag Edit the file containing the <i>tag</i> and position the editor at its definition. -v Start up in display editing state using vi. You can achieve the same effect by simply typing the -vi command itself. -V Verbose. Any non-tty input will be echoed on standard error. This may be useful when processing editor commands within shell scripts. -x Encryption option; when used, vi simulates the X command of ex and prompts the user for a key. This key is used to encrypt and decrypt text using the algorithm of the crypt command. The X command makes an educated guess to determine whether text read in is encrypted or not. The temporary buffer file is encrypted also, using a transformed version of the key typed in for the -x option.

- wn** Set the default window size to *n*. This is useful when using the editor over a slow speed line.
- C** Encryption option; same as the **-x** option, except that **vi** simulates the **C** command of **ex**. The **C** command is like the **X** command of **ex**, except that all text read in is assumed to have been encrypted.
- +command | -c command** Begin editing by executing the specified editor *command* (usually a search or positioning command).

The *filename* argument indicates one or more files to be edited.

The *view* invocation is the same as **vi** except that the **readonly** flag is set.

The *vedit* invocation is intended for beginners. It is the same as **vi** except that the **report** flag is set to 1, the **showmode** and **novice** flags are set, and **magic** is turned off. These defaults make it easier to learn how to use **vi**.

vi Modes

- | | |
|-----------|---|
| Command | Normal and initial mode. Other modes return to command mode upon completion. ESC (escape) is used to cancel a partial command. |
| Input | Entered by setting any of the following options: a A i I o O c C s S R . Arbitrary text may then be entered. Input mode is normally terminated with ESC character, or, abnormally, with an interrupt. |
| Last line | Reading input for : ? or ! ; terminate by typing a carriage return; an interrupt cancels termination. |

COMMAND SUMMARY

Sample commands

In the descriptions, **CR** stands for carriage return and **ESC** stands for the escape key.

- | | |
|-----------------|--|
| ← ↓ ↑ → | arrow keys move the cursor |
| h j k l | same as arrow keys |
| itextESC | insert <i>text</i> |
| cwnewESC | change word to <i>new</i> |
| easESC | pluralize word (end of word; append s ;
escape from input state) |
| x | delete a character |
| dw | delete a word |
| dd | delete a line |
| 3dd | delete 3 lines |
| u | undo previous change |
| ZZ | exit vi , saving changes |
| :q!CR | quit, discarding changes |
| /textCR | search for <i>text</i> |
| ^U ^D | scroll up or down |
| :cmdCR | any ex or ed command |

Counts before vi commands

Numbers may be typed as a prefix to some commands. They are interpreted in one of these ways.

line/column number **z G |**

	scroll amount	^D ^U
	repeat effect	most of the rest
Interrupting, canceling	ESC	end insert or incomplete cmd
	DEL	(delete or rubout) interrupts
File manipulation	ZZ	if file modified, write and exit; otherwise, exit
	:wCR	write back changes
	:w!CR	forced write, if permission originally not valid
	:qCR	quit
	:q!CR	quit, discard changes
	:e nameCR	edit file <i>name</i>
	:e!CR	reedit, discard changes
	:e + nameCR	edit, starting at end
	:e +nCR	edit starting at line <i>n</i>
	:e #CR	edit alternate file
	:e! #CR	edit alternate file, discard changes
	:w nameCR	write file <i>name</i>
	:w! nameCR	overwrite file <i>name</i>
	:shCR	run shell, then return
	:! cmdCR	run <i>cmd</i> , then return
	:nCR	edit next file in arglist
	:n argsCR	specify new arglist
	^G	show current file and line
	:ta tagCR	position cursor to <i>tag</i>

In general, any **ex** or **ed** command (such as *substitute* or *global*) may be typed, preceded by a colon and followed by a carriage return.

Positioning within file	^F	forward screen
	^B	backward screen
	^D	scroll down half screen
	^U	scroll up half screen
	nG	go to the beginning of the specified line (end default), where <i>n</i> is a line number
	/pat	next line matching <i>pat</i>
	?pat	previous line matching <i>pat</i>
	n	repeat last / or ? command
	N	reverse last / or ? command
	/pat/+n	<i>n</i> th line after <i>pat</i>
	?pat?-n	<i>n</i> th line before <i>pat</i>
]]	next section/function
	[[previous section/function
	(beginning of sentence
)	end of sentence
	{	beginning of paragraph

	}	end of paragraph
	%	find matching () { or }
Adjusting the screen	^L	clear and redraw window
	^R	clear and redraw window if ^L is → key
	zCR	redraw screen with current line at top of window
	z-CR	redraw screen with current line at bottom of window
	z.CR	redraw screen with current line at center of window
	/pat/z-CR	move <i>pat</i> line to bottom of window
	zn.CR	use <i>n</i> -line window
	^E	scroll window down 1 line
	^Y	scroll window up 1 line
Marking and returning	``	move cursor to previous context
	''	move cursor to first non-white space in line
	mx	mark current position with the ASCII lower-case letter <i>x</i>
	`x	move cursor to mark <i>x</i>
	^x	move cursor to first non-white space in line marked by <i>x</i>
Line positioning	H	top line on screen
	L	last line on screen
	M	middle line on screen
	+	next line, at first non-white
	-	previous line, at first non-white
	CR	return, same as +
	↓ or j	next line, same column
	↑ or k	previous line, same column
Character positioning	^	first non white-space character
	0	beginning of line
	\$	end of line
	l or →	forward
	h or ←	backward
	^H	same as ← (backspace)
	space	same as → (space bar)
	fx	find next <i>x</i>
	Fx	find previous <i>x</i>
	tx	move to character prior to next <i>x</i>
	Tx	move to character following previous <i>x</i>
	;	repeat last f , F , t , or T
	,	repeat inverse of last f , F , t , or T
	<i>n</i>	move to column <i>n</i>
	%	find matching ({) or }

**Words, sentences,
paragraphs**

w forward a word
b back a word
e end of word
) to next sentence
} to next paragraph
(back a sentence
{ back a paragraph
W forward a blank-delimited word
B back a blank-delimited word
E end of a blank-delimited word

**Corrections during
insert**

^H erase last character (backspace)
^W erase last word
erase your erase character, same as **^H** (backspace)
kill your kill character, erase this line of input
**** quotes your erase and kill characters
ESC ends insertion, back to command mode
CTRL-C interrupt, suspends insert mode
^D backtab one character; reset left margin of *autoindent*
^^D caret (^) followed by control-d (^D);
backtab to beginning of line;
do not reset left margin of *autoindent*
0^D backtab to beginning of line; reset left margin of *autoindent*
^V quote non-printable character

Insert and replace

a append after cursor
A append at end of line
i insert before cursor
I insert before first non-blank
o open line below
O open above
rx replace single char with x
RtextESC replace characters

Operators

Operators are followed by a cursor motion, and affect all text that would have been moved over. For example, since **w** moves over a word, **dw** deletes the word that would be moved over. Double the operator, for example, **dd** to affect whole lines.

d delete
c change
y yank lines to buffer
< left shift
> right shift
! filter through command

**Miscellaneous
Operations**

C change rest of line (**c\$**)
D delete rest of line (**d\$**)
s substitute chars (**cl**)
S substitute lines (**cc**)
J join lines
x delete characters (**dl**)
X delete characters before cursor (**dh**)
Y yank lines (**yy**)

Yank and Put

Put inserts the text most recently deleted or yanked; however, if a buffer is named (using the ASCII lower-case letters **a - z**), the text in that buffer is put instead.

3yy yank 3 lines
3yl yank 3 characters
p put back text after cursor
P put back text before cursor
"xp put from buffer *x*
"xy yank to buffer *x*
"xd delete into buffer *x*

Undo, Redo, Retrieve

u undo last change
U restore current line
. repeat last change
"dp retrieve *d*'th last delete

AUTHOR

vi and **ex** were developed by The University of California, Berkeley California, Computer Science Division, Department of Electrical Engineering and Computer Science.

ENVIRONMENT

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

LC_CTYPE

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

LC_TIME

Determines how **vi** handles date and time formats. In the "C" locale, date and time handling follows the U.S. rules.

FILES	<p>/tmp default directory where temporary work files are placed; it can be changed using the directory option (see the ex(1) set command)</p> <p>/usr/share/lib/terminfo/?/* compiled terminal description database</p> <p>/usr/lib/.COREterm/?/* subset of compiled terminal description database</p>
SEE ALSO	<p>intro(1), ed(1), edit(1), ex(1), environ(5)</p> <p><i>Solaris Advanced User's Guide</i></p>
NOTES	<p>Two options, although they continue to be supported, have been replaced in the documentation by options that follow the Command Syntax Standard (see intro(1)). A -r option that is not followed with an option-argument has been replaced by -L and +command has been replaced by -c command.</p> <p>The message file too large to recover with -r option, which is seen when a file is loaded, indicates that the file can be edited and saved successfully, but if the editing session is lost, recovery of the file with the -r option will not be possible.</p> <p>The encryption options are provided with the Security Administration Utilities package, which is available only in the United States.</p> <p>The editing environment defaults to certain configuration options. When an editing session is initiated, vi attempts to read the EXINIT environment variable. If it exists, the editor uses the values defined in EXINIT, otherwise the values set in \$HOME/.exrc are used. If \$HOME/.exrc does not exist, the default values are used.</p> <p>To use a copy of .exrc located in the current directory other than \$HOME, set the exrc option in EXINIT or \$HOME/.exrc. Options set in EXINIT can be turned off in a local .exrc only if exrc is set in EXINIT or \$HOME/.exrc.</p> <p>Tampering with entries in /usr/share/lib/terminfo/?/* or /usr/share/lib/terminfo/?/* (for example, changing or removing an entry) can affect programs such as vi that expect the entry to be present and correct. In particular, removing the "dumb" terminal may cause unexpected problems.</p> <p>Software tabs using ~T work only immediately after the <i>autoindent</i>.</p> <p>Left and right shifts on intelligent terminals do not make use of insert and delete character operations in the terminal.</p>

NAME	vipw – edit the password file
SYNOPSIS	/usr/ucb/vipw
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>vipw edits the password file while setting the appropriate locks, and does any necessary processing after the password file is unlocked. If the password file is already being edited, then you will be told to try again later. The vi(1) editor will be used unless the environment variable VISUAL or EDITOR indicates an alternate editor.</p> <p>vipw performs a number of consistency checks on the password entry for root, and will not allow a password file with a “mangled” root entry to be installed. It also checks the /etc/shells file to verify the login shell for root.</p>
FILES	/etc/ptmp /etc/shells
SEE ALSO	passwd(1) , vi(1) , passwd(4)

NAME	volcancel – cancel user's request for removable media that is not currently in drive
SYNOPSIS	<code>/usr/lib/vold/volcancel [-n] [volume]</code>
DESCRIPTION	<p>volcancel cancels a user's request to access a particular floppy or CD-ROM file system. This command is useful when the removable media containing the file system is not currently in the drive.</p> <p>Use the path <code>/vol/rdisk/name_of_volume</code> to specify the volume. If called without a volume name to cancel, volcancel checks for Volume Management running.</p>
OPTIONS	-n Display the nickname to the device name translation table.
EXAMPLES	<p>To cancel a request to access an unnamed CD-ROM, use</p> <pre>example% /usr/lib/vold/volcancel vol/rdisk/unnamed_cdrom</pre> <p>To check if volume management is running, use:</p> <pre>example% /usr/lib/vold/volcancel echo volmgmt not running</pre>
SEE ALSO	rmmount(1M) , volcheck(1) , vold(1M) , volmissing(1) , rmmount.conf(4) , vold.conf(4) , volfs(7)

NAME	volcheck – check for media in a drive. Default checks all floppy media.
SYNOPSIS	volcheck [-v] [-i secs] [-t secs] <i>pathname</i>
DESCRIPTION	volcheck tells Volume Management to look at each <i>pathname</i> in sequence and determine if new media has been inserted in the drive. The default action is to volcheck all floppy drives pointed to by volume management.
OPTIONS	-v Verbose. -t secs Check the named device(s) for the next <i>secs</i> seconds. The maximum number of seconds allowed is 28800, which is 8 hours. The frequency of checking is specified by -i . There is no default total time. -i secs Set the frequency of device checking to <i>secs</i> seconds. The default is 2 seconds. The minimum frequency is 1 second.
EXAMPLES	example% volcheck -v /dev/diskette /dev/diskette has media asks Volume Management to examine the floppy drive for new media. example% volcheck -i 2 -t 600 /dev/diskette1 & asks Volume Management if there is a floppy in the floppy drive every 2 seconds for 600 seconds (10 minutes).
FILES	/dev/volctl Volume Management control port
SEE ALSO	eject(1), volcancel(1), volmissing(1) rmmount(1M), vold(1M), rmmount.conf(4), vold.conf(4), volfs(7)
WARNINGS	Due to a hardware limitation in many floppy drives, the act of checking for media causes mechanical action in the floppy drive. Continuous polling of the floppy drive will cause the drive to wear out. It is recommended that polling the drive only be performed during periods of high use.

NAME	volmissing – notify user that volume requested is not in the CD-ROM or floppy drive
SYNOPSIS	<code>/usr/lib/vold/volmissing [-c] [-p] [-s] [-m <i>alias</i>]</code>
DESCRIPTION	<p>volmissing informs a user when a requested volume is not available. Depending on the option selected, users are notified through their console window, syslogd(1M), or a mail message.</p> <p>volmissing -p is the default action taken by vold(1M), the Volume Management daemon, when it needs to notify a user that the requested volume is not available. If you want to change this default event, modify the <code>/etc/vold.conf</code> file. See vold.conf(4).</p> <p>You can change the notification method for your system by editing the vold.conf configuration file and providing a new option for volmissing in the notify entry under the Events category.</p>
OPTIONS	<p>-c Send a message to the user's console requesting the volume be inserted. To end the notification without inserting the requested volume, use volcancel(1).</p> <p>-p All volmissing events will be handled through a GUI, provided a window system is running on the console. If this option is specified, and no window system is running, all messages go to the system console.</p> <p>-s Send one message to the syslogd(1M).</p> <p>-m <i>alias</i> Send a mail message to the specified mail alias about the missing volume.</p>
FILES	<p><code>/etc/vold.conf</code> Volume Management daemon configuration file. Directs the Volume Management daemon to control certain devices, and causes action to be taken when specific criteria is met.</p> <p><code>/usr/lib/vold/volmissing_popup</code> Pop-up used when the -p option is supplied and a window system is running.</p>
SEE ALSO	volcancel (1), volcheck (1), rmmount (1M), syslogd (1M), vold (1M), rmmount.conf (4), vold.conf (4), volfs (7)

NAME	vsig – synchronize a co-process with the controlling FMLI application
SYNOPSIS	vsig
AVAILABILITY	SUNWesu
DESCRIPTION	The vsig executable sends a SIGUSR2 signal to the controlling FMLI process. This signal/alarm causes FMLI to execute the FMLI built-in command checkworld which causes all posted objects with a reread descriptor evaluating to TRUE to be reread. vsig takes no arguments.
EXAMPLES	<p>The following is a segment of a shell program:</p> <pre> echo "Sending this string to an FMLI process" vsig</pre> <p>The vsig executable will flush the output buffer <i>before</i> it sends the SIGUSR2 signal to make sure the string is actually in the pipe created by the cocreate function.</p>
SEE ALSO	coproc(1F) , kill(1) , kill(2) , signal(3C)
NOTES	Because vsig synchronize with FMLI, it should be used rather than kill to send a SIGUSR2 signal to FMLI.

NAME	w – who is logged in, and what are they doing
SYNOPSIS	w [-hlsuw] [<i>user</i>]
AVAILABILITY	SUNWcsu
DESCRIPTION	<p>The w command displays a summary of the current activity on the system, including what each user is doing. The heading line shows the current time, the length of time the system has been up, the number of users logged into the system and the average number of jobs in the run queue over the last 1, 5 and 15 minutes.</p> <p>The fields displayed are: the users login name, the name of the tty the user is on, the time of day the user logged on (in <i>hours:minutes</i>), the idle time—that is, the number of minutes since the user last typed anything (in <i>hours:minutes</i>), the CPU time used by all processes and their children on that terminal (in <i>minutes:seconds</i>), the CPU time used by the currently active processes (in <i>minutes:seconds</i>), the name and arguments of the current process.</p> <p>If a <i>user</i> name is included, output is restricted to that user.</p>
OPTIONS	<p>-h Suppress the heading.</p> <p>-l Produce a long form of output, which is the default.</p> <p>-s Produce a short form of output. In the short form, the tty is abbreviated, the login time and CPU times are left off, as are the arguments to commands.</p> <p>-u Produces the heading line which shows the current time, the length of time the system has been up, the number of users logged into the system, and the average number of jobs in the run queue over the last 1, 5 and 15 minutes.</p> <p>-w Produces a long form of output, which is also the same as the default.</p>
EXAMPLE	<pre>example% w 10:54am up 27 day(s), 57 mins, 1 user, load average: 0.28, 0.26, 0.22 User tty login@ idle JCPU PCPU what ralph console 7:10am 1 10:05 4:31 w</pre>
ENVIRONMENT	<p>If any of the LC_* variables (LC_CTYPE, LC_MESSAGES, LC_TIME, LC_COLLATE, LC_NUMERIC, and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of tar for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how tar behaves.</p> <p>LC_CTYPE</p> <p>Determines how tar handles characters. When LC_CTYPE is set to a valid value, tar can display and handle text and filenames containing valid characters for that locale. tar can display and handle Extended Unix code (EUC) characters where any individual character can be 1, 2, or 3 bytes wide. tar can also handle EUC</p>

characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid.

LC_MESSAGES

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

LC_TIME

Determines how **tar** handles date and time formats. In the "C" locale, date and time handling follow the U.S. rules.

FILES

/var/adm/utmp

SEE ALSO

ps(1), **who(1)**, **whodo(1M)**, **utmp(4)**

NOTES

The notion of the "current process" is muddy. The current algorithm is 'the highest numbered process on the terminal that is not ignoring interrupts, or, if there is none, the highest numbered process on the terminal'. This fails, for example, in critical sections of programs like the shell and editor, or when faulty programs running in the background fork and fail to ignore interrupts. In cases where no process can be found, **w** prints -.

The CPU time is only an estimate, in particular, if someone leaves a background process running after logging out, the person currently on that terminal is "charged" with the time.

Background processes are not shown, even though they account for much of the load on the system.

Sometimes processes, typically those in the background, are printed with null or garbaged arguments. In these cases, the name of the command is printed in parentheses.

w does not know about the conventions for detecting background jobs. It will sometimes find a background job instead of the right one.

NAME	wait – shell built-in function to wait for other jobs or processes
SYNOPSIS	
sh	wait [<i>n</i>] wait [% <i>jobid</i> ...]
cs h	wait [<i>n</i>]
ksh	† wait [<i>job</i>]
DESCRIPTION	
sh	<p>Wait for your background process whose process id is <i>n</i> and report its termination status. If <i>n</i> is omitted, all your shell's currently active background processes are waited for and the return code will be zero. wait accepts a job identifier, when Job Control is enabled, and the argument, <i>jobid</i>, is preceded by a percent-sign.</p> <p>The shell itself executes wait, without creating a new process. If you get the error message cannot fork, too many processes, try using the wait command to clean up your background processes. If this doesn't help, the system process table is probably full or you have too many active foreground processes. (There is a limit to the number of process ids associated with your login, and to the number the system can keep track of.)</p> <p>Not all the processes of a 3- or more-stage pipeline are children of the shell, and thus cannot be waited for.</p> <p>If <i>n</i> is not an active process id, all your shell's currently active background processes are waited for and the return code will be zero.</p>
cs h	<p>Wait for your background process whose process id is <i>n</i> and report its termination status. If <i>n</i> is omitted, all your shell's currently active background processes are waited for and the return code will be zero.</p> <p>The shell itself executes wait, without creating a new process. If you get the error message cannot fork, too many processes, try using the wait command to clean up your background processes. If this doesn't help, the system process table is probably full or you have too many active foreground processes. (There is a limit to the number of process ids associated with your login, and to the number the system can keep track of.)</p> <p>Not all the processes of a 3- or more-stage pipeline are children of the shell, and thus cannot be waited for.</p> <p>If <i>n</i> is not an active process id, all your shell's currently active background processes are waited for and the return code will be zero.</p>
ksh	<p>Wait for the specified <i>job</i> and report its termination status. If <i>job</i> is not given then all currently active child processes are waited for. The exit status from this command is that of the process waited for.</p>

On this man page, **ksh**(1) commands that are preceded by one or two † (daggers) are treated specially in the following ways:

1. Variable assignment lists preceding the command remain in effect when the command completes.
2. I/O redirections are processed after variable assignments.
3. Errors cause a script that contains them to abort.
4. Words, following a command preceded by †† that are in the format of a variable assignment, are expanded with the same rules as a variable assignment. This means that tilde substitution is performed after the = sign and word splitting and file name generation are not performed.

SEE ALSO**cs****h**(1), **jobs**(1), **ksh**(1), **sh**(1)

NAME	wc – display a count of lines, words and characters in a file
SYNOPSIS	wc [-cClw] [<i>name</i> ...]
AVAILABILITY	SUNWcsu
DESCRIPTION	wc counts lines, words, and characters in the named files, or in the standard input if no <i>names</i> appear. It also keeps a total count for all named files. A word is a string of characters delimited by a SPACE, TAB, or by any other character in the library function iswspace() (see iswalph(3I)).
OPTIONS	When <i>name</i> is specified on the command line, the names are printed along with the counts. If no option is specified the default is -lwc (count lines, words, and bytes.) -c Count bytes. -C Count characters. -l Count lines. -w Count words delimited by white space characters or new line characters. Delimiting characters are Extended Unix Code (EUC) characters from any code set defined by iswspace() .
ENVIRONMENT	If any of the LC_* variables (LC_CTYPE , LC_MESSAGES , LC_TIME , LC_COLLATE , LC_NUMERIC , and LC_MONETARY) (see environ(5)) are not set in the environment, the operational behavior of wc for each corresponding locale category is determined by the value of the LANG environment variable. If LC_ALL is set, its contents are used to override both the LANG and the other LC_* variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how wc behaves. LC_CTYPE Determines how wc handles characters. When LC_CTYPE is set to a valid value, wc can display and handle text and filenames containing valid characters for that locale. wc can display and handle EUC characters where any individual character can be 1, 2, or 3 bytes wide. wc can also handle EUC characters of 1, 2, or more column widths. In the "C" locale, only characters from ISO 8859-1 are valid. LC_MESSAGES Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).
SEE ALSO	setlocale(3C) , iswalph(3I) , environ(5)

NAME	what – extract SCCS version information from a file
SYNOPSIS	what [-s] <i>filename</i> . . .
DESCRIPTION	what searches each <i>filename</i> for occurrences of the pattern @(#) that the SCCS get command (see sccs-get (1)) substitutes for the %Z% ID keyword, and prints what follows up to a " , > , NEWLINE , \ , or null character.
OPTIONS	-s Stop after the first occurrence of the pattern.
EXAMPLES	For example, if a C program in file program.c contains <pre>char sccsid[] = "@(#)identification information";</pre> and program.c is compiled to yield program.o and a.out , the command: <pre>example% what program.c program.o a.out</pre> produces: <pre>program.c: identification information program.o: identification information a.out: identification information</pre>
SEE ALSO	sccs (1), sccs-admin (1), sccs-cdc (1), sccs-comb (1), sccs-delta (1), sccs-get (1), sccs-help (1), sccs-prs (1), sccs-prt (1), sccs-rmdel (1), sccs-sact (1), sccs-sccsdiff (1), sccs-unget (1), sccs-val (1), sccsfile (4) <i>Programming Utilities Guide</i>
DIAGNOSTICS	Use the SCCS help command for explanations (see sccs-help (1)).
BUGS	There is a remote possibility that a spurious occurrence of the '@(#)' pattern could be found by what .

NAME	whatis – display a one-line summary about a keyword
SYNOPSIS	whatis <i>command</i> ...
AVAILABILITY	SUNWdoc
DESCRIPTION	<p>whatis looks up a given <i>command</i> and displays the header line from the manual section. You can then run the man(1) command to get more information. If the line starts 'name(<i>section</i>) ...' you can do 'man -s section name' to get the documentation for it. Try 'whatis ed' and then you should do 'man -s 1 ed' to get the manual page for ed(1).</p> <p>whatis is actually just the -f option to the man(1) command.</p> <p>whatis uses the /usr/share/man/windex database. This database is created by catman(1M). If this database does not exist, whatis will fail.</p>
FILES	/usr/share/man/windex table of contents and keyword database
SEE ALSO	apropos (1), man (1), catman (1M)

NAME	whereis – locate the binary, source, and manual page files for a command
SYNOPSIS	<code>/usr/ucb/whereis [-bmsu] [-BMS directory... -f] filename ...</code>
AVAILABILITY	SUNWscpu
DESCRIPTION	<p>whereis locates source/binary and manuals sections for specified files. The supplied names are first stripped of leading pathname components and any (single) trailing extension of the form <i>.ext</i>, for example, <i>.c</i>. Prefixes of <i>s.</i> resulting from use of source code control are also dealt with. whereis then attempts to locate the desired program in a list of standard places:</p> <ul style="list-style-type: none"> <code>/usr/bin</code> <code>/usr/bin</code> <code>/usr/5bin</code> <code>/usr/games</code> <code>/usr/hosts</code> <code>/usr/include</code> <code>/usr/local</code> <code>/usr/etc</code> <code>/usr/lib</code> <code>/usr/share/man</code> <code>/usr/src</code> <code>/usr/ucb</code>
OPTIONS	<ul style="list-style-type: none"> -b Search only for binaries. -m Search only for manual sections. -s Search only for sources. -u Search for unusual entries. A file is said to be unusual if it does not have one entry of each requested type. Thus 'whereis -m -u *' asks for those files in the current directory which have no documentation. -B Change or otherwise limit the places where whereis searches for binaries. -M Change or otherwise limit the places where whereis searches for manual sections. -S Change or otherwise limit the places where whereis searches for sources. -f Terminate the last directory list and signals the start of file names, and <i>must</i> be used when any of the -B, -M, or -S options are used.

- EXAMPLES** Find all files in **/usr/bin** which are not documented in **/usr/share/man/man1** with source in **/usr/src/cmd**:
- ```
example% cd /usr/ucb
example% whereis -u -M /usr/share/man/man1 -S /usr/src/cmd -f *
```
- FILES** **/usr/src/\***  
**/usr/{doc,man}/\***  
**/etc, /usr/{lib,bin,ucb,old,new,local}**
- SEE ALSO** **chdir(2)**
- BUGS** Since **whereis** uses **chdir(2)** to run faster, pathnames given with the **-M**, **-S**, or **-B** must be full; that is, they must begin with a **'/'**.

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
|---------------------|------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>         | which – locate a command; display its pathname or alias                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>SYNOPSIS</b>     | <b>which</b> [ <i>filename</i> ] . . .                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>AVAILABILITY</b> | SUNWcsu                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>DESCRIPTION</b>  | <b>which</b> takes a list of names and looks for the files which would be executed had these names been given as commands. Each argument is expanded if it is aliased, and searched for along the user's path. Both aliases and path are taken from the user's <b>.cshrc</b> file.                                                                                                                                                                                                 |
| <b>FILES</b>        | <b>~/.cshrc</b> source of aliases and path values                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>SEE ALSO</b>     | <b>csh(1)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>DIAGNOSTICS</b>  | A diagnostic is given for names which are aliased to more than a single word, or if an executable file with the argument name was not found in the path.                                                                                                                                                                                                                                                                                                                           |
| <b>BUGS</b>         | Only aliases and paths from <b>~/.cshrc</b> are used; importing from the current environment is not attempted. Must be executed by <b>csh(1)</b> , since only <b>csh</b> knows about aliases.<br>To compensate for <b>~/.cshrc</b> files in which aliases depend upon the <b>prompt</b> variable being set, <b>which</b> sets this variable. If the <b>~/.cshrc</b> produces output or prompts for input when <b>prompt</b> is set, <b>which</b> may produce some strange results. |

|                    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|--------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>        | while, until – shell built-in functions to repetitively execute a set of actions while/until conditions are evaluated TRUE                                                                                                                                                                                                                                                                                                                                                 |
| <b>SYNOPSIS</b>    |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| sh                 | <b>while</b> <i>conditions</i> ; <b>do</b> <i>actions</i> ; <b>done</b><br><b>until</b> <i>conditions</i> ; <b>do</b> <i>actions</i> ; <b>done</b>                                                                                                                                                                                                                                                                                                                         |
| csch               | <b>while</b> ( <i>expr</i> )<br>...<br><b>end</b>                                                                                                                                                                                                                                                                                                                                                                                                                          |
| ksh                | <b>while</b> <i>conditions</i> ; <b>do</b> <i>actions</i> ; <b>done</b><br><b>until</b> <i>conditions</i> ; <b>do</b> <i>actions</i> ; <b>done</b>                                                                                                                                                                                                                                                                                                                         |
| <b>DESCRIPTION</b> |                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| sh                 | A <b>while</b> command repeatedly executes the <b>while</b> <i>conditions</i> and, if the exit status of the last command in the <i>conditions</i> list is zero, executes the <b>do</b> <i>actions</i> ; otherwise the loop terminates. If no commands in the <b>do</b> <i>actions</i> are executed, then the <b>while</b> command returns a zero exit status; <b>until</b> may be used in place of <b>while</b> to negate the loop termination test.                      |
| csch               | While <i>expr</i> is true (evaluates to nonzero), repeat commands between the <b>while</b> and the matching <b>end</b> statement. <b>break</b> and <b>continue</b> may be used to terminate or continue the loop prematurely. The <b>while</b> and <b>end</b> must appear alone on their input lines. If the shell's input is a terminal, it prompts for commands with a question-mark until the <b>end</b> command is entered and then performs the commands in the loop. |
| ksh                | A <b>while</b> command repeatedly executes the <b>while</b> <i>conditions</i> and, if the exit status of the last command in the <i>conditions</i> list is zero, executes the <b>do</b> <i>actions</i> ; otherwise the loop terminates. If no commands in the <b>do</b> <i>actions</i> are executed, then the <b>while</b> command returns a zero exit status; <b>until</b> may be used in place of <b>while</b> to negate the loop termination test.                      |
| <b>SEE ALSO</b>    | <b>break</b> (1), <b>csch</b> (1), <b>ksh</b> (1), <b>sh</b> (1)                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>NOTES</b>       | Both the Bourne shell, <b>sh</b> , and the Korn shell, <b>ksh</b> , can use the semicolon and the carriage return interchangeably in their syntax of the <b>if</b> , <b>for</b> , and <b>while</b> built-in commands.                                                                                                                                                                                                                                                      |

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>         | who – who is on the system                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>SYNOPSIS</b>     | <p><b>who</b> [ <b>-uTIHqpdbrtas</b> ] [ <i>filename</i> ]</p> <p><b>who -qn</b> <i>x</i> [ <i>filename</i> ]</p> <p><b>who am i</b></p> <p><b>who am I</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>AVAILABILITY</b> | SUNWcsu                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>DESCRIPTION</b>  | <p><b>who</b> can list the user's name, terminal line, login time, elapsed time since activity occurred on the line, and the process-ID of the command interpreter (shell) for each current UNIX system user. It examines the <b>/var/adm/utmp</b> file to obtain its information. If <i>filename</i> is given, that file (which must be in <b>utmp(4)</b> format) is examined. Usually, <i>filename</i> will be <b>/var/adm/wtmp</b>, which contains a history of all the logins since the file was last created.</p> <p><b>who</b> with the <b>am i</b> or <b>am I</b> option identifies the invoking user.</p> <p>The general format for output is:</p> <pre style="margin-left: 40px;">name [state] line time [idle] [pid] [comment] [exit]</pre> <p>The <i>name</i>, <i>line</i>, and <i>time</i> information is produced by all options except <b>-q</b>; the <i>state</i> information is produced only by <b>-T</b>; the <i>idle</i> and <i>pid</i> information is produced only by <b>-u</b> and <b>-l</b>; and the <i>comment</i> and <i>exit</i> information is produced only by <b>-a</b>. The information produced for <b>-p</b>, <b>-d</b>, and <b>-r</b> is explained during the discussion of each option, below.</p>                  |
| <b>OPTIONS</b>      | <p>With options, <b>who</b> can list logins, logoffs, reboots, and changes to the system clock, as well as other processes spawned by the <b>init</b> process. These options are:</p> <p><b>-u</b> This option lists only those users who are currently logged in. The <i>name</i> is the user's login name. The <i>line</i> is the name of the line as found in the directory <b>/dev</b>. The <i>time</i> is the time that the user logged in. The <i>idle</i> column contains the number of hours and minutes since activity last occurred on that particular line. A dot (.) indicates that the terminal has seen activity in the last minute and is therefore "current". If more than twenty-four hours have elapsed or the line has not been used since boot time, the entry is marked <b>old</b>. This field is useful when trying to determine whether a person is working at the terminal or not. The <i>pid</i> is the process-ID of the user's shell. The <i>comment</i> is the comment field associated with this line as found in <b>/sbin/inittab</b> (see <b>inittab(4)</b>). This can contain information about where the terminal is located, the telephone number of the dataset, type of terminal if hard-wired, and so forth.</p> |

- T** This option is the same as the **-s** option, except that the *state* of the terminal line is printed. The *state* describes whether someone else can write to that terminal. A ' + ' appears if the terminal is writable by anyone; a ' - ' appears if it is not. **root** can write to all lines having a ' + ' or a ' - ' in the *state* field. If a bad line is encountered, a ' ? ' is printed.
- l** This option lists only those lines on which the system is waiting for someone to login. The *name* field is LOGIN in such cases. Other fields are the same as for user entries except that the *state* field does not exist.
- H** This option will print column headings above the regular output.
- q** This is a quick **who**, displaying only the names and the number of users currently logged on. When this option is used, all other options are ignored.
- p** This option lists any other process which is currently active and has been previously spawned by **init**. The *name* field is the name of the program executed by **init** as found in **/sbin/inittab**. The *state*, *line*, and *idle* fields have no meaning. The *comment* field shows the *id* field of the line from **/sbin/inittab** that spawned this process. See **inittab(4)**.
- d** This option displays all processes that have expired and not been respawned by **init**. The *exit* field appears for dead processes and contains the termination and exit values (as returned by **wait(3B)**), of the dead process. This can be useful in determining why a process terminated.
- b** This option indicates the time and date of the last reboot.
- r** This option indicates the current *run-level* of the **init** process. In addition, it produces the process termination status, process id, and process exit status (see **utmp(4)**) under the *idle*, *pid*, and *comment* headings, respectively.
- t** This option indicates the last change to the system clock (via the **date** command) by **root**. See **su(1M)** and **date(1)**.
- a** This option processes **/var/adm/utmp** or the named *filename* with all options turned on.
- s** This option is the default and lists only the *name*, *line*, and *time* fields.
- nx** This option takes a numeric argument, *x*, which specifies the number of users to display per line. *x* must be at least 1. The **-n** option must be used with **-q**.

Note to the super-user: after a shutdown to the single-user state, **who** returns a prompt; the reason is that since **/var/adm/utmp** is updated at login time and there is no login in single-user state, **who** cannot report accurately on this state. **who am i**, however, returns the correct information.

**ENVIRONMENT**

If any of the **LC\_\*** variables (**LC\_CTYPE**, **LC\_MESSAGES**, **LC\_TIME**, **LC\_COLLATE**, **LC\_NUMERIC**, and **LC\_MONETARY**) (see **environ(5)**) are not set in the environment, the operational behavior of **who** for each corresponding locale category is determined by the value of the **LANG** environment variable. If **LC\_ALL** is set, its contents are used to override both the **LANG** and the other **LC\_\*** variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how **who** behaves.

**LC\_CTYPE**

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

**LC\_TIME**

Determines how **who** handles date and time formats. In the "C" locale, date and time handling follows the U.S. rules.

**FILES**

**/sbin/inittab**  
**/var/adm/utmp**  
**/var/adm/wtmp**

**SEE ALSO**

**date(1)**, **login(1)**, **mesg(1)**, **init(1M)**, **su(1M)**, **wait(3B)**, **inittab(4)**, **utmp(4)**, **environ(5)**

|                     |                                                                                                                                                                                                                                                                                                                                                                                       |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>         | whoami – display the effective current username                                                                                                                                                                                                                                                                                                                                       |
| <b>SYNOPSIS</b>     | <b>/usr/ucb/whoami</b>                                                                                                                                                                                                                                                                                                                                                                |
| <b>AVAILABILITY</b> | SUNWscpu                                                                                                                                                                                                                                                                                                                                                                              |
| <b>DESCRIPTION</b>  | <b>whoami</b> displays the login name corresponding to the current effective user ID. If you have used <b>su</b> to temporarily adopt another user, <b>whoami</b> will report the login name associated with that user ID. <b>whoami</b> gets its information from the <b>getuid</b> and <b>getpwuid</b> library routines (see <b>getuid</b> and <b>getpwnam(3C)</b> , respectively). |
| <b>FILES</b>        | <b>/etc/passwd</b> username data base                                                                                                                                                                                                                                                                                                                                                 |
| <b>SEE ALSO</b>     | <b>su(1M)</b> , <b>who(1)</b> , <b>getuid(2)</b> , <b>getpwnam(3C)</b>                                                                                                                                                                                                                                                                                                                |

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>         | whois – Internet user name directory service                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>SYNOPSIS</b>     | <b>whois</b> [ <b>-h</b> <i>host</i> ] <i>identifier</i>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                               |
| <b>AVAILABILITY</b> | SUNWcsu                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>DESCRIPTION</b>  | <p><b>whois</b> searches for an Internet directory entry for an <i>identifier</i> which is either a name (such as “Smith”) or a handle (such as “SRI-NIC”). To force a name-only search, precede the name with a period; to force a handle-only search, precede the handle with an exclamation point.</p> <p>To search for a group or organization entry, precede the argument with * (an asterisk). The entire membership list of the group will be displayed with the record.</p> <p>You may of course use an exclamation point and asterisk, or a period and asterisk together.</p> |
| <b>EXAMPLES</b>     | <p>The command:</p> <p style="padding-left: 40px;"><b>example% whois Smith</b></p> <p>looks for the name or handle SMITH.</p> <p>The command:</p> <p style="padding-left: 40px;"><b>example% whois !SRI-NIC</b></p> <p>looks for the handle SRI-NIC only.</p> <p>The command:</p> <p style="padding-left: 40px;"><b>example% whois .Smith, John</b></p> <p>looks for the name JOHN SMITH only.</p> <p>Adding ... to the name or handle argument will match anything from that point; that is, <b>ZU ...</b> will match ZUL, ZUM, and so on.</p>                                        |

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                         |
|---------------------|-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>         | write – write to another user                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                           |
| <b>SYNOPSIS</b>     | <b>write</b> <i>user</i> [ <i>line</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
| <b>AVAILABILITY</b> | SUNWcsu                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |
| <b>DESCRIPTION</b>  | <p><b>write</b> copies lines from your terminal to that of another user. When first called, it sends the message:</p> <p style="padding-left: 40px;"><b>Message from</b> <i>yourname</i> (<i>tty??</i>) [ <i>date</i> ]..</p> <p>to the person you want to talk to. When it has successfully completed the connection, it also sends two bells to your own terminal to indicate that what you are typing is being sent.</p> <p>The recipient of the message should write back at this point. Communication continues until an end of file is read from the terminal, an interrupt is sent, or the recipient has executed "mesg n". At that point <b>write</b> writes EOT on the other terminal and exits.</p> <p>If you want to write to a user who is logged in more than once, the <b>line</b> argument may be used to indicate which line or terminal to send to (for example, <b>term/12</b>); otherwise, the first writable instance of the user found in <b>/var/adm/utmp</b> is assumed and the following message posted:</p> <p style="padding-left: 40px;"><b>user is logged on more than one place.</b><br/> <b>You are connected to</b> "<i>terminal</i>".<br/> <b>Other locations are:</b><br/> <i>terminal</i></p> <p>Permission to write may be denied or granted by use of the <b>mesg</b> command. Writing to others is normally allowed by default. Certain commands, such as the <b>pr</b> command, disallow messages in order to prevent interference with their output. However, if the user has super-user privilege, messages can be forced onto a write-inhibited terminal.</p> <p>If the character <b>!</b> is found at the beginning of a line, <b>write</b> calls the shell to execute the rest of the line as a command.</p> <p><b>write</b> runs <b>setgid()</b> [see <b>setuid(2)</b>] to the group ID <b>tty</b>, in order to have write permissions on other user's terminals.</p> <p><b>write</b> will detect non-printable characters before sending them to the user's terminal. Control characters will appear as a '^' followed by the appropriate ASCII character; characters with the high-order bit set will appear in "meta" notation. For example, '\003' is displayed as '^C' and '\372' as '^M-z'.</p> <p>The following protocol is suggested for using <b>write</b>: when you first <b>write</b> to another user, wait for them to <b>write</b> back before starting to send. Each person should end a message with a distinctive signal (that is, <b>(o)</b> for "over") so that the other person knows when to reply. The signal <b>(oo)</b> (for "over and out") is suggested when conversation is to be terminated.</p> |

**ENVIRONMENT**

If any of the **LC\_\*** variables (**LC\_CTYPE**, **LC\_MESSAGES**, **LC\_TIME**, **LC\_COLLATE**, **LC\_NUMERIC**, and **LC\_MONETARY**) (see **environ(5)**) are not set in the environment, the operational behavior of **write** for each corresponding locale category is determined by the value of the **LANG** environment variable. If **LC\_ALL** is set, its contents are used to override both the **LANG** and the other **LC\_\*** variables. If none of the above variables is set in the environment, the "C" (U.S. style) locale determines how **write** behaves.

**LC\_CTYPE**

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

**LC\_MESSAGES**

Determines how diagnostic and informative messages are presented. This includes the language and style of the messages, and the correct form of affirmative and negative responses. In the "C" locale, the messages are presented in the default form found in the program itself (in most cases, U.S. English).

**LC\_TIME**

Determines how **write** handles date and time formats. In the "C" locale, date and time handling follows the U.S. rules.

**FILES**

**/var/adm/utmp** to find user  
**/usr/bin/sh** to execute !

**SEE ALSO**

**mail(1)**, **mesg(1)**, **pr(1)**, **sh(1)**, **who(1)**, **setuid(2)**, **environ(5)**

**DIAGNOSTICS**

**user is not logged on** The person you are trying to **write** to is not logged on.  
**Permission denied** The person you are trying to **write** to denies that permission (with **mesg**).  
**Warning: cannot respond, set mesg -y**  
Your terminal is set to **mesg n** and the recipient cannot respond to you.  
**Can no longer write to user**  
The recipient has denied permission (**mesg n**) after you had started writing.

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>         | xargs – construct argument list(s) and execute command                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>SYNOPSIS</b>     | <b>xargs</b> [ <i>flags</i> ] [ <i>command</i> [ <i>initial-arguments</i> ] ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>AVAILABILITY</b> | SUNWcsu                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |
| <b>DESCRIPTION</b>  | <p><b>xargs</b> combines the fixed <i>initial-arguments</i> with arguments read from standard input to execute the specified <i>command</i> one or more times. The number of arguments read for each <i>command</i> invocation and the manner in which they are combined are determined by the flags specified.</p> <p><i>command</i>, which may be a shell file, is searched for, using one's \$PATH. If <i>command</i> is omitted, <b>/usr/bin/echo</b> is used.</p> <p>Arguments read in from standard input are defined to be contiguous strings of characters delimited by one or more blanks, tabs, or new-lines; empty lines are always discarded. Blanks and tabs may be embedded as part of an argument if escaped or quoted. Characters enclosed in quotes (single or double) are taken literally, and the delimiting quotes are removed. Outside of quoted strings a backslash (\) escapes the next character.</p> <p>Each argument list is constructed starting with the <i>initial-arguments</i>, followed by some number of arguments read from standard input (Exception: see <b>-i</b> flag). Flags <b>-i</b>, <b>-l</b>, and <b>-n</b> determine how arguments are selected for each command invocation. When none of these flags are coded, the <i>initial-arguments</i> are followed by arguments read continuously from standard input until an internal buffer is full, and then <i>command</i> is executed with the accumulated args. This process is repeated until there are no more args. When there are flag conflicts (for example, <b>-l</b> vs. <b>-n</b>), the last flag has precedence.</p> |
| <b>OPTIONS</b>      | <p><b>-l</b><i>number</i>      <i>command</i> is executed for each non-empty <i>number</i> lines of arguments from standard input. The last invocation of <i>command</i> will be with fewer lines of arguments if fewer than <i>number</i> remain. A line is considered to end with the first new-line <i>unless</i> the last character of the line is a blank or a tab; a trailing blank/tab signals continuation through the next non-empty line. If <i>number</i> is omitted, 1 is assumed. Option <b>-x</b> is forced.</p> <p><b>-i</b><i>replstr</i>      Insert mode: <i>command</i> is executed for each line from standard input, taking the entire line as a single arg, inserting it in <i>initial-arguments</i> for each occurrence of <i>replstr</i>. A maximum of 5 arguments in <i>initial-arguments</i> may each contain one or more instances of <i>replstr</i>. Blanks and tabs at the beginning of each line are thrown away. Constructed arguments may not grow larger than 255 characters, and option <b>-x</b> is also forced. {} is assumed for <i>replstr</i> if not specified.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                 |

|                 |                                                                                                                                                                                                                                                                                                                                                                                                                                               |
|-----------------|-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>-nnumber</b> | Execute <i>command</i> using as many standard input arguments as possible, up to <i>number</i> arguments maximum. Fewer arguments are used if their total size is greater than <i>size</i> characters, and for the last invocation if there are fewer than <i>number</i> arguments remaining. If option <b>-x</b> is also coded, each <i>number</i> arguments must fit in the <i>size</i> limitation, else <b>xargs</b> terminates execution. |
| <b>-t</b>       | Trace mode: The <i>command</i> and each constructed argument list are echoed to file descriptor 2 just prior to their execution.                                                                                                                                                                                                                                                                                                              |
| <b>-p</b>       | Prompt mode: The user is asked whether to execute <i>command</i> each invocation. Trace mode ( <b>-t</b> ) is turned on to print the command instance to be executed, followed by a <i>?... prompt</i> . A reply of <i>y</i> (optionally followed by anything) executes the command; anything else, including just a carriage return, skips that particular invocation of <i>command</i> .                                                    |
| <b>-x</b>       | <b>xargs</b> will terminate if any argument list would be greater than <i>size</i> characters; <b>-x</b> is forced by the options <b>-i</b> and <b>-l</b> . When neither of the options <b>-i</b> , <b>-l</b> , or <b>-n</b> are coded, the total length of all arguments must be within the <i>size</i> limit.                                                                                                                               |
| <b>-ssize</b>   | The maximum total size of each argument list is set to <i>size</i> characters; <i>size</i> must be a positive integer less than or equal to 470. If <b>-s</b> is not coded, 470 is taken as the default. Note that the character count for <i>size</i> includes one extra character for each argument and the count of characters in the command name.                                                                                        |
| <b>-eofstr</b>  | <i>eofstr</i> is taken as the logical end-of-file string. Underbar ( <i>_</i> ) is assumed for the logical EOF string if <b>-e</b> is not coded. The value <b>-e</b> with no <i>eofstr</i> coded turns off the logical EOF string capability (underbar is taken literally). <b>xargs</b> reads standard input until either end-of-file or the logical EOF string is encountered.                                                              |

**xargs** terminates if either it receives a return code of **-1** from, or if it cannot execute, *command*. When *command* is a shell program, it should explicitly **exit** (see **sh(1)**) with an appropriate value to avoid accidentally returning with **-1**.

## EXAMPLES

The following examples moves all files from directory \$1 to directory \$2, and echo each move command just before doing it:

```
example% ls $1 | xargs -i -t mv $1/{ } $2/{ }
```

The following example combines the output of the parenthesized commands onto one line, which is then echoed to the end of file **log**:

```
example% (logname; date; echo $0 $*) | xargs >>log
```

The user is asked which files in the current directory are to be archived and archives them into **arch** (1.) one at a time, or (2.) many at a time.

1. **ls | xargs -p -l ar r arch**
2. **ls | xargs -p -l | xargs ar r arch**

The following example executes **diff(1)** with successive pairs of arguments originally typed as shell arguments:

```
example% echo $* | xargs -n2 diff
```

**SEE ALSO****sh(1)**

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>         | xgettext – extract gettext call strings from C programs                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>SYNOPSIS</b>     | <b>xgettext</b> [ <b>-ns</b> ] [ <b>-a</b> [ <b>-x</b> <i>exclude-file</i> ] ] [ <b>-c</b> <i>comment-tag</i> ] [ <b>-d</b> <i>default-domain</i> ] [ <b>-j</b> ]<br>[ <b>-m</b> <i>prefix</i> ] [ <b>-M</b> <i>suffix</i> ] [ <b>-p</b> <i>pathname</i> ] -   <i>filename</i> ...<br><b>xgettext -h</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>AVAILABILITY</b> | SUNWloc                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>DESCRIPTION</b>  | <p><b>xgettext</b> is used to automate the creation of portable message files (<b>.po</b>). A <b>.po</b> file contains copies of “C” strings that are found in ANSI C source code in <i>filename</i> or the standard input if ‘-’ is specified on the command line. The <b>.po</b> file can be used as input to the <b>msgfmt(1)</b> utility, which produces a binary form of the message file that can be used by application during run-time.</p> <p><b>xgettext</b> writes <i>msgid</i> strings from <b>gettext(3I)</b> calls in <i>filename</i> to the default output file <b>messages.po</b>. The default output file name can be changed by <b>-d</b> option. <i>msgid</i> strings in <b>dgettext()</b> calls are written to the output file <i>domainname.po</i> where <i>domainname</i> is the first parameter to the <b>dgettext()</b> call.</p> <p>By default, <b>xgettext</b> creates a <b>.po</b> file in the current working directory, and each entry is in the same order the strings are extracted from <i>filenames</i>. When the <b>-p</b> option is specified, the <b>.po</b> file is created in the <i>pathname</i> directory. An existing <b>.po</b> file is overwritten.</p> <p>Duplicate <i>msgids</i> are written to the <b>.po</b> file as comment lines. When the <b>-s</b> option is specified, the <b>.po</b> is sorted by the <i>msgid</i> string, and all duplicated <i>msgids</i> are removed. All <i>msgstr</i> directives in the <b>.po</b> file are empty unless the <b>-m</b> option is used.</p> |
| <b>OPTIONS</b>      | <p><b>-n</b> Add comment lines to the output file indicating file name and line number in the source file where each extracted string is encountered. These lines appear before each <i>msgid</i> in the following format:</p> <pre style="margin-left: 40px;">#<br/># File: filename, line: line-number</pre> <p><b>-s</b> Generate output sorted by <i>msgids</i> with all duplicate <i>msgids</i> removed.</p> <p><b>-a</b> Extract all strings, not just those found in <b>gettext(3I)</b>, and <b>dgettext()</b> calls. Only one <b>.po</b> file is created.</p> <p><b>-c</b> <i>comment-tag</i> The comment block beginning with <i>comment-tag</i> as the first token of the comment block is added to the output <b>.po</b> file as # delimited comments. For multiple domains, <b>xgettext</b> directs comments and messages to the prevailing text domain.</p> <p><b>-d</b> <i>default-domain</i> Rename default output file from <b>messages.po</b> to <i>default-domain .po</i>.</p> <p><b>-j</b> Join messages with existing message files. If a <b>.po</b> file does not exist, it is created. If a <b>.po</b> file does exist, new messages are appended. Any duplicate <b>msgids</b> are commented out in the resulting <b>.po</b> file. Domain directives in the existing <b>.po</b> file are ignored. Results not guaranteed if the existing message file has been edited.</p>                                                                                                                                     |

- m *prefix*** Fill in the *msgstr* with *prefix*. This is useful for debugging purposes. To make *msgstr* identical to *msgid*, use an empty string ("") for *prefix*.
- M *suffix*** Fill in the *msgstr* with *suffix*. This is useful for debugging purposes.
- p *pathname*** Specify the directory where the output files will be placed. This option overrides the current working directory.
- x *exclude-file*** Specify a **.po** file that contains a list of *msgid*s that are not to be extracted from the input files. The format of *exclude-file* is identical to the **.po** file. However, only the *msgid* directive line in *exclude-file* is used. All other lines are simply ignored. The **-x** option can only be used with the **-a** option.
- h** Print a help message on the standard output.

**SEE ALSO** [msgfmt\(1\)](#), [gettext\(3I\)](#)

**NOTES** **xgettext** is not able to extract cast strings, for example ANSI C casts of literal strings to (**const char \***). This is unnecessary anyway, since the prototypes in **<libintl.h>** already specify this type.

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>         | xstr – extract strings from C programs to implement shared strings                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>SYNOPSIS</b>     | <pre>xstr -c filename [ -v ] [ -I array ] xstr [ -I array ] xstr filename [ -v ] [ -I array ]</pre>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>AVAILABILITY</b> | SUNWcsu                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>DESCRIPTION</b>  | <p><b>xstr</b> maintains a file called <b>strings</b> into which strings in component parts of a large program are hashed. These strings are replaced with references to this common area. This serves to implement shared constant strings, which are most useful if they are also read-only.</p> <p>The command:</p> <pre>example% xstr -c filename</pre> <p>extracts the strings from the C source in <i>name</i>, replacing string references by expressions of the form <b>&amp;xstr[<i>number</i>]</b> for some <i>number</i>. An appropriate declaration of <b>xstr</b> is prepended to the file. The resulting C text is placed in the file <b>x.c</b>, to then be compiled. The strings from this file are placed in the <b>strings</b> data base if they are not there already. Repeated strings and strings which are suffixes of existing strings do not cause changes to the data base.</p> <p>After all components of a large program have been compiled, a file declaring the common <b>xstr</b> space called <b>xs.c</b> can be created by a command of the form:</p> <pre>example% xstr</pre> <p>This <b>xs.c</b> file should then be compiled and loaded with the rest of the program. If possible, the array can be made read-only (shared) saving space and swap overhead.</p> <p><b>xstr</b> can also be used on a single file. A command:</p> <pre>example% xstr filename</pre> <p>creates files <b>x.c</b> and <b>xs.c</b> as before, without using or affecting any <b>strings</b> file in the same directory.</p> <p>It may be useful to run <b>xstr</b> after the C preprocessor if any macro definitions yield strings or if there is conditional code which contains strings which may not, in fact, be needed. <b>xstr</b> reads from the standard input when the argument <b>'-'</b> is given. An appropriate command sequence for running <b>xstr</b> after the C preprocessor is:</p> <pre>example% cc -E name.c   xstr -c - example% cc -c x.c example% mv x.o name.o</pre> <p><b>xstr</b> does not touch the file <b>strings</b> unless new items are added; thus <b>make(1S)</b> can avoid remaking <b>xs.o</b> unless truly necessary.</p> |

|                 |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|-----------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>OPTIONS</b>  | <p><b>-c filename</b> Take C source text from <i>filename</i>.</p> <p><b>-v</b> Verbose: display a progress report indicating where new or duplicate strings were found.</p> <p><b>-I array</b> Specify the named <i>array</i> in program references to abstracted strings. The default array name is <b>xstr</b>.</p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                     |
| <b>FILES</b>    | <p><b>strings</b> data base of strings</p> <p><b>x.c</b> massaged C source</p> <p><b>xs.c</b> C source for definition of array "xstr*(rq</p> <p><b>/tmp/xs*</b> temp file when <b>xstr filename</b> doesn't touch <b>strings</b></p>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                       |
| <b>SEE ALSO</b> | <b>make(1S)</b>                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                            |
| <b>BUGS</b>     | If a string is a suffix of another string in the data base, but the shorter string is seen first by <b>xstr</b> both strings will be placed in the data base, when just placing the longer one there would do.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>NOTES</b>    | <p>Be aware that <b>xstr</b> indiscriminately replaces all strings with expressions of the form <b>&amp;xstr[number]</b> regardless of the way the original C code might have used the string. For example, you will encounter a problem with code that uses <b>sizeof()</b> to determine the length of a literal string because <b>xstr</b> will replace the literal string with a pointer that most likely will have a different size than the string's. To circumvent this problem:</p> <ul style="list-style-type: none"> <li>• use <b>strlen()</b> instead of <b>sizeof()</b>; note that <b>sizeof()</b> returns the size of the array (including the null byte at the end), whereas <b>strlen()</b> doesn't count the null byte. The equivalent of <b>sizeof("xxx")</b> really is <b>(strlen("xxx")+1)</b>.</li> <li>• use <b>#define</b> for operands of <b>sizeof()</b> and use the <b>define</b>'d version. <b>xstr</b> ignores <b>#define</b> statements. Make sure you run <b>xstr</b> on <i>filename</i> before you run it on the preprocessor.</li> </ul> <p>You will also encounter a problem when declaring an initialized character array of the form</p> <pre>char x[] = "xxx";</pre> <p><b>xstr</b> will replace <i>xxx</i> with an expression of the form <b>&amp;xstr[number]</b> which will not compile. To circumvent this problem, use <b>static char *x = "xxx"</b> instead of <b>static char x[] = "xxx"</b>.</p> |

**NAME** yacc – yet another compiler-compiler

**SYNOPSIS** yacc [ -vVdlt ] [ -Q [ y | n ] ] [ -p *driver-file* ] *filename*

**DESCRIPTION** The **yacc** command converts a context-free grammar into a set of tables for a simple automaton that executes an LALR(1) parsing algorithm. The grammar may be ambiguous; specified precedence rules are used to break ambiguities.

The output file, **y.tab.c**, must be compiled by the C compiler to produce a function **yyparse()**. This program must be loaded with the lexical analyzer program, **yylex()**, as well as **main()** and **yyerror()**, an error handling routine. These routines must be supplied by the user; the **lex(1)** command is useful for creating lexical analyzers usable by **yacc**.

**OPTIONS**

- d** Generates the file **y.tab.h** with the **#define** statements that associate the **yacc** user-assigned “token codes” with the user-declared “token names.” This association allows source files other than **y.tab.c** to access the token codes.
- l** Specifies that the code produced in **y.tab.c** will not contain any **#line** constructs. This option should only be used after the grammar and the associated actions are fully debugged.
- t** Compiles runtime debugging code by default. Runtime debugging code is always generated in **y.tab.c** under conditional compilation control. By default, this code is not included when **y.tab.c** is compiled. Whether or not the **-t** option is used, the runtime debugging code is under the control of **YYDEBUG**, a preprocessor symbol. If **YYDEBUG** has a non-zero value, then the debugging code is included. If its value is zero, then the code will not be included. The size and execution time of a program produced without the runtime debugging code will be smaller and slightly faster.
- v** Prepares the file **y.output**, which contains a description of the parsing tables and a report on conflicts generated by ambiguities in the grammar.
- V** Prints on the standard error output the version information for **yacc**.
- Q[y | n]** The **-Qy** option puts the version stamping information in **y.tab.c**. This allows you to know what version of **yacc** built the file. The **-Qn** option (the default) writes no version information.
- p driver-file** Allows you to specify the parser of your choice instead of **/usr/ccs/bin/yaccpar**. For example, you can specify:  

```
yacc -p ~/myparser parser.y
```

**ENVIRONMENT** The environment variables **LC\_CTYPE** and **LANG** control the character classification throughout **yacc**. On entry to **yacc**, these environment variables are checked in the following order: **LC\_CTYPE** and **LANG**. When a valid value is found, remaining environment variables for character classification are ignored. For example, a new setting for **LANG** does not override the current valid character classification rules of **LC\_CTYPE**.

When none of the values are valid, the shell character classification defaults to the POSIX.1 C locale.

**yacc** can handle characters from EUC primary and supplementary codesets as one-token symbols. EUC codes may only be single character quoted terminal symbols. **yacc** expects **yylex()** to return a wide character (**wchar\_t**) value for these one-token symbols.

**FILES**

**y.output**  
**y.tab.c**  
**y.tab.h** defines for token names  
**yacc.acts** temporary file  
**yacc.debug** temporary file  
**yacc.tmp** temporary file  
**yaccpar** parser prototype for C programs

**SEE ALSO**

**lex(1)**  
The **yacc** chapter in the *Programming Utilities Guide* manual.

**DIAGNOSTICS**

The number of reduce-reduce and shift-reduce conflicts is reported on the standard error output; a more detailed report is found in the **y.output** file. Similarly, if some rules are not reachable from the start symbol, this instance is also reported.

**NOTES**

Because file names are fixed, at most one **yacc** process can be active in a given directory at a given time.

|                     |                                                                                                                                                                                                                                                                                                                           |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>         | <b>ypcat</b> – print values in a NIS database                                                                                                                                                                                                                                                                             |
| <b>SYNOPSIS</b>     | <b>ypcat</b> [ <b>-k</b> ] [ <b>-d</b> <i>ypdomain</i> ] <i>mname</i>                                                                                                                                                                                                                                                     |
| <b>AVAILABILITY</b> | SUNWnisu                                                                                                                                                                                                                                                                                                                  |
| <b>DESCRIPTION</b>  | <p>The <b>ypcat</b> command prints out values in the NIS name service map specified by <i>mname</i>, which may be either a map name or a map nickname. Since <b>ypcat</b> uses the NIS network services, no NIS server is specified.</p> <p>Refer to <b>ypfiles(4)</b> for an overview of the NIS name service.</p>       |
| <b>OPTIONS</b>      | <p><b>-k</b>                Display the keys for those maps in which the values are null or the key is not part of the value. None of the maps derived from files that have an ASCII version in <b>/etc</b> fall into this class.</p> <p><b>-d</b> <i>ypdomain</i>    Specify a domain other than the default domain.</p> |
| <b>SEE ALSO</b>     | <b>ypmatch(1)</b> , <b>ypfiles(4)</b>                                                                                                                                                                                                                                                                                     |

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                 |
|---------------------|---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>         | ypmatch – print the value of one or more keys from a NIS map                                                                                                                                                                                                                                                                                                                                                                    |
| <b>SYNOPSIS</b>     | <b>ypmatch</b> [ <b>-k</b> ] [ <b>-t</b> ] [ <b>-d domain</b> ] <i>key</i> [ <i>key...</i> ] <i>mname</i><br><b>ypmatch -x</b>                                                                                                                                                                                                                                                                                                  |
| <b>AVAILABILITY</b> | SUNWnisu                                                                                                                                                                                                                                                                                                                                                                                                                        |
| <b>DESCRIPTION</b>  | <b>ypmatch</b> prints the values associated with one or more keys from the NIS's name services map specified by <i>mname</i> , which may be either a <i>mapname</i> or a map nickname ( <i>mnames</i> ). Multiple keys can be specified; all keys will be searched for in the same map. The keys must be the same case and length. No pattern matching is available. If a key is not matched, a diagnostic message is produced. |
| <b>OPTIONS</b>      | <b>-k</b> Before printing the value of a key, print the key itself, followed by a ':' (colon).<br><b>-t</b> This option inhibits map nickname translation.<br><b>-d domain</b> Specify a domain other than the default domain.<br><b>-x</b> Display the map nickname table. This lists the nicknames ( <i>mnames</i> ) the command knows of, and indicates the <i>mapname</i> associated with each nickname.                    |
| <b>SEE ALSO</b>     | <b>ypcat(1)</b> , <b>ypfiles(4)</b>                                                                                                                                                                                                                                                                                                                                                                                             |

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                          |
|---------------------|----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>         | yppasswd – change your network password in the NIS database                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                              |
| <b>SYNOPSIS</b>     | <b>yppasswd</b> [ <i>username</i> ]                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                      |
| <b>AVAILABILITY</b> | SUNWcsu                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>DESCRIPTION</b>  | <p><b>yppasswd</b> changes the network password associated with the user <i>username</i> in the Network Information Service (NIS) database. If the user has done a <b>keylogin</b>(1), and a publickey/secretkey pair exists for the user in the NIS <b>publickey.byname</b> map, <b>yppasswd</b> also re-encrypts the secretkey with the new password. The NIS password may be different from the local one on your own machine. Use <b>passwd</b>(1) to change the password information on the local machine, and <b>nispasswd</b>(1) to change the password information stored in Network Information Service, Version 3 database (NIS+).</p> <p><b>yppasswd</b> prompts for the old NIS password, and then for the new one. You must type in the old password correctly for the change to take effect. The new password must be typed twice, to forestall mistakes.</p> <p>New passwords must be at least four characters long, if they use a sufficiently rich alphabet, and at least six characters long if monospace. These rules are relaxed if you are insistent enough. Only the owner of the name or the super-user may change a password; in either case you must prove you know the old password.</p> <p>The NIS password daemon, <b>rpc.yppasswdd</b> must be running on your NIS server in order for the new password to take effect.</p> |
| <b>WARNINGS</b>     | Even after the user has successfully changed his or her password using this command, the subsequent <b>login</b> (1) using the new password will be successful only if the user's password and shadow information is obtained from NIS, (see <b>getpwnam</b> (3C), <b>getspnam</b> (3C), and <b>nsswitch.conf</b> (4)).                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                  |
| <b>SEE ALSO</b>     | <b>keylogin</b> (1), <b>login</b> (1), <b>passwd</b> (1), <b>nispasswd</b> (1), <b>getpwnam</b> (3C), <b>getspnam</b> (3C), <b>secure_rpc</b> (3N), <b>nsswitch.conf</b> (4)                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                             |
| <b>BUGS</b>         | The update protocol passes all the information to the server in one RPC call, without ever looking at it. Thus if you type in your old password incorrectly, you will not be notified until after you have entered your new password.                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                                    |

|                     |                                                                                                                                                                                                                                                                                                                                                                                                                                            |
|---------------------|--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------|
| <b>NAME</b>         | ypwhich – return name of NIS server or map master                                                                                                                                                                                                                                                                                                                                                                                          |
| <b>SYNOPSIS</b>     | <b>ypwhich</b> [ <b>-d</b> <i>domain</i> ] [[ <b>-t</b> ] <b>-m</b> [ <i>mname</i> ] <i>hostname</i> ]<br><b>ypwhich</b> <b>-x</b>                                                                                                                                                                                                                                                                                                         |
| <b>AVAILABILITY</b> | SUNWnisu                                                                                                                                                                                                                                                                                                                                                                                                                                   |
| <b>DESCRIPTION</b>  | <b>ypwhich</b> returns the name of the NIS server that supplies the NIS name services to a NIS client, or which is the master for a map. If invoked without arguments, it gives the NIS server for the local machine. If <i>hostname</i> is specified, that machine is queried to find out which NIS master it is using.<br><br>Refer to <b>ypfiles(4)</b> for an overview of the NIS name services.                                       |
| <b>OPTIONS</b>      | <b>-d</b> <i>domain</i> Use <i>domain</i> instead of the default domain.<br><b>-t</b> This option inhibits map nickname translation.<br><b>-m</b> <i>mname</i> Find the master NIS server for a map. No <i>hostname</i> can be specified with <b>-m</b> . <i>mname</i> can be a mapname, or a nickname for a map. When <i>mname</i> is omitted, produce a list of available maps.<br><b>-x</b> Display the map nickname translation table. |
| <b>SEE ALSO</b>     | <b>ypfiles(4)</b>                                                                                                                                                                                                                                                                                                                                                                                                                          |

# Index

---

## *Special Characters*

- .mo files
  - message object files — msgfmt, 1-532
- .po files
  - portable object files — msgfmt, 1-532

## **A**

- accounting
  - search and print files — acctcom, 1-25
- acctcom — search and print process accounting files, 1-25
- adb — debugger, 1-27
  - \$ Modifier, 1-30
  - : Modifier, 1-30
  - ? and / Modifiers, 1-29
  - ?, /, and = Modifiers, 1-29
  - Binary Operators, 1-28
  - Commands, 1-28
  - Expressions, 1-27
  - Unary Operators, 1-28
  - Variables, 1-28
  - Verbs, 1-28
- addbib — create or extend bibliography, 1-33
- adds /dev entries to give SunOS 4.1 compatible names to SunOS 5.x devices — ucblinks, 1B-909
- alias — shell built-in functions to create your own pseudonym or shorthand for a command or

- series of commands, 1-35
- apropos — locate commands by keyword, 1-36
- ar — maintain portable archive or library, 1-37
- arch — display architecture of current host, 1B-39
- archive
  - maintain a portable one across all machines — ar, 1-37
- archives
  - create tape archives, and add or extract files — tar, 1-849
- argument list(s) and execute command, construct — xargs, 1-972
- as — assembler, 1-40
- assembler
  - as, 1-40
- at — execute commands at a later time, 1-42
- atq — display the jobs queued to run at specified times, 1-46
- atrm — remove jobs spooled by at or batch, 1-47
- audio file formats
  - convert — audioconvert, 1-48
- audio files
  - play — audioplay, 1-51
  - record — audiorecord, 1-53
- audioconvert — convert audio file formats, 1-48
- audioplay — play audio files, 1-51

---

audiorecord — record an audio file, 1-53  
authentication and authorization for network environment  
— kerberos, 1-356  
awk — pattern scanning and processing language, 1-56

**B**

banner — make posters, 1-60  
basename — display portions of pathnames, 1B-62, 1-61  
batch — execute commands at a later time, 1-42  
bc — arbitrary precision arithmetic language, 1-63  
bdiff — display line-by-line differences between pairs of large text files, 1-66  
bg — shell built-in functions to control process execution, 1-348  
bibliography  
create an inverted index to a bibliographic database — indexbib, 1-342  
create or extend — addbib, 1-33  
expand and insert references from a bibliographic database — refer, 1-699  
find references in a bibliographic database — lookbib, 1-433  
format and print a bibliographic database — roffbib, 1-713  
sort a bibliographic database — sortbib, 1-808  
biff — mail notifier, 1B-67  
binary file transmission  
decode binary file — uudecode, 1C-924  
encode binary file — uuencode, 1C-924  
binary files  
find printable strings — strings, 1-819  
locate — whereis, 1B-961  
block count  
for a file — sum, 1-832  
blocks, count a in file — sum, 1B-833  
Bourne shell  
— sh, 1-780  
Bourne shell variables  
— CDPATH, 1-783

Bourne shell variables, *continued*  
— HOME, 1-783  
— IFS, 1-783  
— LC\_ALL, 1-784  
— LC\_CTYPE, 1-784  
— LC\_MESSAGES, 1-784  
— MAIL, 1-783  
— MAILCHECK, 1-783  
— MAILPATH, 1-783  
— PATH, 1-783  
— PS1, 1-783  
— PS2, 1-783  
— SHACCT, 1-784  
— SHELL, 1-784  
break — shell built-in functions to escape from or advance within a controlling while, for, foreach, or until loop, 1-68  
build programs — make

**C**

C  
create a tags file for use with ex and vi — ctags, 1-170  
C compiler, 1B-76  
C language  
C preprocessor — cpp, 1-135  
C language program  
resolve and remove ifdefed lines from C program source — unifdef, 1-914  
C programming language  
create C error messages — mkstr, 1B-526  
extract strings from C code — xstr, 1-977  
formats program in nice style using troff — vgrind, 1-939  
C shell  
aliases — csh, 1-148  
built-in commands — csh  
command and filename substitution — csh, 1-150  
command execution — csh, 1-152  
command line parsing — csh, 1-146  
command substitution — csh, 1-150  
control flow — csh, 1-152  
environment variables and shell variables —

---

    csh,  
C shell, *continued*  
    1-162  
    event designators — csh, 1-146  
    expressions and operators — csh, 1-151  
    filename completion — csh, 1-145  
    filename substitution — csh, 1-151  
    history substitution — csh, 1-146  
    I/O redirection — csh, 1-148  
    initialization and termination — csh, 1-144  
    interactive operation — csh, 1-144  
    job control — csh, 1-153  
    lexical structure — csh, 1-145  
    modifiers — csh, 1-147  
    noninteractive operation — csh, 1-144  
    quick substitution — csh, 1-148  
    signal handling — csh, 1-153  
    status reporting — csh, 1-154  
    variable substitution — csh, 1-149  
    word designators — csh, 1-147

C shell commands  
— %, 1-161  
: — csh, 1-154  
— @, 1-161  
— alias, 1-154  
— bg, 1-154  
— break, 1-154  
— breaksw, 1-154  
— case, 1-154  
— cd, 1-154  
— chdir, 1-154  
— continue, 1-154  
— default, 1-154  
— dirs, 1-155  
— echo, 1-155  
— else, 1-156  
— end, 1-155  
— endif, 1-156  
— eval, 1-155  
— exec, 1-155  
— exit, 1-155  
— fg, 1-155  
— foreach, 1-155  
— glob, 1-155  
— goto, 1-155

### C shell commands, *continued*

— hashstat, 1-155  
— history, 1-156  
— if, 1-156  
— jobs, 1-156  
— kill, 1-156  
— limit, 1-156  
— login, 1-157  
— logout, 1-157  
— nice, 1-157  
— nohup, 1-157  
— notify, 1-157  
— onintr, 1-158  
— popd, 1-158  
— pushd, 1-158  
— rehash, 1-158  
— repeat, 1-158  
— set, 1-158  
— setenv, 1-158  
— shift, 1-159  
— source, 1-160  
— stop, 1-160  
— suspend, 1-160  
— switch, 1-160  
— time, 1-160  
— umask, 1-160  
— unalias, 1-160  
— unhash, 1-161  
— unlimit, 1-161  
— unset, 1-161  
— unsetenv, 1-161  
— wait, 1-161  
— while, 1-161

cal — display a calendar, 1-69

calculator, desk

— dc, 1-183

calendar — reminder service, 1-70

display — cal, 1-69

call-graph, display profile data — gprof, 1-320

cancel — cancel requests to an LP print service, 1-435

cancel user's request for removable media that is not currently in drive — volcancel, 1-950

case — shell built-in functions to choose from among a list of actions, 1-72

---

**cat** — concatenate and display files, 1-74  
**cc** — C compiler, 1B-76  
**cd** — shell built-in functions to change the current working directory, 1-77  
**CDPATH** variable — **sh**, 1-783  
**character translation** — **tr**, 1B-890, 1-888  
**chdir** — shell built-in functions to change the current working directory, 1-77  
**check for media in a drive. Default checks for floppy.** — **volcheck**, 1-951  
**check spelling** — **spell**, 1-810  
**checkeq** — check eqn constructs, 1-234  
**checknr** — check nroff/troff files, 1-79  
**chgrp** — change the group ownership of a file, 1-80  
**chmod** — change the permissions mode of a file, 1-83  
**chown** — change owner of file, 1-87  
**chown** — change owner of file, 1B-89  
**clear** — clear terminal screen, 1-115  
**cmp** — compare two files, 1-116  
**cocheck** — (FMLI utility) communicate with a process, 1F-124  
**cocreate** — (FMLI utility) communicate with a process, 1F-124  
**code formatter**  
     formats program in nice style using troff — **vgrind**, 1-939  
**code set**  
     conversion utility — **iconv**, 1-335  
**codestroy** — (FMLI utility) communicate with a process, 1F-124  
**cof2elf** — COFF to ELF object file translation, 1-118  
**COFF to ELF object file translation** — **cof2elf**, 1-118  
**col** — filters reverse line-feeds from two-column nroff text, 1-119  
**comm** — select or reject lines common to two sorted files, 1-121  
**command**  
     describe — **whatis**, 1-960  
**command options**  
     command options, *continued*  
         parse — **getopt**, 1-310, 1-312  
**commands**  
     change priority of — **nice**, 1-551  
     display the last commands executed, in reverse order — **lastcomm**, 1-399  
     locate a command; display its pathname or alias — **which**, 1-963  
     locate by keyword — **apropos**, 1-36  
**communications**  
     connect to remote system — **cu**, 1C-172, 1-875  
     decode binary files — **uudecode**, 1C-924  
     encode binary files — **uuencode**, 1C-924  
     public UNIX-to-UNIX copy — **uupick**, 1C-929  
     system to system command execution — **uux**, 1C-931  
     talk to another user — **talk**, 1-848  
     UNIX-to-UNIX copy — **uucp**, 1C-921  
     user interface to a remote system using the TELNET protocol — **telnet**, 1-860  
     UUCP list of names — **uuname**, 1C-921  
     UUCP log — **uulog**, 1C-921  
     write to another user — **write**, 1-970  
**compiler generators**  
     **lex** — lexical analyzer generator, 1-408  
**compilers**  
     C compiler — **cc**, 1B-76  
     regular expression compile — **regcmp**, 1-701  
     RPC protocol compiler — **rpcgen**, 1-715  
**compress** — compress files, 1-122  
**concatenate**  
     files and display them — **cat**, 1-74  
**connect to remote system**  
     — **cu**, 1C-172  
**continue** — shell built-in functions to escape from or advance within a controlling while, for, foreach, or until loop, 1-68  
**control line printer** — **lpc**  
**convert units** — **units**, 1-917  
**coproc** — (FMLI utility) communicate with a process, 1F-124  
**copy**  
     archives — **cpio**, 1-130  
     files — **cp**, 1-128

core image  
   of running processes — `gcore`, 1-305  
`coreceive` — (FMLI utility) communicate with a process, 1F-124  
`cosend` — (FMLI utility) communicate with a process, 1F-124  
 count blocks in file — `sum`, 1B-833  
 count lines, words, characters in file — `wc`, 1-958  
`cp` — copy files, 1-128  
`cpio` — copy archives, 1-130  
`cpp` — C preprocessor, 1-135  
 create  
   bibliography — `addbib`, 1-33  
`crontab` — user crontab file, 1-141  
`crypt` — encrypt, 1-143  
`csh` — shell command interpreter with a C-like syntax  
`csplit` — split a file with respect to a given context, 1-166  
`ct` — spawn login to a remote terminal, 1C-168  
`ctags` — create a tags file for use with `ex` and `vi`, 1-170  
`cu` — connect to remote system, 1C-172  
 curve, smooth  
   interpolate — `spline`, 1-811  
`cut` — cut out selected fields of each line of a file, 1-177

## D

`date` — display date and/or set date, 1-180  
   prompts for a date — `ckdat`, 1-90  
   provides error message for date — `errdate`, 1-90  
   provides help message for date — `helpdate`, 1-90  
   validates a date — `valdate`, 1-90  
`dead.letter`  
   mail interrupted during input — `mail`, 1-469  
 debug tools  
   debugger — `adb`, 1-27  
 decode binary file — `uudecode`, 1C-924  
 decode files  
   — `crypt`, 1-143

`decrypt` — `crypt`, 1-143  
 dependencies, dynamic  
   of executable files or shared objects — `ldd`, 1-406  
`deroff` — remove `nroff`, `troff`, `tbl` and `eqn` constructs, 1-185  
`describe` command — `whatis`, 1-960  
 desk calculator  
   — `dc`, 1-183  
 devices  
   eject media device from drive — `eject`, 1-230  
`df` — display status of disk space on file systems, 1B-186  
 dictionary, system  
   find words — `look`, 1-432  
`diff` — display line-by-line differences between pairs of text files, 1-187  
   3-way — `diff3`, 1-190  
   big — `bdiff`, 1-66  
`diff` command  
   side-by-side — `sdiff`, 1-767  
`diff3` — display line-by-line differences between three text files, 1-190  
`diffmk` — mark differences between versions of a `troff` input file, 1-192  
`dircmp` — compares contents of directories, 1-193  
 directories  
   compare contents — `dircmp`, 1-193  
   list contents — `ls`, 1-454  
   list contents of — `ls`, 1B-459  
   make — `mkdir`, 1-522  
   make link to — `ln`, 1B-418  
   print working directory name — `pwd`, 1-688  
   remove — `rmdir`, 1-711  
`dirname` — delivers all but last level of pathname, 1-61  
`dirs` — shell built-in functions to change the current working directory, 1-77  
`dis` — object code disassembler, 1-194  
`disable` — disable LP printers, 1-232  
 disassembler  
   object code — `dis`, 1-194  
 display

---

display, *continued*

- a list of all valid user names — `dispuid`, 1-197
- architecture of current host — `arch`, 1B-39
- call-graph profile data — `gprof`, 1-320
- contents of directory — `ls`, 1-454
- current news — `news`, 1-549
  - `date`, 1-180
- disk usage — `du`, 1B-204
- dynamic dependencies of executable files or shared objects — `ldd`, 1-406
- effective user name — `whoami`, 1B-968
- file names — `ls`, 1B-459
- first few lines of files — `head`, 1-331
- group membership of user — `groups`, 1B-328, 1-327
- how long the system has been up — `uptime`, 1-919
- identifier of current host — `hostid`, 1B-333
- last commands executed, in reverse order — `lastcomm`, 1-399
- list of all valid group names — `disgid`, 1-196
- login and logout information about users and terminals — `last`, 1-398
- name of current host — `hostname`, 1B-334
- name of the user running the process — `logname`, 1-430
- printer queue — `lpq`, 1B-444
- process status — `ps`, 1B-685
- processor type of current host — `mach`, 1B-466
- size of page of memory — `pagesize`, 1B-614
- status of disk space on file system — `df`, 1B-186
- status of local hosts — `ruptime`, 1-725
- status of network hosts — `rup`, 1-723
- users on system — `users`, 1B-920
- working directory name — `pwd`, 1-688

display editor — `vi`, 1-942

display or change font information in the RAM of the video card on an x86 system in text mode — `loadfont`, 1-421

displays package parameter values — `pkgparam`, 1-638

display profile data — `prof`, 1-678

document production

document production, *continued*

- check `nroff`/`troff` files — `checknr`, 1-79
- check spelling — `spell`, 1-810
- create an inverted index to a bibliographic database — `indexbib`, 1-342
- create or extend bibliography — `addbib`, 1-33
- eliminate `.so`'s from `nroff` input — `soelim`, 1-803
- expand and insert references from a bibliographic database — `refer`, 1-699
- filters reverse line-feeds from two-column `nroff` text — `col`, 1-119
- find references in a bibliographic database — `lookbib`, 1-433
- format and print a bibliographic database — `roffbib`, 1-713
- format documents for display or line-printer — `nroff`, 1-605
- format tables for `nroff` or `troff` — `tbl`, 1-856
- mark differences between versions of a `troff` input file — `difffmk`, 1-192
- remove `nroff`, `troff`, `tbl` and `eqn` constructs — `deroff`, 1-185
- simple text formatters — `fmt`, 1-286
- sort a bibliographic database — `sortbib`, 1-808
- `troff` postprocessor for PostScript printers — `dpost`, 1-201
- typeset mathematics — `eqn`, 1-234
- typeset or format documents — `troff`, 1-893

DOS

- convert text file from DOS format to ISO format — `dos2unix`, 1-198
- convert text file from ISO format to DOS format — `unix2dos`, 1-918

`dos2unix` — convert text file from DOS format to ISO format, 1-198

download — host resident PostScript font downloader, 1-199

`dpost` — `troff` postprocessor for PostScript printers, 1-201

draw graph — `graph`, 1-323

`du` — display disk usage per directory or file, 1B-204

---

`dump` — dump selected parts of an object file, 1-206  
`dumpcs` — show codeset table for the current locale, 1-208  
`dumpkeys` — dump keyboard translation tables, 1-423

## E

`echo` — (FMLI utility) put string on virtual output, 1F-212, 1-209  
`echo` — echo arguments to standard output, 1B-211  
`ed` — text editor, 1-213  
`edit` — text editor, 1-224  
`egrep` — search a file for a pattern using full regular expressions, 1-228  
`eject` — eject media device from drive, 1-230  
`enable` — enable LP printers, 1-232  
`encode binary file` — `uuencode`, 1C-924  
`encode files`  
— `crypt`, 1-143  
`encryption key, user`  
change — `chkey`, 1-82  
`env` — obtain or alter environment variables for command execution, 1-233  
`environment`  
display variables — `printenv`, 1B-668  
set terminal characteristics — `tset`, 1B-900  
`environment variables`  
set or alter for command execution — `env`, 1-233  
`environment variables, global`  
FMLI, 1F-778  
`eqn`  
remove `nroff`, `troff`, `tbl` and `eqn` constructs — `deroff`, 1-185  
`eqn` — mathematical typesetting, 1-234  
`equations`  
typeset mathematics — `eqn`, 1-234  
`error` — analyze error messages, 1-237  
`eval` — shell built-in functions to execute other commands, 1-246  
`ex` — text editor, 1-240  
`exec` — shell built-in functions to execute other

commands, 1-246  
`exit` — shell built-in functions to enable the execution of the shell to advance beyond its sequence of steps, 1-248  
`expand` — expand TAB characters to SPACE characters, 1-250  
`export` — shell built-in functions to determine the characteristics for environmental variables of the current shell and its descendents, 1-773  
`exportfs` — translates `exportfs` options to `share/unshare` commands, 1B-251  
`expr` — evaluate arguments as an expression, 1-252, 1B-255  
`expression evaluation` — `expr`, 1B-255  
`extr` — extract strings from source files, 1-258  
`extract strings from C code` — `xstr`, 1-977

## F

`face` — executable for the Framed Access Command Environment Interface, 1-261  
`factor` — obtain the prime factors of a number, 1-262  
`false` — provide truth values, 1-895  
`fastboot` — reboot system without checking disks, 1B-263  
`fasthalt` — halt system without checking disks, 1B-263  
`fc` — shell built-in functions to re-use previous command-lines from the current shell, 1-332  
`fdformat` — floppy diskette format  
format floppy diskette, 1-264  
`fg` — shell built-in functions to control process execution, 1-348  
`fgrep` — search a file for a character string, 1-267  
`file`  
change ownership — `chown`, 1B-89  
determine type of — `file`, 1B-271  
display names — `ls`, 1B-459  
files perusal filter for CRTs — `pg`, 1-630  
make link to — `ln`, 1B-418  
print — `lpr`, 1B-446  
strip affixes — `basename`, 1B-62

---

file, *continued*

sum — sum and count blocks in file, 1B-833

file — get file type, 1B-271

file system

- display status of disk space — `df`, 1B-186
- make hard or symbolic links to files — `ln`, 1-416
- where am I — `pwd`, 1-688

file transfer program

- `ftp`, 1-295

file — determine file type, 1-269

- change owner of file — `chown`, 1-87
- change the permissions mode of a file — `chmod`, 1-83
- compare two files — `cmp`, 1-116
- compress — `compress`, 1-122
- compress files — `pack`, 1-611
- concatenate and display — `cat`, 1-74
- copy — `cp`, 1-128
- copy archives — `cpio`, 1-130
- `crypt` — encrypt/decrypt, 1-143
- cut out selected fields of each line of a file — `cut`, 1-177
- display a count of lines, words and characters in a file — `wc`, 1-958
- display first few lines — `head`, 1-331
- display last part — `tail`, 1-846
- display line-by-line differences between pairs of large text files — `bdiff`, 1-66
- display line-by-line differences between pairs of text files — `diff`, 1-187
- display line-by-line differences between three text files — `diff3`, 1-190
- display uncompressed files but leaves compressed files intact — `zcat`, 1-122
- expand compressed files — `unpack`, 1-611
- extract SCCS version information from a file — `what`, 1-959
- `find`, 1-273
- mark differences between versions of a troff input file — `diffmk`, 1-192
- merge same lines of several files or subsequent lines of one file — `paste`, 1-619
- move — `mv`, 1-537
- print checksum and block count for a file —

sum,

file — determine file type, *continued* 1-832

- print differences between two files side-by-side — `sdiff`, 1-767
- remove — `rm`, 1-711
- search a file for a character string — `fgrep`, 1-267
- search a file for a pattern — `grep`, 1-325
- search for a pattern using full regular expressions — `egrep`, 1-228
- select or reject lines common to two sorted files — `comm`, 1-121
- sort or merge — `sort`, 1-804
- split a file into pieces — `split`, 1-812
- split a file with respect to a given context — `csplit`, 1-166
- strip affixes from pathnames — `basename`, 1-61
- transfer to and from a remote machine — `tftp`, 1-869
- uncompress — `uncompress`, 1-122
- update access time and/or modification time of a file — `touch`, 1-883

find — find files, 1-273

floppy diskette format — `fdformat`, 1-264

fmlcut — (FMLI utility) cut out columns from a table or fields from each line of a file, 1F-279

fmlexpr — (FMLI utility) evaluate arguments as an expression, 1F-281

fmlgrep — (FMLI utility) search afile for a pattern, 1F-283

FMLI

- `cocheck` — communicate with a process, 1F-124
- `cocreate` — communicate with a process, 1F-124
- `codestroy` — communicate with a process, 1F-124
- `coproc` — communicate with a process, 1F-124
- `coreceive` — communicate with a process, 1F-124
- `cosend` — communicate with a process, 1F-124
- `echo` — put string on virtual output, 1F-212
- `fmlcut` — cut out columns from a table or

- 
- fields
  - FMLI, *continued*
  - from each line of a file, 1F-279
  - fmlexpr — evaluate arguments as an expression, 1F-281
  - fmlgrep — search afile for a pattern, 1F-283
  - fmlfi — invoke fmlfi, 1-284
  - getfrm — returns the current frameID number, 1F-308
  - getitems — returns a list of currently marked menu items, 1F-309
  - indicator — displays application specific alarms or working indicator, or both, on FMLI banner line, 1F-341
  - message — puts arguments on FMLI message line, 1F-520
  - pathconv — converts an alias to its pathname, 1F-621
  - readfile, longline — reads file, gets longest line, 1F-697
  - regex — match patterns against a string, 1F-703
  - reinit — changes the descriptors in the initialization file, 1F-705
  - reset — (FMLI utility) changes the entry in a field of a form to its default value, 1F-708
  - run — runs a program, 1F-722
  - set, unset — set and unset local or global environment variables, 1F-778
  - setcolor — redefine or create a color, 1F-779
  - shell — run a command using shell, 1F-795
  - test — evaluates the expression expression, 1F-867
  - vsig — synchronize a co-process with the controlling FMLI application, 1F-953
  - fmt — simple text formatters, 1-286
  - fold — fold long lines, 1-292
  - fonts
    - prepends host resident PostScript fonts to files — download, 1-199
  - for — shell built-in functions to repeatedly execute action(s) for a selected number of times, 1-293
  - foreach — shell built-in functions to repeatedly execute action(s) for a selected number of
    - times, 1-293
  - formatters, text
    - fmt, 1-286
  - Forms and Menu Language Interpreter, See FMLI
  - FORTTRAN
    - create a tags file for use with ex and vi — ctags, 1-170
  - Framed Access Command Environment, see face
  - frameID number (FMLI utility) — getfrm, 1F-308
  - from — sender of mail messages, 1B-294
  - ftp — file transfer program, 1-295
  - function — shell built-in command to define a function which is usable within this shell, 1-304
- G**
- gencat — generate a formatted message catalog, 1-306
  - getfrm — (FMLI utility) returns the current frameID number, 1F-308
  - getitems — (FMLI utility) returns a list of currently marked menu items, 1F-309
  - getopt — parse command options, 1-310, 1-312
  - getoptcvt — parse command options, 1-312
  - getoptcvt — parse command options, 1-312, 1-314
  - gettext — retrieve text string from message database, 1-316, 1-317
  - glob — shell built-in function to expand a word list, 1-319
  - goto — shell built-in functions to enable the execution of the shell to advance beyond its sequence of steps, 1-248
  - gprof — call-graph profile, 1-320
  - graph — draw graph, 1-323
  - graphics
    - interpolate smooth curve — spline, 1-811
  - graphics filters for plotters — plot, 1B-643
  - grep
    - search a file for a character string — fgrep, 1-267
    - search a file for a pattern — grep, 1-325
    - search a file for a pattern using full regular

expressions

**grep**, *continued*  
 — **egrep**, 1-228

**group IDs**  
 change the group ownership of a file — **chgrp**, 1-80  
 display a list of all valid group names — **disgid**, 1-196  
 prompts for group ID — **ckgid**, 1-92  
 provides error message for group ID — **errgid**, 1-92  
 provides help message for group ID — **helpgid**, 1-92  
 validates group ID — **valgid**, 1-92

**groups** — display group membership, 1B-328, 1-327

**grpck** — check group database entries, 1B-329

## H

**halt system without checking disks** — **fasthalt**, 1B-263

**hash** — shell built-in functions to evaluate the internal hash table of the contents of directories, 1-330

**hashstat** — shell built-in functions to evaluate the internal hash table of the contents of directories, 1-330

**head** — display first few lines of files, 1-331

**history** — shell built-in functions to re-use previous command-lines from the current shell, 1-332

**HOME variable** — **sh**, 1-783

**host machines, local**  
 show status — **ruptime**, 1-725  
 who's logged in — **rwho**, 1-728

**host machines, remote**  
 display status of network hosts (RPC version) — **rup**, 1-723  
 who's logged in — **rusers**, 1-727

**host resident PostScript font downloader** — **download**, 1-199

**hostid** — display host ID, 1B-333

**hostname** — display host name, 1B-334

## I

**i386** — get processor type truth value, 1-467

**iAPX286** — get processor type truth value, 1-467

**if** — shell built-in functions to evaluate expression(s) or to make execution of actions dependent upon the evaluation of expression(s), 1-337

**IFS variable** — **sh**, 1-783

**indicator** — (FMLI utility) displays application specific alarms or working indicator, or both, on FMLI banner line, 1F-341

**indxbib** — create an inverted index to a bibliographic database, 1-342

**install** — install files, 1B-343

**integer**  
 prompts for an integer — **ckint**, 1-94  
 provides error message for integer — **errint**, 1-94  
 provides help message for integer — **helpint**, 1-94  
 validates an integer — **valint**, 1-94

**integer, range**  
 prompts for an integer within a specified range — **ckrange**, 1-104  
 provides error message for integer within a specified range — **errrange**, 1-104  
 provides help message for integer within a specified range — **helprange**, 1-104  
 validate an integer within a specified range — **valrange**, 1-104

**Internet**  
 transfer files to and from a remote machine — **tftp**, 1-869  
 transfer of files to and from remote network sites — **ftp**, 1-295  
 user name directory service — **whois**, 1-969

**interprocess communication**  
 remove a message queue, semaphore set, or shared memory ID — **ipcrm**, 1-344  
 report status — **ipcs**, 1-345

**ipcrm** — remove a message queue, semaphore set, or shared memory ID, 1-344

**ipcs** — report inter-process communication facilities status, 1-345

## J

- jobs — shell built-in functions to control process execution, 1-348
- join — relational database operator, 1-352
- jsh — the job control shell command interpreter

## K

- kbd — manipulate the state of keyboard or display the type of keyboard, 1-354
- Kerberos login utility
  - kinit, 1-362
- Kerberos system
  - introduction — Kerberos, 1-356
- Kerberos ticket-granting-ticket
  - fetch and store using service key — ksrvtgt, 1-397
- Kerberos tickets
  - destroy — kdestroy, 1-355
  - list currently held — klist, 1-363
- keyboard
  - load and dump keyboard translation tables — loadkeys, dumpkeys, 1-423
  - manipulate the state of keyboard or display the type of keyboard — kbd, 1-354
- keylogin — decrypt and store secret key with keyserv, 1-358
- keylogout — delete stored secret key with keyserv, 1-359
- keywords
  - prompts for and validates a keyword — ckkeywd, 1-99
- kill — terminate a process by default, 1-360
- KornShell
  - aliasing — ksh, 1-366
  - arithmetic evaluation — ksh, 1-373
  - blank interpretation — ksh, 1-372
  - command execution — ksh, 1-378
  - command re-entry — ksh, 1-379
  - command substitution — ksh, 1-367
  - commands — ksh, 1-364
  - comments — ksh, 1-366
  - conditional expressions — ksh, 1-374
  - definitions — ksh, 1-364

## KornShell, *continued*

- emacs editing mode — ksh, 1-380
  - environment — ksh, 1-376
  - file name generation — ksh, 1-372
  - functions — ksh, 1-377
  - I/O — ksh, 1-375
  - in-line editing options — ksh, 1-379
  - invocation — ksh, 1-394
  - jobs — ksh, 1-377, 1-350
  - parameter substitution — ksh, 1-368
  - process substitution — ksh, 1-367
  - prompting — ksh, 1-374
  - quoting — ksh, 1-373
  - restricted command and programming language — rksh, 1-364
  - signals — ksh, 1-378
  - special commands — ksh, 1-385
  - tilde substitution — ksh, 1-367
  - vi editing mode — ksh, 1-382
- ksh — KornShell, a standard command and programming language, 1-364

## L

- languages
  - C compiler — cc, 1B-76
  - C preprocessor — cpp, 1-135
  - create C error messages — mkstr, 1B-526
  - extract strings from C code — xstr, 1-977
  - lex — generate lexical analyzer, 1-408
- last — display login and logout information about users and terminals, 1-398
- lastcomm — display the last commands executed, in reverse order, 1-399
- LC\_ALL variable — sh, 1-784
- LC\_CTYPE variable — sh, 1-784
- LC\_MESSAGES variable — sh, 1-784
- ld — link editor for object files, 1-400
- ld — link editor, 1B-405
- ldd — list dynamic dependencies of executable files or shared objects, 1-406
- let — shell built-in function to evaluate one or more arithmetic expressions, 1-407
- LEX
  - create a tags file for use with ex and vi —

---

ctags,  
**LEX, *continued***  
 1-170  
 lex — generate lexical analyzer, 1-408  
 library archive  
   find ordering relation for an object or library  
   archive — `lorder`, 1-434  
 limit — shell built-in functions to set/get limita-  
   tions on the system resources available to the  
   current shell and it's descendents., 1-412  
 line — read one line from standard input and  
   write to standard output, 1-414  
 line numbering filter — `nl`, 1-599  
 line printer control — `lpc`  
 link  
   make hard or symbolic links to files — `ln`,  
   1-416  
 link editor — `ld`, 1B-405, 1-400  
 list  
   contents of directory — `ls`, 1-454  
   file names — `ls`, 1B-459  
 list of service grades  
   print — `uuglist`, 1C-925  
 list, sorted  
   find lines — `look`, 1-432  
 listusers — list user login information, 1-415  
 ln — make hard or symbolic links to files, 1-416  
 loadfont — display or change font information in  
   the RAM of the video card on an x86 system in  
   text mode, 1-421  
 loadkeys — load keyboard translation tables,  
   1-423  
 log, system  
   add entries — `logger`, 1-424  
 logger — add entries to the system log, 1-424,  
   1B-425  
 login  
   change login password and password attributes  
   — `passwd`, 1-615  
   display effective user name — `whoami`, 1B-968  
   display login and logout information about  
   users and terminals — `last`, 1-398  
   get the name of the user running the process —  
   logname,  
   1-430  
   list user login information — `listusers`,  
   1-415  
   remote — `rlogin`, 1-709  
   spawn login to a remote terminal — `ct`,  
   1C-168  
   who is logged in, and what are they doing —  
   w, 1-954  
 login — sign on to the system, 1-426  
 login environment  
   display variables — `printenv`, 1B-668  
   set terminal characteristics — `tset`, 1B-900  
 login password  
   change in NIS — `yppasswd`, 1-983  
 logname — get the name of the user running the  
   process, 1-430  
 logout — shell built-in function to exit from a  
   login session, 1-431  
   display login and logout information about  
   users and terminals — `last`, 1-398  
 look — find words in the system dictionary or  
   lines in a sorted list, 1-432  
 lookbib — find references in a bibliographic data-  
   base, 1-433  
 lorder — find ordering relation for an object or  
   library archive, 1-434  
 lp — send requests to an LP print service, 1-435  
 LP print service  
   print information about the status — `lpstat`,  
   1-451  
   send/cancel requests to an LP print service —  
   lp, cancel, 1-435  
 lpc — line printer control, 1B-441  
 lpq — display printer queue, 1B-444  
 lpr — print files, 1B-446  
 lprm — remove print jobs, 1B-449  
 lpstat — print information about the status of the  
   LP print service, 1-451  
 lptest — generate printer test pattern, 1B-453  
 ls — list contents of directory, 1-454, 1B-459

---

## M

m4 — macro processor, 1-462  
mach — display processor type of current host, 1B-466  
machid — get processor type truth value, 1-467  
machine IDs  
  get processor type truth value — machid, 1-467  
macro processor — m4, 1-462  
magnetic tape  
  backspace files — mt, 1-535  
  backspace records — mt, 1-535  
  copy — tcopy, 1-858  
  erase — mt, 1-535  
  forward space files — mt, 1-535  
  forward space records — mt, 1-535  
  get unit status — mt, 1-535  
  manipulate — mt, 1-535  
  place unit off-line — mt, 1-535  
  retension — mt, 1-535  
  rewind — mt, 1-535  
  skip backward files — mt, 1-535  
  skip backward records — mt, 1-535  
  skip forward files — mt, 1-535  
  skip forward records — mt, 1-535  
  write EOF mark on — mt, 1-535  
mail — read mail or send mail to users  
  automatic replies — vacation, 1-934  
mail services  
  mail notifier — biff, 1B-67  
  sender of mail messages — from, 1B-294  
mail utilities  
  create aliases database — newaliases, 1-544  
  statistics — mailstats, 1-474  
MAIL variable — sh, 1-783  
mailbox  
  storage for incoming mail — mailx, 1-475  
MAILCHECK variable — sh, 1-783  
MAIL variable — sh, 1-783  
mailstats — mail delivery statistics, 1-474  
mailx — interactive message processing system  
mailx commands  
  — !, 1-478  
  — #, 1-478

mailx commands, *continued*  
  — =, 1-478  
  — ?, 1-478  
  — |, 1-482  
  — alias, 1-478  
  — alternates, 1-478  
  — cd, 1-478  
  — chdir, 1-478  
  — Copy, 1-479  
  — delete, 1-479  
  — discard, 1-479  
  — dp, 1-479  
  — dt, 1-479  
  — echo, 1-479  
  — edit, 1-479  
  — else, 1-480  
  — endif, 1-480  
  — exit, 1-479  
  — field, 1-479  
  — file, 1-479  
  — folder, 1-479  
  — Followup, 1-480  
  — from, 1-480  
  — group, 1-478  
  — headers, 1-480  
  — help, 1-480  
  — hold, 1-480, 1-482  
  — if, 1-480  
  — ignore, 1-479  
  — inc, 1-480  
  — list, 1-481  
  — load, 1-481  
  — mail, 1-481  
  — mbox, 1-481  
  — More, 1-481  
  — New, 1-481  
  — next, 1-482  
  — Page, 1-481  
  — pipe, 1-482  
  — preserve, 1-480, 1-482  
  — Print, 1-482, 1-484  
  — Put, 1-482  
  — quit, 1-482  
  — Reply, 1-483, 1-482  
  — replyall, 1-483

- 
- mailx commands, *continued*
    - replysender, 1-482
    - Respond, 1-483, 1-482
    - retain, 1-483
    - Save, 1-483
    - set, 1-483
    - shell, 1-484
    - size, 1-484
    - source, 1-484
    - top, 1-484
    - touch, 1-484
    - Type, 1-482, 1-484
    - unalias, 1-484
    - undelete, 1-484
    - undiscard, 1-484
    - ungroup, 1-484
    - unignore, 1-484
    - Unread, 1-481, 1-484
    - unretain, 1-484
    - unset, 1-485
    - version, 1-485
    - visual, 1-485
    - write, 1-485
    - xit, 1-485
    - z, 1-485
  - maintain groups of programs — *sysV-make*
  - maintain programs — *make*
  - make — maintain, update, and regenerate related programs and files
    - Appending to a Macro, 1S-500
    - Bourne Shell Constructs, 1S-507
    - Clearing Special Targets, 1S-499
    - Command Dependencies, 1S-499
    - Command Execution, 1S-507
    - Command Substitutions, 1S-508
    - Conditional Macro Definitions, 1S-502
    - Dynamic Macros, 1S-501
    - Global, 1S-496
    - Hidden Dependencies, 1S-499
    - Implicit Rules, 1S-504
    - implicit rules, list of <make/make.rules> — <make, 1S-506
    - Library Maintenance, 1S-507
    - Macros, 1S-497, 1S-499
    - Makefile Target Entries, 1S-496
    - make — maintain, update, and regenerate related programs and files, *continued*
    - Pattern Matching Rules, 1S-504
    - Pattern Replacement Macro References, 1S-500
    - Predefined Macros, 1S-502
    - Reading Makefiles and the Environment, 1S-495
    - Rules, 1S-498
    - Signals, 1S-508
    - Special Characters, 1S-496
    - Special-Function Targets, 1S-498
    - Special-Purpose Macros, 1S-500
    - Suffix Replacement Macro References, 1S-500
    - Suffix Rules, 1S-504
    - System V version of *make* — *sysV-make*
    - Targets and Dependencies, 1S-496
    - The Suffixes List, 1S-507
  - make — build programs
  - man — online display of reference pages, 1-512
  - manual page files
    - locate — *whereis*, 1B-961
  - manual pages
    - describe command — *whatis*, 1-960
  - matrix display program for PostScript printers — *postmd*, 1-652
  - mbox
    - storage file for read mail — *mailx*, 1-475
  - mconnect — open connection to remote mail server, 1-516
  - mcs — manipulate the comment section of an object file, 1-517
  - menu item
    - builds a menu and prompts user to choose one item from menu — *ckitem*, 1-96
    - provides error message for menu item — *erritem*, 1-96
    - provides help message for menu item — *helpitem*, 1-96
  - menu items, FMLI
    - returns a list of — *getitems*, 1F-309
  - mesg — permit or deny messages via *write*, 1-519
  - message — puts arguments on FMLI message line, 1F-520
  - messages
    - create message object file — *msgfmt*, 1-532

---

messages, *continued*

creating portable object files — `msgfmt`, 1-532  
display contents of, or search for a text string  
in, message data bases — `srchtxt`,  
1-813  
display on stderr or system console —  
`fmtmsg`, 1-288  
editing messages — `msgfmt`, 1-532  
extract `gettext` call strings — `xgettext`,  
1-975  
generate a formatted message catalog — `gen-  
cat`, 1-306  
permit or deny messages via `write` — `msg`,  
1-519  
retrieve text string from message database —  
`gettext`, 1-316  
setting the domain — `msgfmt`, 1-532  
setting the message identifier — `msgfmt`, 1-532  
setting the message string — `msgfmt`, 1-532  
`mkdir` — make directories, 1-522  
`mkmsgs` — create message files for use by `gettext`,  
1-524  
`mkstr` — create C error messages, 1B-526  
`more` — browse through a text file, 1-528  
`msgfmt` — create message object file, 1-532  
`mt` — manipulate magnetic tape, 1-535  
`mv` — move files, 1-537

## N

`nawk` — pattern scanning and processing language,  
1-539  
`neqn` — mathematical typesetting, 1-234  
`newaliases` — make mail aliases database, 1-544  
`newform` — change the format of a text file, 1-545  
`newgrp` — shell built-in function to allow new  
group permissions to the user, 1-548  
`news` — print news items, 1-549  
NFS, secure  
decrypt and store secret key with `keyser` —  
`keylogin`, 1-358  
delete stored secret key with `keyser`— `keylo-  
gout`, 1-359  
`nice` — run a command at low priority, 1-551  
change process nice value — `renice`, 1B-706

## NIS, See NIS+

change login password in — `yppasswd`, 1-983  
print the value of one or more keys from a NIS  
map — `ypmatch`, 1-982  
print values in a NIS database — `ypcat`, 1-981  
return name of NIS server or map master —  
`ypwhich`, 1-984

## NIS+

Authentication — `nis+`, 1-558  
Authorization — `nis+`, 1-558  
change access rights on a NIS+ object —  
`nischmod`, 1-569  
change password information — `nispasswd`,  
1-587  
change the group owner of a NIS+ object —  
`nischgrp`, 1-567  
change the owner of a NIS+ object —  
`nischown`, 1-571  
change the time to live of a NIS+ object —  
`nischttl`, 1-573  
Concatenation Path — `nis+`, 1-555  
create NIS+ directories — `nismkdir`, 1-585  
Directories and Domains — `nis+`, 1-558  
Directory Authorization — `nis+`, 1-559  
display NIS+ defaults — `nisdefaults`, 1-575  
display NIS+ error messages — `niserror`,  
1-577  
display tables — `niscat`, 1-565  
Grammar — `nis+`, 1-554  
group administration — `nisgrpadm`, 1-578  
Group Names — `nis+`, 1-557  
Indexed Names — `nis+`, 1-554  
list the contents of a NIS+ directory — `nisls`,  
1-582  
Name Expansion — `nis+`, 1-554  
Namespaces — `nis+`, 1-556  
NIS+ Administrative Commands — `nis+`,  
1-561  
NIS+ Directory Object — `nis+`, 1-552  
NIS+ Files and Directories — `nis+`, 1-563  
NIS+ Group Object — `nis+`, 1-553  
NIS+ Link Object — `nis+`, 1-553  
NIS+ Programming API — `nis+`, 1-562  
NIS+ Table Object — `nis+`, 1-552  
NIS+ User Commands — `nis+`, 1-561

---

NIS+, *continued*

Principal Names — `nis+`, 1-556  
remove directories — `nisrmdir`, 1-591  
remove objects — `nism`, 1-590  
return the state of the NIS+ namespace using a conditional expression — `nistest`, 1-597  
Simple Names — `nis+`, 1-553  
symbolically link NIS+ objects — `nisl`, 1-580  
table administration tool — `nistbladm`, 1-593  
Table Authorization — `nis+`, 1-560  
utilities for searching NIS+ tables —  
`nismatch`, `nisgrep`, 1-583  
`niscat` — display NIS+ tables, 1-565  
`nischgrp` — change the group owner of a NIS+ object, 1-567  
`nischmod` — change access rights on a NIS+ object, 1-569  
`nischown` — change the owner of a NIS+ object, 1-571  
`nischttl` — change the time to live of a NIS+ object, 1-573  
`nisdefaults` — display NIS+ defaults, 1-575  
`niserror` — display NIS+ error messages, 1-577  
`nisgrep` — utility for searching NIS+ tables, 1-583  
`nisgrpadm` — NIS+ group administration command, 1-578  
`nisl` — symbolically link NIS+ objects, 1-580  
`nisl` — list the contents of a NIS+ directory, 1-582  
`nismatch` — utility for searching NIS+ tables, 1-583  
`nismkdir` — create a NIS+ directory, 1-585  
`nism` — remove NIS+ objects, 1-590  
`nisrmdir` — remove a NIS+ directory, 1-591  
`nistbladm` — administer NIS+ tables, 1-593  
`nistest` — return the state of the NIS+ namespace using a conditional expression, 1-597  
`nl` — number lines, 1-599  
`nm` — print name list of an object file, 1-601  
`nohup` — run a command immune to hangups and quits, 1-603  
`notify` — shell built-in functions to control process execution, 1-348

`notify` user that volume requested is not in the CD-ROM or floppy drive — `volmissing`, 1-952  
`nroff` — format documents for display or line-printer, 1-605  
`nroff` utilities  
check `nroff` and `troff` files — `checknr`, 1-79  
eliminate `.so`'s from `nroff` input — `soelim`, 1-803  
filters reverse line-feeds from two-column `nroff` text — `col`, 1-119  
format tables — `tbl`, 1-856  
remove `nroff`, `troff`, `tbl` and `eqn` constructs — `deroff`, 1-185

## O

object archive  
find ordering relation for an object or library archive — `lorder`, 1-434  
object files  
find printable strings — `strings`, 1-819  
manipulate the comment section — `mcs`, 1-517  
print section sizes in bytes — `size`, 1-800  
strip symbol table, debugging and line number information — `strip`, 1-820  
octal dump  
— `od`, 1-607  
`od` — octal dump, 1-607  
`on` — execute a command on a remote system, but with the local environment, 1-609  
`onintr` — shell built-in functions to respond to (hardware) signals, 1-891  
online reference pages — `man`, 1-512

## P

`pack` — compress files, 1-611  
`page` — page through a text file, 1-528  
`pagesize` — display size of a page of memory, 1B-614  
Pascal  
create a tags file for use with `ex` and `vi` — `ctags`, 1-170  
`passwd` — change login password and password attributes, 1-615

---

password  
 change in NIS — `yppasswd`, 1-983

password file  
 edit — `vipw`, 1B-949

passwords  
 change login password and password attributes  
 — `passwd`, 1-615

paste — merge same lines of several files or subsequent lines of one file, 1-619

PATH variable — `sh`, 1-783

pathconv — search FMLI criteria for filename, 1F-621

pathname  
 prompts for a pathname — `ckpath`, 1-101  
 provides error message for pathname — `errpath`, 1-101  
 provides help message for pathname — `help-path`, 1-101  
 validates pathname — `valpath`, 1-101

pcat — compress files, 1-611

pcmapkeys — set keyboard extended map and scancode translation for the PC console in text mode, 1-622

pdp11 — get processor type truth value, 1-467

performance monitoring  
 display call-graph profile data — `gprof`, 1-320  
 resource usage for a command — `rusage`, 1B-726  
 time a command — `time`, 1-871  
 time a command; report process data and system activity — `timex`, 1-873

pg — files perusal filter for CRTs, 1-630

pkginfo — display software package information, 1-634

pkgmk — produce an installable package, 1-636

pkgparam — displays package parameter values, 1-638

pkgproto — generate prototype file entries for input to `pkgmk` command, 1-639

pkgtrans — translate package format, 1-641

plot — graphics filters for plotters, 1B-643

plotters  
 graphics filters — `plot`, 1B-643

popd — shell built-in functions to change the current working directory, 1-77

postplot — PostScript translator for `plot(4)` graphics files, 1-655

postdaisy — PostScript translator for Diablo 630 daisy-wheel files, 1-645

postdmd — PostScript translator for DMD bitmap files, 1-647

postio — serial interface for PostScript printers, 1-649

postmd — matrix display program for PostScript printers, 1-652

postprint — PostScript translator for text files, 1-657

postprocessors  
 troff for PostScript printers — `dpost`, 1-201

postreverse — reverse the page order in a PostScript file, 1-659

PostScript  
 matrix display program — `postmd`, 1-652  
 prepends host resident PostScript fonts to files — `download`, 1-199  
 reverse the page order in a PostScript file — `postreverse`, 1-659  
 serial interface — `postio`, 1-649  
 translator for Diablo 630 daisy-wheel files — `postdaisy`, 1-645  
 translator for DMD bitmap files — `postdmd`, 1-647  
 translator for `plot(4)` graphics files — `postplot`, 1-655  
 translator for Tektronix 4014 files — `posttek`, 1-661  
 translator for text files — `postprint`, 1-657  
 troff postprocessor for PostScript printers — `dpost`, 1-201

PostScript translator for Diablo 630 daisy-wheel files — `postdaisy`, 1-645

PostScript translator for MD bitmap files — `postdmd`, 1-647

PostScript translator for Tektronix 4014 files — `posttek`, 1-661

PostScript translator for text files — `postprint`,

1-657

posttek — PostScript translator for Tektronix 4014 files, 1-661

pr — print files, 1-663

prime factors  
 obtain for a number — factor, 1-262

print — shell built-in function to output characters to the screen or window, 1-667  
 formatted output — printf, 1-669  
 print files — pr, 1-663

print files — lpr, 1B-446  
 prepends host resident PostScript fonts to files — download, 1-199

printenv — display environment variables, 1B-668

printer, line  
 control — lpc, 1B-441

printers  
 display queue — lpq, 1B-444  
 print information about the status — lpstat, 1-451  
 remove jobs from queue — lprm, 1B-449  
 send/cancel requests — lp, cancel, 1-435  
 test — lptest, 1B-453

printers, LP  
 — disable, 1-232  
 — enable, 1-232

print — print formatted output, 1-669

process accounting  
 search and print files — acctcom, 1-25  
 time a command; report process data and system activity — timex, 1-873

process scheduler  
 display or set scheduling parameters of specified process(es) — priocntl, 1-671

process status  
 report — ps, 1-681

process, running  
 change priority — renice, 1B-706

processes  
 display status — ps, 1B-685  
 get core images of running processes — gcore, 1-305  
 terminate a process by default — kill, 1-360

processors  
 display type — mach, 1B-466

prof — display profile data, 1-678

profile  
 display call-graph — gprof, 1-320

programming languages  
 analyze and disperse compiler error messages — error, 1-237  
 C compiler — cc, 1B-76  
 C preprocessor — cpp, 1-135  
 extract strings from C code — xstr, 1-977  
 formats program in nice style using troff — vgrind, 1-939  
 lex — generate lexical analyzer, 1-408

programming tools  
 arbitrary precision arithmetic language — bc, 1-63  
 assembler — as, 1-40  
 create a tags file for use with ex and vi — ctags, 1-170  
 create C error messages — mkstr, 1B-526  
 debugger — adb, 1-27  
 display call-graph profile data — gprof, 1-320  
 dump selected parts of an object file — dump, 1-206  
 find printable strings in an object or binary file — strings, 1-819  
 — install, 1B-343  
 lex — generate lexical analyzer, 1-408  
 link editor — ld, 1B-405  
 link editor for object files — ld, 1-400  
 macro processor — m4, 1-462  
 make — build programs  
 object code disassembler — dis, 1-194  
 print name list of an object file — nm, 1-601  
 print section sizes in bytes of object files — size, 1-800  
 regular expression compile — regcmp, 1-701  
 resolve and remove ifdef ed lines from C program source — undef, 1-914  
 resource usage for a command — rusage, 1B-726  
 RPC protocol compiler — rpcgen, 1-715  
 Source Code Control System — sccs, 1-759  
 strip symbol table, debugging and line number

---

information  
programming tools, *continued*  
from an object file — `strip`, 1-820  
time a command — `time`, 1-871  
update access time and/or modification time of  
a file — `touch`, 1-883  
`ps` — display process status, 1B-685  
PS1 variable — `sh`, 1-783  
PS2 variable — `sh`, 1-783  
`pushd` — shell built-in functions to change the  
current working directory, 1-77  
`pwd` — print working directory name, 1-688

## Q

`queue`, printer  
display — `lpq`, 1B-444  
`queues`  
display the jobs queued to run at specified  
times — `atq`, 1-46  
remove jobs spooled by `at` or `batch` —  
`atrm`, 1-47

## R

`rcp` — remote file copy, 1-689  
`rdist` — remote file distribution, 1-691  
`read` — shell built-in function to receive from stan-  
dard input (keyboard), 1-696  
`readfile`, `longline` — (FMLI utility) reads file,  
gets longest line, 1F-697  
`readonly` — shell built-in function to protect the  
value of the given variable from reassignment,  
1-698  
reboot system without checking disks — `fast-`  
`boot`, 1B-263  
`red` — text editor, 1-213  
`refer` — expand and insert references from a  
bibliographic database, 1-699  
`regcmp` — regular expression compile, 1-701  
regenerate groups of programs — `sysV-make`  
regenerate programs — `make`  
`regex` — (FMLI utility) match patterns against a  
string, 1F-703  
`rehash` — shell built-in functions to evaluate the

internal hash table of the contents of direc-  
tories, 1-330  
`reinit` — (FMLI utility) changes the descriptors in  
the initialization file, 1F-705  
relational database  
— `join`, 1-352  
reminder services  
— `calendar`, 1-70  
mail notifier — `biff`, 1B-67  
remote shell — `rsh`, 1-719  
remote system  
connect — `tip`, 1-875  
connect to — `cu`, 1C-172  
execute a command on a remote system, but  
with the local environment — `on`,  
1-609  
file copy — `rcp`, 1-689  
file distribution — `rdist`, 1-691  
remote login — `rlogin`, 1-709  
shell — `rsh`, 1-719  
show status — `rup`, 1C-724, 1-723  
spawn login — `ct`, 1C-168  
system to system command execution — `uux`,  
1C-931  
transfer files to and from — `tftp`, 1-869  
who's logged in on remote machines —  
`rusers`, 1-727  
`renice` — alter priority of running processes,  
1B-706  
`repeat` — shell built-in function to execute a com-  
mand more than once, 1-707  
repeated lines in file  
`repot` — `uniq`, 1-915  
`reset` — (FMLI utility) changes the entry in a field  
of a form to its default value, 1F-708  
`reset` — reset terminal bits, 1B-900  
reset terminal bits — `reset`, 1B-900  
`return` — shell built-in functions to enable the exe-  
cution of the shell to advance beyond its  
sequence of steps, 1-248  
reverse page order  
PostScript file — `postreverse`, 1-659  
reverse the page order in a PostScript file — `pos-`

treverse, 1-659  
 rksh — KornShell, restricted command and programming language, 1-364  
 rlogin — remote login, 1-709  
 rm — remove files, 1-711  
 rmail — only permits sending of mail, 1-469  
 rmdir — remove directories, 1-711  
 roffbib — format and print bibliographic database, 1-713  
 RPC  
   display host status of remote machines — `rup`, 1C-724  
   display status of network hosts — `rup`, 1-723  
   protocol compiler — `rpcgen`, 1-715  
 RPC Language  
   RPC protocol compiler — `rpcgen`, 1-715  
 RPC, secure  
   decrypt and store secret key with `keyserv` — `keylogin`, 1-358  
   delete stored secret key with `keyserv` — `keylogout`, 1-359  
`rpcgen` — RPC protocol compiler, 1-715  
 rsh — remote shell, 1-719  
 run — (FMLI utility) runs a program, 1F-722  
 run a command immune to hangups and quits — `nohup`, 1-603  
`rup` — display status of network hosts, 1C-724, 1-723  
`ruptime` — display status of local hosts, 1-725  
`rusage` — resource usage for a command, 1B-726  
`rusers` — who's logged in on remote machines, 1-727  
`rwho` — who's logged in on local machines, 1-728

## S

sag — system activity graph, 1-729  
 sar — system activity reporter, 1-731  
 SCCS  
   extract SCCS version information from a file — `what`, 1-959  
 sccs — Source Code Control System, 1-759  
 SCCS commands  
   admin — create and administer SCCS history

files,  
 SCCS commands, *continued*  
   1-734  
   cdc — change the delta commentary of an SCCS delta, 1-738  
   comb — combine deltas, 1-740  
   delta — change the delta commentary of an SCCS delta, 1-741  
   get — retrieve a version of an SCCS file, 1-743  
   help — help regarding SCCS error or warning messages, 1-748  
   prt — display delta table information from an SCCS file, 1-752  
   rmdel — remove a delta from an SCCS file, 1-754  
   sact — show editing activity status of an SCCS file, 1-755  
   sccs-prs — display selected portions of an SCCS history, 1-749  
   sccsdiff — compare versions of SCCS file, 1-756  
   unget — unget SCCS file, 1-757  
   val — validate SCCS file, 1-758  
 SCCS delta  
   change commentary — `sccs-cdc`, 1-738  
   combine — `sccs-comb`, 1-740  
   create — `delta`, 1-741  
   remove — `rmdel`, 1-754  
 SCCS delta table  
   print form an SCCS file — `sccs-prt`, 1-752  
 SCCS files  
   compare versions — `sccs-sccsdiff`, 1-756  
   retrieve a version of a file — `sccs-get`, 1-743  
   show editing activity status — `sccs-sact`, 1-755  
   undo a previous get of an SCCS file — `sccs-unget`, 1-757  
   validate — `sccs-val`, 1-758  
 SCCS help  
   regarding SCCS error or warning messages — `sccs-help`, 1-748  
 SCCS history  
   display selected portions — `sccs-prs`, 1-749  
 SCCS history files  
   create and administer — `sccs-admin`, 1-734

---

**sccs-admin** — create and administer SCCS history files, 1-734  
**sccs-cdc** — change the delta commentary of an SCCS delta, 1-738  
**sccs-comb** — combine deltas, 1-740  
**sccs-delta** — change the delta commentary of an SCCS delta, 1-741  
**sccs-get** — retrieve a version of an SCCS file, 1-743  
**sccs-help** — help regarding SCCS error or warning messages, 1-748  
**sccs-prs** — display selected portions of an SCCS history, 1-749  
**sccs-prt** — display delta table information from an SCCS file, 1-752  
**sccs-rmdel** — remove delta from SCCS file, 1-754  
**sccs-sact** — show editing activity status of an SCCS file, 1-755  
**sccs-sccsdiff** — compare versions of SCCS file, 1-756  
**sccs-unget** — unget SCCS file, 1-757  
**sccs-val** — validate SCCS file, 1-758  
**screen-oriented editor** — **vi**, 1-942  
**script** — make script of terminal session, 1-766  
**sdiff** — print differences between two files side-by-side, 1-767  
**sed** — stream editor, 1-769  
**select** — shell built-in functions to choose from among a list of actions, 1-72  
**serial interface for PostScript printers** — **postio**, 1-649  
**set** — shell built-in functions to determine the characteristics for environmental variables of the current shell and its descendents, 1-773  
**set keyboard extended map and scancode translation for the PC console in text mode** — **pcmap-keys**, 1-622  
**set, unset** — (FLMI utility) set and unset local or global environment variables, 1F-778  
**setcolor** — (FMLI utility) redefine or create a color, 1F-779  
**setenv** — shell built-in functions to determine the characteristics for environmental variables of the current shell and its descendents, 1-773  
**sh** — the standard shell command interpreter  
**SHACCT variable** — **sh**, 1-784  
**shell** — (FMLI utility) run a command using shell, 1F-795  
**shell command interpreter built-in functions** — **shell\_builtins**, 1-796  
**shell command interpreter builtin-functions**  
— **alias**, 1-35  
— **bg**, 1-348  
— **break**, 1-68  
— **case**, 1-72  
— **cd**, 1-77  
— **chdir**, 1-77  
— **continue**, 1-68  
— **dirs**, 1-77  
— **eval**, 1-246  
— **exit**, 1-248  
— **fc**, 1-332  
— **fg**, 1-348  
— **for**, 1-293  
— **foreach**, 1-293  
— **function**, 1-304  
— **getopts**, 1-314  
— **glob**, 1-319  
— **hash**, 1-330  
— **hashstat**, 1-330  
— **history**, 1-332  
— **if**, 1-337  
— **jobs**, 1-348  
— **kill**, 1-360  
— **let**, 1-407  
— **limit**, 1-412  
— **logout**, 1-431  
— **newgrp**, 1-548  
— **notify**, 1-348  
— **onintr**, 1-891  
— **popd**, 1-77  
— **print**, 1-667  
— **pushd**, 1-77  
— **read**, 1-696  
— **readonly**, 1-698  
— **rehash**, 1-330  
— **repeat**, 1-707

---

shell command interpreter builtin-functions, *continued*

- return, 1-248
- select, 1-72
- set, 1-773
- setenv, 1-773
- shift, 1-798
- source, 1-246
- stop, 1-348
- suspend, 1-834
- switch, 1-72
- test, 1-337
- times, 1-872
- trap, 1-891
- type, 1-907
- typeset, 1-907
- ulimit, 1-412
- umask, 1-911
- unalias, 1-35
- unhash, 1-330
- unlimit, 1-412
- unset, 1-773
- unsetenv, 1-773
- until, 1-964
- wait, 1-956
- whence, 1-907
- while, 1-964

shell programming

- echo arguments — echo, 1-209
- read one line from standard input and write to standard output — line, 1-414

shell scripts

- display size of page memory — pagesize, 1B-614
- provide truth values — true, false, 1-895

SHELL variable — sh, 1-784

shell variables, in Bourne shell

shell\_builtins — shell command interpreter built-in functions, 1-796

shells

- C shell — csh
- remote — rsh, 1-719
- the job control shell command interpreter — jsh, 1-780
- the standard shell command interpreter — sh,

1-780

shells, *continued*

- shift — shell built-in function to traverse either a shell's argument list or a list of field-separated words, 1-798
- show codeset table for the current locale — dumpcs, 1-208
- shutdown — shut down multiuser operation, 1B-799
- Simple Mail Transfer Protocol
  - connection to remote mailserver — mconnect, 1-516
- size — print section sizes in bytes of object files, 1-800
- sleep — suspend execution for an interval, 1-802
- SMTP, See Simple Mail Transfer Protocol
- soelim — eliminate .so's from nroff input, 1-803
- software package, displays package parameter values — pkgparam, 1-638
  - display information — pkginfo, 1-634
  - generate prototype file entries for input to pkgmk command — pkgproto, 1-639
  - produce an installable package — pkgmk, 1-636
  - translate package format — pkgtrans, 1-641
- sort — sort and/or merge files, 1-804
- sort, topological
  - items mentioned in input — tsort, 1-905
- sortbib — sort bibliographic database, 1-808
- source — shell built-in functions to execute other commands, 1-246
- Source Code Control System, See SCCS
- source files
  - locate — whereis, 1B-961
- sparc — get processor type truth value, 1-467
- spell — check spelling, 1-810
- spline — interpolate smooth curve, 1-811
- split — split a file into pieces, 1-812
- srchtxt — display contents of, or search for a text string in, message data bases, 1-813
- standard output
  - replicate — tee, 1-859
- statistics

statistics, *continued*

collected by sendmail — mailstats, 1-474

stop — shell built-in functions to control process execution, 1-348

strchg — change stream configuration, 1-816

strconf — query stream configuration, 1-816

stream editor — sed, 1-769

STREAMS

change or query stream configuration —  
strchg, strconf, 1-816

string

prompt for defined string answer — ckstr,  
1-107

provide an error message for defined string  
answer — errstr, 1-107

provide an help message for defined string  
answer — helpstr, 1-107

validate a defined string answer — valstr,  
1-107

strings — find printable strings in object or  
binary file, 1-819

strip — strip symbol table, debugging and line  
number information from an object file, 1-820

stty — set the options for a terminal, 1-821, 1B-827

sum — print checksum and block count for a file,  
1-832, 1B-833

sun — get processor type truth value, 1-467

SunOS/BSD Source Compatibility Package

— stty, 1B-827

SunOS/BSD Source Compatibility Package com-  
mands

— arch, 1B-39

— basename, 1B-62

— biff, 1B-67

— cc, 1B-76

— chown, 1B-89

— df, 1B-186

— du, 1B-204

— echo, 1B-211

— expr, 1B-255

— fastboot, 1B-263

— file, 1B-271

— from, 1B-294

— groups, 1B-328

SunOS/BSD Source Compatibility Package com-  
mands, *continued*

— grpck, 1B-329

— hostid, 1B-333

— hostname, 1B-334

— install, 1B-343

— ld, 1B-405

— ln, 1B-418

— logger, 1B-425

— lpc, 1B-441

— lpq, 1B-444

— lpr, 1B-446

— lprm, 1B-449

— lptest, 1B-453

— ls, 1B-459

— mach, 1B-466

— mkstr, 1B-526

— pagesize, 1B-614

— plot, 1B-643

— printenv, 1B-668

— ps, 1B-685

— renice, 1B-706

— rusage, 1B-726

— shutdown, 1B-799

— sum, 1B-833

— test, 1B-865

— tr, 1B-890

— tset, 1B-900

— users, 1B-920

— vipw, 1B-949

— whereis, 1B-961

— whoami, 1B-968

suspend — shell built-in function to halt the  
current shell, 1-834

suspend execution of command

— sleep, 1-802

switch — shell built-in functions to choose from  
among a list of actions, 1-72

symorder — update symbol table ordering, 1-835

system activity

graphical representation — sag, 1-729

reporter — sar, 1-731

time a command; report process data and sys-  
tem activity — timex, 1-873

system administration

---

system administration, *continued*  
— install, 1B-343

system call and signals  
trace — truss, 1-896

system log  
add entries — logger, 1-424

system name  
print — uname, 1-912

system to system command execution — uux,  
1C-931

system to system copy — uucp, 1C-921

system to system copy, public — uucp, 1C-929

system uptime  
display — uptime, 1-919

sysV-make — maintain, update, and regenerate  
groups of programs

## T

TAB characters  
expand to SPACE characters, and vice versa —  
expand, unexpand, 1-250

tables  
format for nroff or troff — tbl, 1-856

tabs — set tabs on a terminal, 1-842

tail — display last part of file, 1-846

talk — talk to another user, 1-848

tape  
backspace files — mt, 1-535  
backspace records — mt, 1-535  
erase — mt, 1-535  
forward space files — mt, 1-535  
forward space records — mt, 1-535  
get unit status — mt, 1-535  
place unit off-line — mt, 1-535  
retension — mt, 1-535  
rewind — mt, 1-535  
skip backward files — mt, 1-535  
skip backward records — mt, 1-535  
skip forward files — mt, 1-535  
skip forward records — mt, 1-535  
write EOF mark on — mt, 1-535

tape archives  
create — tar, 1-849

tape, magnetic

tape, magnetic, *continued*  
copy, blocking preserved — tcopy, 1-858  
manipulate — mt, 1-535  
scan — tcopy, 1-858

tar — create tape archives, and add or extract files,  
1-849

tbl — format tables for nroff or troff, 1-856  
remove nroff, troff, tbl and eqn con-  
structs — deroff, 1-185

tcopy — copy a magnetic tape, 1-858

tee — replicate the standard output, 1-859

telnet — user interface to a remote system using  
the TELNET protocol, 1-860

TELNET protocol  
user interface to a remote system using the  
TELNET protocol — telnet, 1-860

terminal  
set options — stty, 1-821  
set tabs — tabs, 1-842

terminal screen  
— clear, 1-115

terminal session  
make script— script, 1-766

terminals  
get name — tty, 1-906  
initialize a terminal or query terminfo database  
— tput, 1-884  
reset bits — reset, 1B-900  
set characteristics — stty, 1B-827, 1B-900

terminate a process by default — kill, 1-360

terminfo database  
initialize a terminal or query terminfo database  
— tput, 1-884

test — (FMLI utility) evaluates the expression  
expression, 1F-867, 1-337

test — condition evaluation, 1B-865

text editing  
screen-oriented (visual) display editor based on  
ex — vi, 1-942  
stream editor — sed, 1-769

text editor  
— ed, 1-213  
— edit, 1-224  
— ex, 1-240

---

text files  
 browse or page through a text file — `more`,  
 page, 1-528  
 change format — `newform`, 1-545

text formatter  
 format documents for display or line-printer —  
`nroff`, 1-605

text processing utilities  
 check spelling — `spell`, 1-810  
 concatenate and display files — `cat`, 1-74  
 display last part of file — `tail`, 1-846  
 pattern scanning and processing language —  
`awk`, 1-56, 1-539  
 report repeated lines in a file — `uniq`, 1-915  
 search a file for a character string — `fgrep`,  
 1-267  
 search a file for a pattern — `grep`, 1-325  
 search a file for a pattern using full regular  
 expressions — `egrep`, 1-228  
 sort and/or merge files — `sort`, 1-804  
 split a file into pieces — `split`, 1-812  
 translate characters — `tr`, 1B-890, 1-888  
 underline text — `ul`, 1-910

text retrieval tools  
 create message files for use by `gettxt` —  
`mkmsgs`, 1-524  
 retrieve text string from message database —  
`gettxt`, 1-317

tftp — trivial file transfer program, 1-869

tilde escape commands for mail  
 — `mail`, 1-485

time — time a command, 1-871  
 prompts for time — `cktime`, 1-109  
 provides error message for time — `errtime`,  
 1-109  
 provides help message for time — `helptime`,  
 1-109  
 validates time — `valtime`, 1-109

timed event services  
 display the jobs queued to run at specified  
 times — `atq`, 1-46  
 execute commands at a later time — `at`,  
`batch`, 1-42  
 reminder service — `calendar`, 1-70  
 remove jobs spooled by `at` or `batch` —  
`atrm`,

timed event services, *continued*  
 1-47  
 user crontab file — `crontab`, 1-141

times — shell built-in function to report time  
 usages of the current shell, 1-872

timex — time a command; report process data and  
 system activity, 1-873

tip — connect to remote system, 1-875

touch — update access time and/or modification  
 time of a file, 1-883

tput — initialize a terminal or query terminfo data-  
 base, 1-884

tr — translate characters, 1B-890, 1-888

translate characters — `tr`, 1B-890, 1-888

translates `exportfs` options to `share/unshare` com-  
 mands — `exportfs`, 1B-251

trap — shell built-in functions to respond to  
 (hardware) signals, 1-891

Trivial File Transfer Protocol, See TFTP

troff — typeset or format documents, 1-893

troff postprocessor for PostScript printers —  
`dpost`, 1-201

troff utilities  
 check `nroff` and `troff` files — `checknr`, 1-79  
 eliminate `.so`'s from `nroff` input — `soelim`,  
 1-803  
 filters reverse line-feeds from two-column  
`nroff` text — `col`, 1-119  
 format tables — `tbl`, 1-856  
 postprocessor for PostScript printers — `dpost`,  
 1-201  
 remove `nroff`, `troff`, `tbl` and `eqn` con-  
 structs — `deroff`, 1-185

true — provide truth values, 1-895

truss — trace system calls and signals, 1-896

tset — set terminal characteristics, 1B-900

tsort — topological sort of items mentioned in  
 input, 1-905

tty — time to live value, `nischttl`, 1-573

tty, set characteristics — `stty`, 1B-827, 1B-900  
 set options — `stty`, 1-821

tty — get the name of the terminal, 1-906

`type` — shell built-in functions to set/get attributes and values for shell variables and functions, 1-907

`typeset` — shell built-in functions to set/get attributes and values for shell variables and functions, 1-907

`typeset` documents — `troff`, 1-893

## U

`u370` — get processor type truth value, 1-467

`u3b` — get processor type truth value, 1-467

`u3b15` — get processor type truth value, 1-467

`u3b2` — get processor type truth value, 1-467

`u3b5` — get processor type truth value, 1-467

`ucblinks` — adds `/dev` entries to give SunOS 4.1 compatible names to SunOS 5.x devices, 1B-909

`ul` — underline text, 1-910

`ulimit` — shell built-in functions to set/get limitations on the system resources available to the current shell and its descendents., 1-412

`umask` — shell built-in function to restrict read/write/execute permissions, 1-911

`unalias` — shell built-in functions to create your own pseudonym or shorthand for a command or series of commands, 1-35

`uname` — print name of current system, 1-912

`uncompress` — uncompress files, 1-122

`underline text` — `ul`, 1-910

`unexpand` — unexpand SPACE characters to TAB characters, 1-250

`unhash` — shell built-in functions to evaluate the internal hash table of the contents of directories, 1-330

`unifdef` — resolve and remove `ifdef`ed lines from C program source, 1-914

`uniq` — report repeated lines in a file, 1-915

`units` — converts quantities expressed in standard scales to other scales, 1-917

## UNIX

convert text file from DOS format to ISO format — `dos2unix`, 1-198

## UNIX-to-UNIX commands

## UNIX-to-UNIX commands, *continued*

`uucp` — `uucp`, 1C-921

`uulog` — `uucp`, 1C-921

`uname` — `uucp`, 1C-921

`unix2dos` — convert text file from ISO format to DOS format, 1-918

`unlimit` — shell built-in functions to set/get limitations on the system resources available to the current shell and its descendents., 1-412

`unpack` — expand compressed files, 1-611

`unset` — shell built-in functions to determine the characteristics for environmental variables of the current shell and its descendents, 1-773

`unsetenv` — shell built-in functions to determine the characteristics for environmental variables of the current shell and its descendents, 1-773

`until` — shell built-in functions to repetitively execute a set of actions while/until conditions are evaluated TRUE, 1-964

update groups of programs — `sysV-make`

update programs — `make`

`uptime` — show how long the system has been up, 1-919

## user

get the name of the user running the process — `logname`, 1-430

## user ID

change user IDs of files — `chown`, 1B-89

## user IDs

display a list of all valid user names — `dispuid`, 1-197

prompts for user ID — `ckuid`, 1-111

provides error message for user ID — `erruid`, 1-111

provides help message for user ID — `helpuid`, 1-111

validates user ID — `valuid`, 1-111

## users

display effective name — `whoami`, 1B-968

display group membership — `groups`, 1-327

display information about local and remote users — `finger`, 1-277

list user login information — `listusers`,

1-415

users, *continued*

- talk to another user — talk, 1-848
- who is logged in, and what are they doing — w, 1-954
- who is on the system — who, 1-965
- who's logged in on local machines — rwho, 1-728
- who's logged in on remote machines — rusers, 1-727
- write to another user — write, 1-970

users — display users on system, 1B-920

users, network

- Internet user name directory service — whois, 1-969

uucp — UNIX-to-UNIX copy, 1C-921

UUCP log — uulog, 1C-921

uucp status inquiry

- uustat, 1C-926

uudecode — decode binary file, 1C-924

uuencode — encode binary file, 1C-924

uuglist — print list of service grades available, 1C-925

uulog — UUCP log, 1C-921

uuname — UUCP list of names, 1C-921

uupick — public UNIX-to-UNIX copy, 1C-929

uustat — uucp status inquiry, 1C-926

uuto — public UNIX-to-UNIX copy, 1C-929

uux — system to system command execution, 1C-931

## V

vacation — automatic mail replies, 1-934

vax — get processor type truth value, 1-467

version control

- vc, 1-936

vgrind — formats program in nice style using troff, 1-939

vi — screen-oriented (visual) display editor based on ex, 1-942

vipw — edit password file, 1B-949

volcancel — cancel user's request for removable media that is not currently in drive, 1-950

volcheck — check for media in a drive. Default checks for floppy., 1-951

volmissing — notify user that volume requested is not in the CD-ROM or floppy drive, 1-952

Volume Management

- cancel user's request for removable media that is not currently in drive — volcancel, 1-950
- check for media in a drive. Default checks for floppy. — volcheck, 1-951
- missing volume notification — volmissing, 1-952

vsig — synchronize a co-process with the controlling FMLI application, 1F-953

## W

w — who is logged in, and what are they doing, 1-954

wait — shell built-in function to wait for other jobs or processes, 1-956

wc — display a count of lines, words and characters in a file, 1-958

what — extract SCCS version information from a file, 1-959

whatis — describe command, 1-960

whence — shell built-in functions to set/get attributes and values for shell variables and functions, 1-907

whereis — locate the binary, source and manual page files for a command, 1B-961

which — locate a command; display its pathname or alias, 1-963

while — shell built-in functions to repetitively execute a set of actions while/until conditions are evaluated TRUE, 1-964

who — who is on the system, 1-965

who is logged in — w, 1-954

whoami — display effective user name, 1B-968

whois — Internet user name directory service, 1-969

write — write to another user, 1-970

---

## X

`xargs` — construct argument list(s) and execute command, 1-972  
`xgettext` — extract `gettext` call strings, 1-975  
`xstr` — extract strings from C code, 1-977

## Y

`yacc` — yet another compiler-compiler, 1-979  
YACC  
    create a tags file for use with `ex` and `vi` —  
        `ctags`, 1-170  
`yes/no` answer  
    prompts for yes/no answer — `ckyorn`, 1-113  
    provides error message for yes/no answer —  
        `erryorn`, 1-113  
    provides help message for yes/no answer —  
        `helpyorn`, 1-113  
    validates yes/no answer — `valyorn`, 1-113  
`yet another compiler-compiler` — `yacc`, 1-979  
`ypcat` — print values in a NIS database, 1-981  
`ypmatch` — print the value of one or more keys  
    from a NIS map, 1-982  
`yppasswd` — change your network password in the  
    NIS database, 1-983  
`ypwhich` — return name of NIS server or map mas-  
    ter, 1-984

## Z

`zcat` — displays uncompressed files but leaves  
    compressed files intact, 1-122